Canadian Journal of Philosophy (2020), 50: 3, 339–349
doi:10.1017/can.2019.36

ARTICLE

Higher-order knowledge and sensitivity

Jens Christian Bjerring and Lars Bo Gundersen*

Department of Philosophy, Aarhus University, Aarhus C, Denmark
*Corresponding author. Email: fillg@cas.au.dk

Abstract
It has recently been argued that a sensitivity theory of knowledge cannot account for intuitively appealing instances of higher-order knowledge. In this paper, we argue that it *can* once careful attention is paid to the methods or processes by which we typically form higher-order beliefs. We base our argument on what we take to be a well-motivated and commonsensical view on how higher-order knowledge is typically acquired, and we show how higher-order knowledge is possible in a sensitivity theory once this view is adopted.

Keywords: Higher-order knowledge; sensitivity theory; transmission principle for knowledge; belief-forming methods

1 Introduction

Knowledge is often accompanied by higher-order knowledge. Suppose you happen to know that it is raining just now. If you are in the right reflective mood, you might ponder the question: what exactly do I know? And it seems that not only do you know something about the weather, you also know something about yourself: namely, that one of your beliefs—say, your belief that it is raining—is true. Of course, the question whether this belief of yours is true may never have crossed you mind, and if it did not, your belief would not be accompanied by any higher-order knowledge. Still, given that you are in the right sort of reflective mood—the sort of mood in which such questions *do* cross your mind—if you know that it is raining you also know that your belief that it rains is true. This is as close to a Moorean fact as anything.

Yet DeRose, Kripke, Sosa, Vogel, and Williamson, among others, have argued that a sensitivity theory cannot account for this kind of higher-order knowledge.[1] According to a sensitivity theory, an agent S can know that p only if the following condition is satisfied, where M is the method or process by which one actually comes to believe p:

> (**Sensitivity**): Had p been false, and had S formed her belief whether p via M, S would not have believed p via M.[2]

Consider now some true higher-order belief of yours: say, the higher-order belief that you are not believing falsely that it is raining. Under careful reflection, we may assume, it should be unproblematic to promote such true higher-order beliefs to knowledge. But, goes the argument against the sensitivity theory, such higher-order beliefs are never sensitive.

[1]See, amongst others, DeRose (1995, 22); Kripke (2011/1986, 203); Sosa (1999, 145); Vogel (2000, 610); and Williamson (2000, Ch. 7).

[2]Obviously, there are many details yet to be filled in by the sensitivity theorist, and numerous ways of doing so. The present discussion is neutral on these details, although the versions of the sensitivity theory in Black and Murphy (2012) and Roush (2012) deserve special mention because they both embrace closure.

© The Author(s) 2020. Published by Canadian Journal of Philosophy.

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

For suppose that your higher-order belief *had* been false. In that case, you would falsely believe that it is raining. Yet, you would still *believe* that it is raining, and also that you do not believe *falsely* that it is. So your higher-order belief is not sensitive and, hence, according to (Sensitivity), it does not qualify as knowledge. But if the sensitivity theory cannot account for the Moorean fact that first-order knowledge is often accompanied by higher-order knowledge, it is no wonder that Kripke and Williamson believe that this *argument from higher-order knowledge* gives a serious blow to the sensitivity theory, and that others, such as Sosa and Vogel, go even further and take the argument to be decisive for the theory.[3]

In this paper, we argue that once we pay close attention to the methods that are involved in paradigmatic cases of higher-order knowledge acquisition, higher-order beliefs can be sensitive and hence qualify as knowledge according to the sensitivity theory. In section 2, we present the argument from higher-order knowledge in more detail. While much of the discussion concerning this argument has centered rather narrowly around different logical formalizations of higher-order beliefs, we will approach matters from a different and, we hope, more principled perspective. More specifically, as we argue in section 3, if we are granted what we take to be a well-motivated and commonsensical view on how higher-order knowledge is typically acquired, then higher-order knowledge is possible to attain in a sensitivity theory. In section 4, we return to the argument from higher-order knowledge and argue that a sensitivity theorist—contrary to what Kripke, Sosa, Vogel, and Williamson claim—has nothing to fear from it.

2 The argument from higher-order knowledge

The argument from higher-order knowledge trades on the assumption that (F_1) is the correct formalization of higher-order propositions such as "I am not believing falsely that it is raining" ("B" stands for the belief operator):

(F_1) $\neg(Bp \wedge \neg p)$

Once higher-order beliefs are understood in terms of (F_1), we can make more precise the argument from higher-order knowledge. Following Vogel (2000), consider first some true higher-order belief that an agent S has:

(1) In the actual world, S believes that $\neg(Bp \wedge \neg p)$.

To determine whether the belief that $\neg(Bp \wedge \neg p)$ is sensitive, we need to evaluate the counterfactual "Had $\neg(Bp \wedge \neg p)$ been false, then S would would not have believed that $\neg(Bp \wedge \neg p)$." To do so, we will appeal to the following standard semantics for counterfactuals ('$\Box\!\!\rightarrow$' stands for counterfactual entailment):

(**ConFac**) $p \Box\!\!\rightarrow q$ is true in a possible world w iff either p is impossible, or there is a world in which p and q are true that is closer to w than any world in which p is true but q is false.[4]

[3]In Sosa (1999, 152), for instance, it is declared that "this sort of counterexample […] strikes me as conclusive." For further references, see footnote 1.

[4]See Lewis (1973). We are not claiming that the Lewis/Stalnaker semantics for counterfactuals is unproblematic. Notably, there are severe problems with so-called *counterpossibles*: counterfactual conditionals with impossible antecedents. Since counterpossibles are vacuously true in the standard possible worlds' semantics, it follows that every belief in a necessary proposition is trivially sensitive and, therefore, constitutes knowledge according to a sensitivity account. But intuitively, this is wrong. While the problem of knowledge of necessary truths is a serious problem for the sensitivity theory—and likewise for the standard safety theory for knowledge—this is not the place to discuss it. For more on the semantics of counterpossibles, see for instance Bjerring (2014) and Brogaard and Salerno (2013).

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

Let, then, w be the closest world(s) to the actual one in which $\neg(Bp \wedge \neg p)$ is false. Then, by simple logic, we have:

(2) In w, $\neg\neg(Bp \wedge \neg p)$.
(3) In w, $(Bp \wedge \neg p)$.
(4) In w, Bp.

Let now (Doxastic Optimism) be the trivial sounding principle that whenever S believes that p, she also believes that she truly believes that p:[5]

(**Doxastic Optimism**) $Bp \rightarrow B(\neg(Bp \wedge \neg p))$.

Given (4) and (Doxastic Optimism), we immediately get (5):

(5) In w, $B(\neg(Bp \wedge \neg p))$.

So in the closest world in which (F_1) is false, S still believes that it is true. That is, the counterfactual $(Bp \wedge \neg p) \boxbslash\!\!\rightarrow \neg B(Bp \wedge \neg p)$ is false. Accordingly, when the relevant higher-order propositions are formalized as (F_1), beliefs in such propositions are *not* sensitive. So the sensitivity theory seems to fail: it cannot accommodate our intuitions about paradigmatic instances of higher-order knowledge acquisition.

The argument from higher-order knowledge relies centrally on the assumption that higher-order beliefs should be understood in terms of (F_1): the argument only gets off the ground in step (2) once we negate the target higher-order proposition $\neg(Bp \wedge \neg p)$. If, instead, we formalize the relevant higher-order belief as (F_2), we can show—by an argument analogous to the one above—that higher-order beliefs *are* sensitive:[6]

(F_2) $(Bp \wedge p)$

Suppose first that S truly believes $(Bp \wedge p)$:

(1) In the actual world, S believes that $(Bp \wedge p)$.

To figure out whether the belief that $(Bp \wedge p)$ is sensitive, we need to evaluate the counterfactual $\neg(Bp \wedge p) \boxbslash\!\!\rightarrow \neg B(Bp \wedge p)$. So let w be the closest world(s) to the actual one in which $(Bp \wedge p)$ is false:

(2) In w, $\neg(Bp \wedge p)$.

By simple logic we get (3) from (2):

(3) In w, $\neg Bp \vee \neg p$.

While we know by assumption that Bp is true in the actual world, we can now ask whether Bp or $\neg Bp$ is true in w. Suppose Bp is true in w. By (3), then $\neg p$ is also true in w. So both Bp and $\neg p$ are true in w. Yet, when we keep the agent's *actual* belief-forming method fixed in w—as (Sensitivity) tells us to do—it is not very plausible that both Bp and $\neg p$ should be true in w. For when p is a mundane proposition such as "It is raining" or "There is a chair in front of me," the agent usually forms her

[5]While there might be room for disagreeing with this principle, we shall simply assume it for current purposes. For discussion, see Salerno (2010).

[6]Here we follow the outline in Becker (2006).

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

belief about such propositions in the actual world by employing standard, reliable belief-forming methods such as perception or testimony. But then the world w in which both Bp and $\neg p$ are true would have to be a world in which such standard, reliable belief-forming methods go astray. However, while such a world obviously exists, it is never—contrary to our assumption—among the *closest* worlds to the actual one, and, at least, it is far more remote from actuality than a world in which $(\neg Bp \wedge \neg p)$ or $(\neg Bp \wedge p)$ is true. Hence, in the closest world w, we have:

(4) In w, $\neg Bp$.

We can now contrapose the highly plausible distribution principle $B(Bp \wedge p) \rightarrow Bp$ to obtain (5) and, derivatively, infer (6) from (4) and (5):[7]

(5) In w, $\neg Bp \rightarrow \neg B(Bp \wedge p)$.
(6) In w, $\neg B(Bp \wedge p)$.

It thus follows that in the closest world(s) where (F_2) is false, S does not believe it. That is, the counterfactual $\neg(Bp \wedge p) \mathrel{\square\!\!\rightarrow} \neg B(Bp \wedge p)$ is true. Accordingly, when the relevant higher-order propositions are formalized as (F_2), beliefs in such propositions *are* sensitive.

To judge whether a sensitivity theory can account for higher-order knowledge or not, the crucial question is then whether higher-order beliefs are most plausibly formalized as (F_1) or as (F_2). If we formalize higher-order beliefs as (F_1), most philosophers agree that the sensitivity theory cannot account for plausible instances of higher-order knowlege.[8] If, on the other hand, we formalize higher-order beliefs as (F_2), there seems to be consensus that the sensitivity theory can account for plausible instances of higher-order knowledge.[9] So we seem to be stuck with what Melchior has recently labeled "the heterogeneity problem" for the sensitivity theory:

> Sensitivity accounts of knowledge deliver a very heterogeneous picture of higher-level beliefs about the truth or falsity of our own beliefs. Some beliefs are insensitive, but other beliefs, in closely related but stronger propositions, are sensitive. Thus, sensitivity accounts of knowledge cannot adequately capture higher-level knowledge. (2015, 483)[10]

Essentially, Melchior points out, if we are looking for an explanation of how higher-order beliefs are sensitive, we should expect a uniform account that treats similar higher-order beliefs alike: if the sensitivity theory predicts that the class of higher-order beliefs conforming to (F_2) are sensitive, it should also predict that the logically similar class of higher-order beliefs conforming to (F_1) are sensitive. Yet, since it does not, the sensitivity theory fails.[11]

---

[7]This distribution principle, it seems, only fails in the very odd case where agents form higher-order beliefs about having beliefs they do not in fact have.

[8]Most, but not all, for Wallbridge (2016a) has recently argued that once we pay appropriate attention to the relevant belief-forming methods, then higher-order beliefs formalized as (F_1) do in fact count as sensitive.

[9]Of course, we might quibble about the details of the argument above that shows that higher-order beliefs are sensitive when they are formalized as (F_2). Vogel (2012), for instance, worries that the semantic behavior of counterfactuals with disjunctive antecedents is "extremely wayward" (123). It is well-known that a counterfactual of the form '$A \vee B \mathrel{\square\!\!\rightarrow} C$' sometimes simplifies to '$A \mathrel{\square\!\!\rightarrow} C$' and '$B \mathrel{\square\!\!\rightarrow} C$' but sometimes not. So, according to Vogel, we should be hesitant to rely on arguments that appeal to such counterfactuals. While we can agree with Vogel that the semantics for counterfactuals with disjunctive antecedents is a complicated affair, it makes no difference to the argument for the sensitivity of (F_2)-formalized higher-order beliefs. When it comes to ordinary propositions such as "It is raining" and "I believe that it is raining," the counterfactual '$(\neg Bp \vee p) \mathrel{\square\!\!\rightarrow} \neg B(Bp \wedge p)$' comes out true in the actual world irrespective of whether the counterfactual simplifies or not. For further discussion of these issues, see also Salerno (2010).

[10]Since propositions of the form $Bp \wedge p$ entail propositions of the form $\neg(Bp \wedge \neg p)$, but not vice versa, the former are logically stronger than the latter.

[11]For a reply to Melchior's problem, see Wallbridge (2016b), and for a reply to Wallbridge's reply, see Melchior (2017).

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

In light of the impressive number of philosophers who have claimed that the sensitivity theory fails to account for higher-order knowledge, the received view in the literature is arguably that it so fails. Yet, there are voices to the contrary. Becker (2006) and Salerno (2010) have argued that the sensitivity theory has no problem accounting for higher-order knowledge because (F_2) ought to serve as the correct formalization of higher-order beliefs. Central to their argument—a part of which we will use in section 4—is the observation that whereas $(Bp \wedge p)$ entails that the agent in question has the first-order belief Bp, $\neg(Bp \wedge \neg p)$ does not. But intuitively, if an agent has any higher-order beliefs about which propositions she truly—or not falsely—believes, then it seems as if the agent must also have certain first-order beliefs about those propositions. So, since (F_1) does not entail that the agent has any first-order beliefs, but (F_2) does, Becker and Salerno conclude that (F_2) constitutes a better formalization of higher-order beliefs than (F_1).

We agree with Becker and Salerno that (F_2) ought to serve as the correct formalization of higher-order beliefs. But we want to argue for this view in a new and, we hope, more principled way. Our argument is grounded in what we take to be an intuitively appealing and commonsensical picture of the processes or methods by which people typically come to acquire higher-order knowledge. This picture, as we shall see, makes the relevant method for forming higher-order beliefs a combined one that centrally involves *inferential reasoning*. While it is well-known that inferential reasoning *generally* does not preserve knowledge in a sensitivity theory, we will argue that it does when it comes to the type of inferential reasoning that is involved in paradigmatic cases of higher-order belief acquisition. This, in turn, will show that a sensitivity theory can account satisfactorily for higher-order knowledge and, finally, why (F_2) should serve as the correct formalization of higher-order beliefs.

Moreover, as we argue in section 4, there is no corresponding intuitively appealing and commonsensical specification of the methods that would lead to higher-order knowledge of propositions formalized as (F_1). For when we start reflecting on the methods that would be relevant for forming higher-order beliefs in accordance with (F_1), we end up with a strange method that has some very counterintuitive results. But if so, then it is not surprising that the sensitivity theory gives a heterogeneous picture of higher-order beliefs. For we should not fault a counterintuitive method for yielding insensitive higher-order beliefs—notwithstanding questions of logical formalization.

3 The sensitivity of higher-order knowledge

Typically, the processes or methods involved in acquiring higher-order knowledge are very different from those involved in acquiring first-order knowledge. In particular, first-order knowledge is generally epistemically *prior* to higher-order knowledge. Consider a case in which you acquire higher-order knowledge that your belief that it is raining is true. In the normal run of things, we claim, the following constitutes an intuitively plausible and commensensical story about the relevant higher-order belief formation. First, you acquire knowledge that it is raining through some combination of the usual visual and audio methods. Second, if you are in the right mood, you start reflecting on the fact that you now believe that it is raining, and by going through the following short chain of reflective reasoning, you come to know that this belief of yours is true:

(**Reflection**)

p_1: It is raining, and
p_2: I believe it is raining, so
c: My belief that it is raining is true.

By engaging in this kind of reflective reasoning—most likely an implicit or even unconscious type of reasoning—you are in a position to *transmit* your knowledge of p_1 and p_2 through to c, and thereby in a position to acquire the piece of higher-order knowledge in c. Clearly, you do *not* know c merely

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

in virtue of employing the method that gives you knowledge of p_1. More is needed. You need typically, first, to do a bit of introspective or reflective reasoning to come to know p_2, and then, crucially, a bit of inferential reasoning involving conjunction introduction to transmit your knowledge of p_1 and p_2 through to c. In other words, the method relevant for acquiring higher-order knowledge of c is a *combined* one. It consists partly of the method relevant for gaining knowledge of p_1, which typically involves some combination of standard audio and visual methods such as perception or testimony. It consists partly of the method relevant for gaining knowledge of p_2, which typically involves some combination of more reflective or introspective kinds of methods. And it consists partly of the method relevant for transmitting knowledge of p_1 and p_2 through to c, which typically involves inferential reasoning or, more specifically, basic logical reasoning involving conjunction introduction.

It is well-known that it is not easy to give a precise story about how to individuate belief-forming methods. Yet, to motivate (Reflection), we do not need to make any substantial assumptions about such individuations.[12] We need only appeal to the broad and pretheoretical criteria above that show that the methods relevant for forming first-order beliefs are significantly different from those relevant for forming higher-order beliefs. As long as we have a good enough grip on these different types of methods, it should be clear how basic logical reasoning can bring together the beliefs that result from applying these different methods. Likewise, we need not be overly committal about whether knowledge of p_2 is always gained through a process of introspection. While we might agree that knowledge of p_2 is *typically* gained through some sort of introspection, we can also grant that some other method, which is epistemically just as good as introspection, can help us explain how knowledge of p_2 is acquired. Going forward, however, we will work with a simple picture according to which knowledge of p_1 is gained through perception, knowledge of p_2 is gained through introspection, and knowledge of c is gained through inferential reasoning. But bear in mind that nothing substantial hinges on these simplifications.

Granting (Reflection), it now remains to be shown that the combined method of perception, introspection, and inferential reasoning in the end produces sensitive beliefs. Even among the staunchest critics of the sensitivity theory, it is widely agreed that methods such as perception and introspection do generate knowledge according to the sensitivity theory.[13] However, it is also widely agreed that the sensitivity theorist cannot rely on the method of inferential reasoning as a source of knowledge.[14] According to the sensitivity theory, you may know that you have hands without knowing what follows inferentially from this piece of knowledge: namely, that you are not a handless brain in a vat. So many have concluded that a sensitivity theorist cannot rely on inferential reasoning as a means for transmitting knowledge through a (known) entailment.

Yet, while a sensitivity theorist cannot generally rely on inferential reasoning as a means for extending our knowledge, there are several special instances where it is possible. Nozick himself (1981, 230–39) accepts that knowledge can be transmitted through known extensional generalization, through known equivalence, and through known conjunction introduction. But what is perhaps less well-known is that Nozick in his *Philosophical Explanations* offers a *general* transmission principle which tells us for *any* valid inference, whether knowledge transmits through that inference, and which yields as special cases transmission of knowledge over (known) conjunction introduction and equivalence. We can state this principle as follows ("K" stands for

[12]For a discussion of these issues in relation to the sensitivity theory, see Becker (2012).

[13]See, for instance, Vogel (2012).

[14]But see Baumann (2012), Kripke (2011/1986, Ch. 7), and Zalabardo (2012).

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

the knowledge operator and "→" for a notion of logical entailment that has the force of a strict implication):[15]

(**N-Transmission**)

If
(i) Kp
(ii) $K(p \rightarrow c)$
(iii) $\neg c \,\Box\!\!\rightarrow \neg Bp$
then
Kc

That is, for any valid inference from some set of premises to a conclusion, if an agent knows the premises and knows that the premises entail the conclusion, and if the agent would not have believed the premises had the conclusion been false, then she can transmit that knowledge through the entailment and gain knowledge of the conclusion.

To see how (N-Transmission) can help establish the more restricted transmission principles that Nozick accepts, consider transmission of knowledge through known equivalence: if $K(p)$ and $K(p \leftrightarrow q)$, then $K(q)$. To see that knowledge gets transmitted from knowledge of the two premises to the conclusion, it suffices to show that the following instance of condition (iii) in (N-Transmission) is true in the actual world: $\neg p \,\Box\!\!\rightarrow \neg q$. And it is not hard to see that it is. For when p and q are logically equivalent, then the closest world (to the actual one) in which $\neg p$ is true is also the closest world in which $\neg q$ is true. So, given condition (i), the closest world in which $\neg q$ is true is also a world in which $\neg Bp$ is true. So we know that $\neg q \,\Box\!\!\rightarrow \neg Bp$ is true in the actual world, in which case condition (iii) is satisfied. So knowledge of p can be transmitted through the equivalence and issue in knowledge of q. A similar simple piece of reasoning shows how (N-Transmission) can guarantee that knowledge gets transmitted through known conjunction introduction.

While (N-Transmission) enables us to transmit knowledge through some known inferential patterns, note also that it respects the classical and—among friends of the sensitivity theory, anyway—intuitively plausible instances of transmission failure. The principle tells us that knowledge transmits through (known) conjunction introduction, but it also tells us that it does not transmit through every instance of (known) conjunction elimination. According to Nozick, for instance, I can know the conjunction that I am currently drinking a cup of coffee and that I am not a brain in a vat, but still not be in a position to come to know the latter conjunct by inferring it via conjunction elimination. For in the closest world to the actual one in which I am a brain in a vat, I would still believe that I am not. Likewise, according to Nozick, I may know that I have a hand and know that if I have a hand, then I am not a brain in a vat, and yet be unable to transmit this knowledge through to the conclusion that I am not a brain in a vat. For the following instance of condition (iii) in (N-Transmission) is false in the actual world: "Had I been a brain in a vat, I would not have believed that I have a hand and that my having a hand implies that I am not a brain in a vat." That is, in the closest world to the actual one in which I am a brain in a vat, I would still believe that I have a hand. Of course, there is an intensive and lively debate about whether such consequences of Nozick's view are acceptable.[16] But engaging in that

[15]See Nozick (1981, 231). As formulated, (N-Transmission) admittedly looks more like a closure than a transmission principle. But as Nozick's discussion of the principle makes clear (230–39), it is really a transmission principle. A more careful formulation of (N-Transmission) should be explicit about the fact that one comes to know *c* in virtue of *inferring* it from *p*. Also, a more careful formulation should take into account the intended modal status of the consequence—that one is *in a position* to come to know *c*—and the kinds of complications discussed in Hawthorne (2004). For present purposes, however, we can safely bracket these finer details.

[16]For a general discussion of the various moving pieces in the debate surrounding Nozick's conception of transmission, see (Baumann 2012).

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

debate is not our topic here. Rather, our aim is to show how we can use (N-Transmission) to secure that the combined belief-forming method used in (Reflection) produces sensitive beliefs—and hence how higher-order knowledge is attainable in a sensitivity theory.

Given the setup, it is not hard to see how (N-Transmission) can do the required work. First, since we grant that the sensitivity theorist can appeal to the methods of perception and introspection to account for knowledge of, respectively, p_1 and p_2, we know that condition (i) in (N-Transmission) is satisfied. Second, since the agent in (Reflection) knows that an application of conjunction introduction will lead from the premises p_1 and p_2 to the conclusion c, condition (ii) in (N-Transmission) is satisfied as well. So we merely need to ensure that condition (iii) in (N-Transmission) is satisfied. And it is easily seen that it is. Since we utilize the fact that the agent infers c by applying conjunction introduction on p_1 and p_2, it follows that c is logically equivalent to the conjunction p of the premises p_1 and p_2. But then it is also immediately clear, as we saw above, that $(\neg p \mathrel{\Box\!\!\rightarrow} \neg c)$ is true in the actual world. So all the conditions in (N-Transmission) are satisfied, and we can thus transmit knowledge of p_1 and p_2 through to c. Hence (N-Transmission) vindicates our explanation of how we can acquire higher-order knowledge in the manner suggested by (Reflection).

Derivatively, once we adopt our view on higher-order knowledge acquisition, it becomes clear why (F_2) should serve as the correct formalization of the relevant type of higher-order propositions:

(F_2) $(Bp \wedge p)$

As (Reflection) makes clear, an agent proceeds from knowledge that it is raining *and* from knowledge that she believes that it is raining through an inferential step involving conjunction introduction to knowledge of the higher-proposition that she truly believes that it is raining. So if we are granted our—we submit—intuitively plausible view on how higher-order knowledge is typically attained, we have independent reason for holding that higher-order knowledge should end up with a conjunctive content that conforms to (F_2).

4 Formalizing higher-order beliefs

Suppose you accept our general picture of higher-order knowledge acquisition and grant us that *some* sort of inferential reasoning is involved in obtaining higher-order knowledge. Even then you might object that we have only half-rescued the sensitivity theory. As we saw in section 2, the argument from higher-order knowledge shows that there is a class of higher-order propositions—namely those formalized according to (F_1)—that are never sensitive but are, nevertheless, supposedly easy to come to know. So what is wrong with the argument from higher-order knowledge?

For the argument from higher-order knowledge to be a serious challenge to the sensitivity theory, we need a specification of the belief-forming methods by which agents come to have beliefs in propositions formalized as (F_1). As it is clearly brought out by Nozick's famous grandmother example, a sensitivity theorist should only care about cases where the relevant belief-forming method is specified.[17] Dialectically, it is hence not enough *merely* to point out that there is a class of propositions—namely those formalized as (F_1)—that do not square nicely with the sensitivity theory. Rather, we need a story about the underlying processes or methods that should lead to beliefs in (F_1). In light of the challenges involved in individuating belief-forming methods, we are, of course, not demanding a maximally specific characterization of the relevant belief-forming method. But we can require a characterization that is intuitively plausible, sufficiently general, and independently motivated. Surprisingly, however, not much attention has been given to this issue in the literature.

[17]See Nozick (1981, 179). In Nozick's example, a grandmother has a perceptually based belief that her grandson is well. While this belief is not sensitive *per se*, it nevertheless does qualify as knowledge on a sensitivity picture. F, for the belief in question remains sensitive *relative* to the particular method by which it was initially formed.

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

We have presented a picture of higher-order knowledge acquisition that we trust is both intuitively plausible, sufficiently general, and independently motivated. So let us assume that we accept the general idea behind (Reflection): we come to have higher-order knowledge of c through a piece of inferential reasoning from first-order knowledge of the premises p_1 and p_2. If higher-order propositions conform to (F_2), we may agree—as we did above—that the relevant inferential reasoning involves a step of conjunction introduction. But if higher-order propositions conform to (F_1) instead, the inferential reasoning will look different:

(F_1) $\neg(Bp \wedge \neg p)$

since (F_1) is equivalent to the disjunction $(\neg Bp \vee p)$, in which case the inference from first-order knowledge to higher-order knowledge must involve a step of *disjunction introduction* rather than conjunction introduction. Yet, as expected, since (N-Transmission) does not allow knowledge to be transmitted over an instance of disjunction introduction involving $(\neg Bp \vee p)$, we do not preserve sensitivity when inferring (F_1) from the premises p_1 and p_2.

To evaluate the argument from higher-order knowledge, the question is then how intuitively plausible this "disjunctive" belief-forming method is. We want to suggest that it is quite *im*plausible. Note first the following: when an agent *knows* a proposition p conforming to (F_1), she does *not* need to be aware—let alone know—that she *believes* that p.[18] For (F_1) is equivalent to $(\neg Bp \vee p)$ and, obviously, as we saw earlier, $(\neg Bp \vee p)$ can be true without Bp being true. But then, essentially, the disjunctive belief-forming method can proceed as follows: an agent starts out with knowledge of some proposition p and proceeds—perhaps completely agnostic about whether she believes that p—directly via disjunction introduction to $(\neg Bp \vee p)$ and then via DeMorgan equivalence to $\neg(Bp \wedge \neg p)$. Put differently, if (F_1) is the right formalization of higher-order propositions such as "I truly believe that p" or "I do not falsely believe that p," then the process of forming higher-order beliefs need not involve any knowledge of a premise like p_2: whether or not you recognize that you have the first-order belief that p—whether by introspection or by some other method—is irrelevant for you when it comes to forming higher-order beliefs about propositions that involve those very first-order beliefs.[19]

Yet, when we formalize the relevant higher-order propositions as (F_1), it is easy to see that the disjunctive belief-forming method quickly gives rise to some very counterintuitive ascriptions of higher-order knowledge. To make this vivid, consider a case in which an agent S_1 forms a higher-order belief that involves some other agent S_2. Suppose S_1 knows that p but is completely agonistic as to whether S_2 believes that p. If S_1 forms beliefs in accordance with the disjunctive method, she can now proceed from knowledge of p via disjunction introduction to knowledge of $(\neg B_{S_2}p \vee p)$, and then via DeMorgan equivalence to knowledge of $\neg(B_{S_2}p \wedge \neg p)$. But ask now: is there any reasonably intuitive sense in which S_1 comes to *know* that S_2 *truly believes* that p—or that S_2 does not falsely believe that p—by following this disjunctive belief-forming method? Clearly, the answer is "no." To see this, consider the following application of the method. I happen to know that

[18]Salerno (2010), as mentioned, unfolds this point in great detail and uses it to argue that (F_2) rather than (F_1) constitutes the right formalization of higher-order beliefs.

[19]Note here that the claim is only that a process of higher-order belief formation *need* not involve any recognition of the fact that the agent believes p. It is of course compatible with the view that the agent in fact recognizes and comes to know that she believes p. Yet, if she does, it is also plausible that she would form her belief in the higher-order proposition that she truly believes that p by applying conjunction introduction. For if the agent realizes that p is true and that she believes that p, then it is also plausible that she would use conjunction introduction to infer that her belief that p is true rather than, strangely, ignore this fact and instead proceed via disjunction introduction from knowledge of p alone. More generally, if the agent has reflective knowledge about the fact that she believes that p, then it seems plausible that her preferred method of proceeding to knowledge of the proposition that she truly believes that p would utilize this fact. But then, generally, it is also plausible that she would infer the relevant higher-order proposition using conjunction introduction.

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

Darjeeling is my wife's favorite tea, but can I infer from that piece of knowledge—and that piece of knowledge alone—that some other agent, say Donald Trump, truly believes that Darjeeling is my wife's favorite tea? By disjunction introduction, I surely *can* infer *either* that Darjeeling is my wife's favorite tea *or* that Donald Trump does not believe that it is $[(\neg B_{S_2} p \vee p)]$. So I *can* also infer that it is *not* the case that Darjeeling is *not* my wife's favorite tea *and* that Donald Trump believes that it is $[\neg(B_{S_2} p \wedge \neg p)]$. But have I thereby come to *know* that Donald Trump *truly believes*—or that he does not falsely believe—that my wife's favorite tea is Darjeeling? Of course not. Regardless of what else Trump has beliefs about, he certainly has no beliefs about my wife.

So if (F_1) is the correct formalization of higher-order proposition such as "I truly believe that p," and if the disjunctive method is the relevant one for arriving at knowledge of these higher-order propositions, then we end up with obviously false—or at least highly counterintuitive—knowledge claims about which propositions agents can truly be said to believe. Since our "conjunctive" method requires an agent to know p_2—that is, that she believes the relevant first-order proposition—our belief-forming method faces no such problems. The disjunctive method yields counterintuitive results because it fails to take into consideration the relevant first-order belief, and the reason it fails to take into consideration this first-order belief is due to the fact that (F_1) is to begin with a poor representation of higher-order beliefs.

Insofar as the disjunctive belief-forming method guides the formation of beliefs in higher-order propositions that conform to (F_1), the sensitivity theorist thus has little reason to worry about the argument from higher-order knowledge, and in turn little reason to worry about the fact that beliefs in propositions that conform to (F_1) are not sensitive. By contrast, on the intuitive picture we have presented, you form a higher-order belief about what you truly believe by conjoining facts about your knowledge and facts about your first-order beliefs. It matters little whether or not you come to know the facts about your first-order beliefs through introspection or some other reliable method. Rather, what matters is that you *utilize* the fact that you have certain first-order beliefs to form higher-order beliefs in propositions that involve those first-order beliefs. The picture presented in (Reflection) respects this intuitive datum and it shows, we have argued, why the relevant kind of higher-order propositions should be interpreted in terms of (F_2) rather than (F_1). So while we can grant Melchior (2015) that a sensitivity theory yields a heterogeneous picture of higher-order beliefs, we disagree with him when he claims that it is a problem.[20] In our opinion, the sensitivity theory tells us correctly why beliefs in (F_2) are sensitive and why beliefs in (F_1) are not.

In conclusion, while sensitivity theorists have other hard battles to fight—involving, for instance, inductive knowledge, "epistemic masking and mimicking," and beliefs in necessary propositions—they do not have to worry about higher-order knowledge. When properly understood, beliefs in higher-order propositions are sensitive and may qualify as knowledge.[21]

Acknowledgements. Thanks to two anonymous referees for helpful comments on this article.

Jens Christian Bjerring is an associate professor in philosophy at Aarhus University, Denmark. He was awarded his PhD in philosophy from the Australian National University in 2010, and has special interests in philosophical logic, formal epistemology, and philosophy of artificial intelligence.

[20]Since propositions of the form (F_2) logically entail propositions of the form (F_1), one might worry that a remainder of Melchior's heterogeneous problem exists. But, as far as we can tell, there is no problem here. Since knowledge is not closed under (known) logical entailment, according to (N-Transmission), agents can know propositions of the form $(Bp \wedge p)$ without thereby knowing propositions of the form $\neg(Bp \wedge \neg p)$.

[21]For discussions of the problem of inductive knowledge, see Sosa (1999) and Vogel (2000). For challenges to the conditional account of knowledge arising from masking and mimicking, see Gundersen (2010; 2012). For problems concerning beliefs in necessary propositions, see Bjerring and Skipper (2019) and Blome-Tillmann (2017).

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press

Lars Bo Gundersen is an associate professor in philosophy at Aarhus University, Denmark. He was awarded his PhD in philosophy from the University of St. Andrews in 2001, and has special interests in philosophical logic, epistemology, and philosophy of science.

References

Baumann, P. (2012). Nozick's Defense of Closure. In *The Sensitivity Principle in Epistemology*, edited by K. Becker and T. Black, 11–27. New York: Cambridge University Press.

Becker, K. (2006). Is Counterfactual Reliabilism Compatible with Higher-Order Knowledge? *Dialectica* 60 (1): 79–84.

Becker, K. (2012). Methods and How to Individuate Them. In *The Sensitivity Principle in Epistemology*, edited by K. Becker and T. Black, 79–95. New York: Cambridge University Press.

Bjerring, J. C. (2014). On Counterpossibles. *Philosophical Studies* 168 (2): 327–53.

Bjerring, J. C., and M. Skipper. (2019). A Dynamic Solution to the Problem of Logical Omniscience. *The Journal of Philosophical Logic* 48 (3): 501–21.

Black, T. and P. Murphy. (2012). Sensitivity Meets Explanation: An Improved Counterfactual Condition on Knowledge. In *The Sensitivity Principle in Epistemology*, edited by K. Becker and T. Black, 28–42. New York: Cambridge University Press.

Blome-Tillmann, M. (2017). Sensitivity Actually. *Philosophy and Phenomenological Research* 94 (3): 606–25.

Brogaard, B., and J. Salerno. (2013). Remarks on Counterpossibles. *Synthese* 190: 639–60.

DeRose, K. (1995). Solving the Skeptical Problem. *Philosophical Review* 104 (1): 1–52.

Gundersen, L. (2010). Tracking, Epistemic Dispositions and the Conditional Analysis. *Erkenntnis* 72(3): 353–64.

Gundersen, L. (2012). Knowledge, Cognitive Dispositions and Conditionals. In *The Sensitivity Principle in Epistemology*, edited by K. Becker and T. Black, 66–81. New York: Cambridge University Press.

Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.

Kripke, S. (2011/1986). *Philosophical Troubles*. New York: Oxford University Press.

Lewis, D. (1973). Counterfactuals and Comparative Possibility. *Journal of Philosophical Logic* 2: 418–46.

Melchior, G. (2015). The Heterogeneity Problem for Sensitivity Accounts. *Episteme* 12 (4): 479–96.

Melchior, G. (2017). Sensitivity Has Multiple Heterogeneity Problems: A Reply to Wallbridge. *Philosophia* 45 (4): 1741–47.

Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Roush, S. (2012). Sensitivity and Closure. In *The Sensitivity Principle in Epistemology*, edited by K. Becker and T. Black, 242–68. New York: Cambridge University Press.

Salerno, J. (2010). Truth Tracking and the Problem of Reflective Knowledge. In *Topics in Contemporary Philosophy: Knowledge and Skepticism* vol. 5, edited by J. Keim-Campbell, M. O'Rourke, and H. Silverstein, 73–84. Cambridge, MA: MIT Press.

Sosa, E. (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives* 13: 141–53.

Vogel, J. (2000). Reliabilism Leveled. *Journal of Philosophy* 97 (11): 602–23.

Vogel J. (2012). The Enduring Trouble with Tracking. In *The Sensitivity Principle in Epistemology*, edited by K. Becker and T. Black, 122–51. New York: Cambridge University Press.

Wallbridge, K. (2016a). Sensitivity and Higher-Order Knowledge. *Pacific Philosophical Quarterly* 97 (2). https://doi.org/10.1111/papq.12164.

Wallbridge, K. (2016b). Sensitivity Hasn't Got a Heterogeneity Problem. A Reply to Melchior. *Philosophia* 45 (2): 835–41.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

Zalabardo, J. L. (2012). *Scepticism and Reliable Belief*. Oxford: Oxford University Press.

Cite this article: Bjerring, J. C. and Gundersen, L. B. 2020. Higher-order knowledge and sensitivity. *Canadian Journal of Philosophy* 50: 339–349, doi:10.1017/can.2019.36

https://doi.org/10.1017/can.2019.36 Published online by Cambridge University Press