A Generality Problem for Bootstrapping and Sensitivity

«A Generality Problem for Bootstrapping and Sensitivity»

by Guido Melchior

Source:
Croatian Journal of Philosophy (Croatian Journal of Philosophy), issue: 1 (40) / 2014, pages: 31-47,
on www.ceeol.com.

Croatian Journal of Philosophy
Vol. XIV, No. 40, 2014

A Generality Problem for Bootstrapping and Sensitivity

GUIDO MELCHIOR
Department of Philosophy,
University of Graz

Vogel argues that sensitivity accounts of knowledge are implausible because they entail that we cannot have any higher-level knowledge that our beliefs are true, not false. Becker and Salerno object that Vogel is mistaken because he does not formalize higher-level beliefs adequately. They claim that if formalized correctly, higher-level beliefs are sensitive, and can therefore constitute knowledge. However, these accounts do not consider the belief-forming method as sensitivity accounts require. If we take bootstrapping as the belief-forming method, as the discussed cases suggest, then we face a generality problem. Our higher-level beliefs as formalized by Becker and Salerno turn out to be sensitive according to a wide reading of bootstrapping, but insensitive according to a narrow reading. This particular generality problem does not arise for the alternative accounts of process reliabilism and basis-relative safety. Hence, sensitivity accounts not only deliver opposite results given different formalizations of higher-level beliefs, but also for the same formalization, depending on how we interpret bootstrapping. Therefore, sensitivity accounts do not fail because they make higher-level knowledge impossible, as Vogel argues, and they do not succeed in allowing higher-level knowledge, as Becker and Salerno suggest. Rather, their problem is that they deliver far too heterogeneous results.

Keywords: Sensitivity, bootstrapping, generality problem, higher-level knowledge.

1. *Overview*

I will proceed as follows: In section 2, I present Vogel's argument for his claim that sensitivity accounts of knowledge entail that we cannot have higher-level knowledge that our beliefs are true, not false. In section 3, I analyze Becker's and Salerno's objection to this claim. In section 4, I introduce method-relative sensitivity and bootstrapping and Mooreanism as the relevant belief-forming methods. In section 5, I argue that sensitivity accounts face a generality problem with respect to bootstrapping and Mooreanism. In section 6, I show that the rival accounts of process reliabilism and basis-relative safety do not suffer from this particular generality problem.

2. *Vogel's objection against sensitivity*

Robert Nozick (1981) introduced the notion of sensitivity: S's belief that p is sensitive if and only if S were not to believe that p, if p were not true. Nozick interprets knowledge modally and argues that in cases of knowledge the belief tracks the truth. Nozick provides a first approximation of the definition of knowledge:

S knows that p iff
(1) p is true.
(2) S believes that p.
(3) If p weren't true, S wouldn't believe that p.
(4) If p were true, S would believe that p.

Sensitivity as specified in condition (3) is a necessary condition for knowledge according to Nozick.[1]

Vogel (2000: 611) argues that the view that sensitivity is necessary for knowledge—Vogel calls this view "counterfactual reliabilism"—, is too strong because it "makes it impossible for you to know that any of your beliefs is true, not false". In order to make this point Vogel presents the following inference:

(1) You know Omar has a new pair of shoes.
(2) You know that your belief that Omar has a new pair of shows is true, or at least not false.
(Vogel 2000: 609–610)

Vogel (2000: 610) argues that according to sensitivity accounts, (2) is false, even if (1) is true, i.e. you do know that Omar has a new pair of shoes, though "you do not know that your belief that Omar has new shoes is true, not false." Vogel concludes that sensitivity cannot be necessary for knowledge.

Vogel argues as follows for the claim that (2) is false, if sensitivity is necessary for knowledge: He paraphrases beliefs that one's belief is true or at least not false as:

- $B(\neg(B(p) \wedge \neg p))$

This belief is sensitive if and only if in the nearest possible world w, where $\neg(B(p) \wedge \neg p)$ is false—i.e. where $(B(p) \wedge \neg p)$ is true—S does not

[1] The general view that sensitivity is necessary for knowledge has been heavily criticized in literature, most prominently by Vogel (1987), Sosa (1999), Williamson (2000) and Kripke (2011). These authors offer counter-examples against Nozick by presenting instances of insensitive beliefs that we still regard as constituting knowledge. However, in this paper, I will focus on the particular problem of higher-level knowledge.

believe that $\neg(B(p) \wedge \neg p)$ anymore. In w, S believes that p. Vogel holds that if S believes that p, then S believes that S does not falsely believe that p, so if B(p), then $B(\neg(B(p) \wedge \neg p))$. Hence, in w, S believes that $\neg(B(p) \wedge \neg p)$ although $\neg(B(p) \wedge \neg p)$ is false in w. Therefore, $B(\neg(B(p) \wedge \neg p))$ is insensitive. But if sensitivity is necessary for knowledge, then one does not know that $\neg(B(p) \wedge \neg p)$, even if one knows that p. According to Vogel, this is unacceptable.[2]

To sum up: Vogel claims to have shown that if sensitivity is necessary for knowledge, then S cannot know of any of her beliefs whether they are true, not false. He regards this consequence as unacceptable and concludes that sensitivity is not necessary for knowledge and that "counterfactual reliabilism", which makes this assumption, is too strong.

3. *Vogel's mistake*

Becker (2006, 2007) and Salerno (2010) criticize Vogel for claiming to have shown that sensitivity accounts make it generally impossible for one to know that any of one's beliefs is true, not false. They both point out that '$B(\neg(B(p) \wedge \neg p))$' is only one possible formalization of beliefs that one's belief is true, not false, among others. They present two alternative formalizations:

(1) $B(B(p) \wedge p)$
(2) $B(B(p) \wedge \neg\neg p)$

Becker and Salerno think that (1) is the correct formalization of "My belief that p is true" and that (2) is the correct formalization of "My belief that p is not false", as opposed to '$B(\neg(B(p) \wedge \neg p))$', which Vogel suggests.

It might be a matter of discussion whether (1) and (2) really are correct formalizations of statements about the *truth* or *falsity* of one's beliefs. Rather, they are formalizations of 'I believe that p and it is the case that p' and 'I believe that p and it is not not the case that p', respectively, since they do not contain any claim about 'truth' or 'falsity'. The formalizations of 'My belief that p is true' and 'My belief that p is false' are actually formulations along the following lines:

(1*) $B(B(p) \wedge (p \text{ is true}))$
(2*) $B(B(p) \wedge \neg(p \text{ is false}))$

However, for the following argument, it is not important to decide whether (1) and (2) or (1*) and (2*) are the correct formalizations of "My belief that p is true" and "My belief that p is not false". The crucial point is that the underlying formal structure is a *conjunction* and not the negation of a conjunction as in Vogel's formalization '$B(\neg(B(p) \wedge \neg p))$'.

Notably, higher-level beliefs that my beliefs are true or not false turn out to be *sensitive*, if formalized as in (1) and (2) or (1*) and (2*).

[2] Sosa (1999) makes the same point.

Take S's belief that $B(p) \wedge p$. In evaluating its sensitivity, we have to take the nearest possible world w into account, where this belief is false. In w, $\neg(B(p) \wedge p)$ is true, i.e. it is false that S believes that p and p is true. This negation of a conjunction is equivalent with the disjunction $\neg B(p) \vee \neg p$. Hence $B(B(p) \wedge p)$ is sensitive if and only if in the nearest possible world either S does not hold a first-order belief that p or p is false or both. Becker argues correctly that $B(B(p) \wedge p)$ turns out to be sensitive, if B(p) itself is sensitive which means that sensitivity transmits from B(p) to $B(B(p) \wedge p)$.

> Now, if it is the case that the higher-level belief would be false because in W you do not hold the first-order belief that [p], then in W you also would not believe that you believe that [p]. In that case, your higher-level belief is truth-tracking: If your higher-level belief—that you have a (true) belief that [p]—were false, you would not hold that higher-level belief. On the other hand, if your higher-level belief is false in W because [p] is false, then, assuming as we have that your first-order belief is truth-tracking—that it is a case of knowledge—here again you would not believe that [p], and presumably would not hold your actual world higher-level belief in W. In either case, it is certainly possible for your higher-level belief to constitute knowledge—to satisfy sensitivity. (Becker 2006: 81–82)

Salerno makes the same point in a slightly different way. Salerno (2010: 80–81) argues that if the relevant counterfactual $(\neg B(p) \vee \neg p) \Box\!\!\rightarrow \neg B(B(p) \wedge p)$ is true then the following two counterfactuals are true:

- T1: $\neg p \Box\!\!\rightarrow \neg B(B(p) \wedge p)$
- T2: $\neg B(p) \Box\!\!\rightarrow \neg B(B(p) \wedge p)$

Salerno points out that T2 is trivially satisfied and that it is possible that T1 is true, since T1 captures the idea of knowledge of one's own correctness.

One can also argue that $B(B(p) \wedge p)$ is sensitive as follows: $B(B(p) \wedge p)$ is sensitive if and only if in the nearest possible world w where $B(p) \wedge p$ is false S does not believe that $B(p) \wedge p$ anymore. There are three ways a world might be that would make $B(p) \wedge p$ false and, therefore, $\neg B(p) \vee \neg p$ true:

w_1: '$B(p) \wedge \neg p$' is true: 'I believe that p' is true and p is false
w_2: '$\neg B(p) \wedge \neg p$' is true: 'I believe that p' is false and p is false
w_3: '$\neg B(p) \wedge p$' is true: 'I believe that p' is false and p is true

Using this notion, Becker and Salerno argue that if B(p) is sensitive, then the closest possible world where $B(p) \wedge p$ is false is either w_2 or w_3, but not w_1. One can argue for this claim as follows: If B(p) is sensitive, then in the closest possible world where p is false, S does not believe that p. Hence, w_1 where $B(p) \wedge \neg p$ is true is far off, if B(p) is sensitive. However, there are no reasons to assume that w_2 or w_3 is far off.[3]

[3] Whether w_3 is far off depends on whether $B(\neg(B(p) \wedge \neg p))$ fulfills Nozick's condition (4), often called the 'adherence condition'. This condition states that if p were true, S would believe that p. Condition (4) is fulfilled for B(p), iff in the close

Therefore, w_2 or w_3 is the closest possible world, where $B(p) \wedge p$ is false. It seems reasonable to assume that persons are minimally coherent insofar as they do not believe a conjunction $(a \wedge b)$, if they do not believe one of the conjuncts. Therefore, a minimally coherent person does not believe that $B(p) \wedge p$ in w_2 or w_3, since S does not believe that p.[4] Hence, in the nearest possible world where $B(p) \wedge p$ is false, S does not believe that $B(p) \wedge p$ anymore, if B(p) is sensitive. Therefore, $B(B(p) \wedge p)$ is sensitive, if B(p) is sensitive.

The same results can be obtained for the beliefs (2), (1*) and (2*), i.e. $B(B(p) \wedge \neg\neg p)$, $B(B(p) \wedge$ (p is true)), and $B(B(p) \wedge \neg$(p is false)). Each turn out to be sensitive if B(p) itself is sensitive.

Becker (2006) and Salerno (2010) claim that Vogel does not correctly formulate the sentence "I believe that I do not falsely believe that p" as $B(\neg(B(p) \wedge \neg p))$, since the correct formalization is rather $B(B(p) \wedge \neg\neg p)$. Salerno argues that if S claims to know that one of S's beliefs is not false, then S makes a claim about the truth of a belief that S *actually* has. But $\neg(B(p) \wedge \neg p)$ only states that a conjunction is false, i.e. that the disjunction $\neg B(p) \vee \neg p$ is true. This disjunction of the form 'I do not believe that p or p is false' does not entail that S *actually* beliefs that p, but $B(p) \wedge \neg\neg p$ does. Hence, beliefs that one's belief is true (or not false) are correctly formulated as $B(B(p) \wedge p)$ (or as $B(B(p) \wedge \neg\neg p)$).

I will leave open the question how to correctly formalize "I believe that I do not falsely believe that p", though I think that Salerno makes a good point. However, I think that the mere fact that it is a matter of philosophical discussion whether "I know that I do not falsely believe that p" can be true according to sensitivity accounts is already evidence against sensitivity accounts. Moreover, if it is true that one can know that $B(p) \wedge p$ but not that $\neg(B(p) \wedge \neg p)$, then it is an instance of failure of knowledge-closure, since '$B(p) \wedge p$' is stronger than '$\neg(B(p) \wedge \neg p)$'—a further implausible instance beyond the many others already noted by Nozick's critics.[5]

Interestingly, Vogel (2000: 611, fn. 17) was already aware of the shortcomings of his argument for the claim that sensitivity excludes higher-level knowledge, and describes his argument as "clumsy and roundabout". He argues that he had chosen this line of argument partly to avoid complications regarding counterfactuals with disjunctive an-

possible worlds where p is true, S believes that p. If it is fulfilled, then w_3 is far off, otherwise w_3 is nearby. However, even if w_3 is far off, then w_2 is still nearby and, the nearest possible world where $B(p) \wedge p$ is false.

[4] Exceptions to this kind of minimal coherence are cases of self-deception where persons believe to believe that p without actually believing that p.

[5] Sosa (1999) argues that the fact that it is easily possible that B(p) is sensitive and that $B\neg(B(p) \wedge \neg p)$ is insensitive offers evidence against sensitivity accounts of knowledge. Since $B(p) \wedge p$ is stronger than p, the sensitivity of $B(B(p) \wedge p)$ and the insensitivity of $B\neg(B(p) \wedge \neg p)$ seem even more problematic. Nozick (1981: 227–229) already noticed many instances of failed knowledge-closure, which most regard as highly counter-intuitive and as clear evidence that Nozick's sensitivity account is false. Notably, Nozick is willing to bite the bullet.

tecedents. Avoiding complications is probably not the best motivation for choosing a method, as Salerno (2010: 79) already notes. However, I think that Vogel comes to the right conclusion when he notes that the technical hazards that will arise already offer good reason to be suspicious about sensitivity accounts.

4. *Bootstrapping as a belief-forming method*

Nozick regards his definition of knowledge as outlined in section 2 only as a first approximation since we must also take the belief forming method into account. His refined definition is the following:

S knows, via method (or way of believing) M, that p iff
(1) p is true.
(2) S believes, via method or way of coming to believe M, that p.
(3) If p weren't true and S were to use M to arrive at a belief whether (or not) p, then S wouldn't believe, via M, that p.
(4) If p were true and S were to use M to arrive at a belief whether (or not) p, then S would believe, via M, that p.
(Nozick 1981, 179)[6]

Becker and Salerno argue that some higher-level beliefs about the truth of our beliefs are sensitive and, therefore, can constitute higher-level knowledge. However, neither pays attention to the method of forming these higher-level beliefs. In order to determine whether these higher-level beliefs meet Nozick's refined condition of knowledge we must take the belief forming method into account. I will argue that this deepens the worries about higher-level beliefs.

In the context of criticizing various reliability accounts for their impact on higher-level knowledge, Vogel (2000) introduces the notion of bootstrapping by presenting the case of Roxanne who comes to believe that her gas gauge is reliable by repeatedly looking at the gas gauge. Vogel reconstructs Roxanne's process of reasoning as follows:

Bootstrapping
(1) K(Tank is full at t_1) Reliable Process
(2) K(Gauge reads 'F' at t_1) Perception
(3) K(Gauge reads 'F' at t_1 & Tank is full at t_1) Logical Inference
(4) K(Gauge reads accurately at t_1) Logical Inference
(5) Repeat […]
(6) K(Gauge is reliable) Induction
(Vogel 2008, 519)

[6] Luper-Foy (1984) notes that Nozick's account of knowing via a method faces a technical problem when it comes to *one-sided methods* that recommend believing that p, but do not recommend believing that ¬p. If M is a one-sided method, then condition (3) of Nozick's definition can never be met. In order to solve the problem of one-sided methods, Luper-Foy (1984, 29) suggests replacing Nozick's condition (3) by the following condition that does not suffer from this problem. not-p → not-(S believes that p via M). For this reformulation of sensitivity see also Williamson (2000, 154).

Bootstrapping is not a clearly defined term in the literature. Rather, it is understood as referring to Roxanne's case and to cases that are sufficiently similar to it. The characteristic feature of bootstrapping is—as the notion already indicates—that one acquires knowledge about the reliability of an output producing source through inferences from outputs delivered by this very source. Similar cases of bootstrapping can be constructed for knowledge about the reliability of any information-output-systems, such as newspapers, but also about the *accuracy* of such systems.

In what follows, I will understand bootstrapping in a broader sense that includes, first, any reasoning process that produces beliefs about the reliability or accuracy of a belief-forming source by using beliefs produced by that very source in reasoning and, second, any kind of forming a belief that a belief that p is true or that it is not false partly through inference from this very belief that p. Although Vogel introduced bootstrapping in terms of beliefs about reliability, I am more concerned with beliefs about the truth of beliefs.[7]

It is usually assumed that there is a close connection between bootstrapping and Mooreanism, which can be understood as the claim that we can know that skeptical hypotheses are false partly through inference from our external world knowledge. On a first approximation, we can understand Mooreanism as the thesis that we can acquire knowledge about the truth and reliability of our external world beliefs via inferences from our experience-based beliefs about the external world (and beliefs that we have these beliefs).[8]

Let's now have a closer look at bootstrapping and Mooreanism as *belief forming* methods for $B(B(p) \wedge p)$, which can be understood as deductive inferences from beliefs (1) and (2) to belief (3) as follows:

(1) B(p)
(2) B(I believe that p)
(3) B(I believe that p and p)/B(My belief that p is true)

Below are the beliefs formally structured:

(1) $B(p)$
(2) $B(B(p))$
(3) $B(B(p) \wedge p))$

[7] Vogel (2000) does not establish a connection between higher-level beliefs about the truth of one's own beliefs and bootstrapping. He uses higher-level beliefs in order to argue that sensitivity-based knowledge accounts are too narrow and cases of bootstrapping in order to argue that process reliabilism is too wide. However, I think that Vogel is mistaken not to establish such a connection. There is wide agreement that there is a close connection between the problem of easy knowledge as presented by Cohen (2002 and 2005) and bootstrapping, but the presented problem of easy knowledge is also one of acquiring higher-level knowledge about the truth of one's own beliefs too easily.

[8] For a discussion of the question whether Moore really favored this kind of argument see Baumann (2009).

One can understand the inference from (1) and (2) to (3) in the following way: S believes that p, believes that she believes that p, and comes to believe via inference that her belief that p is true (or that she believes that p and p is the case). Although the process of inference might be one between *propositions*, the step from (1) and (2) to (3) is meant to be a *belief* forming process.

In the case of Mooreanism, p is an external world proposition such as the proposition "there is a computer in front of me", or COMPUTER for short.

If B(COMPUTER) is an experience based belief, then the following belief-forming process is an example of a belief-forming process via Moorean reasoning:

(1) B(COMPUTER)
(2) B(I believe that COMPUTER)
(3) B(My belief that COMPUTER is true)

5. *A generality problem for bootstrapping*

Becker and Salerno show that $B(B(p) \wedge p)$ can turn out to be sensitive if B(p) is sensitive, but only if we do not take the belief forming method into account. In case of bootstrapping and Mooreanism, the relevant belief-forming method is coming to believe (3) via inference from beliefs (1) and (2). However, whether sensitivity transmits from (1) to (3) in these cases depends on how the belief-forming method is specified. Suppose, that S comes to believe that $B(p) \wedge p$ via inference from B(p) and B(B(p)) and that O is the source that delivers B(p). There are at least two different ways to specify S's belief-forming method:

M_a: Inference from S's beliefs delivered from O, whether $B(p) \wedge p$ is true.

M_b: Inference from B(p) and p, whether $B(p) \wedge p$ is true.

M_a is a broader definition than M_b. In case of M_a, the method of coming to believe whether $B(p) \wedge p$ is true is specified as an inference from beliefs that come from the same source as B(p). This can be B(p) or any other of S's beliefs delivered from the same source as B(p). M_b is specified as an inference from the very same beliefs B(p) and B(B(p)). Every belief forming process that is a token of M_b is also a token of M_a, but the opposite entailment relation does not hold. If $B(q_1)$, $B(q_2)$… are delivered by the same source as B(p) and if p is not identical with q_1 or q_2… then the inference from $B(q_1)$, $B(q_2)$… to whether $B(p) \wedge p$ is true is an instance of M_a but not of M_b. As I will later argue in more detail, both interpretations M_a and M_b are in accordance with our intuitive understanding of bootstrapping and Mooreanism.

Things will become more illustrative by reflecting on concrete examples. In Vogel's bootstrapping case, Roxanne's method M_a of coming to believe whether *B(Tank is full at t_1)* $\wedge$ *Tank is full at t_1* is true is any inference from her beliefs about the tank acquired via looking at the gas gauge. Her belief forming method M_b is the inference from the very beliefs that *Tank is full at t_1* and *B(Tank is full at t_1).*

M_a and M_b can collapse but, crucially, they need not. If Roxanne believes that *Tank is full at t_1*, then Roxanne will probably use this very belief in order to come to know whether *B(Tank is full at t_1) $\wedge$ Tank is full at t_1* is true via bootstrapping. In this case, M_a and M_b collapse. However, if Roxanne does *not* believe at t_1 that *Tank is full at t_1*, then Roxanne will use other beliefs about her tank, formed via looking at her gas gauge (e.g. her belief that the tank is empty at t_1) and the two methods do not collapse.

Notably, if B(p) is sensitive, then B(B(p) $\wedge$ p) is *sensitive* if M_a is the belief-forming method, but it is *insensitive*, if M_b is the belief-forming method. The reason is the following:

M_a

If B(p) is sensitive and delivered by source O, then the nearest possible world w, where B(p) $\wedge$ p is false and S utilizes M_a is such that $\neg$B(p) $\wedge$ $\neg$p or $\neg$B(p) $\wedge$ p is true. In both cases, S does not believe that p in w. Hence, S uses inferences from other beliefs delivered from O. S does not come to believe via method M_a that B(p) $\wedge$ p in w and B(B(p) $\wedge$ p) turns out to be *sensitive*.

M_b

The nearest possible world, where B(p) $\wedge$ p is false and S utilizes M_b, has to be one in which S believes that p and believes that B(p). In this case B(p) $\wedge$ p fails to be true because p is false. However, S still comes to believe that B(p) $\wedge$ p via M_b and B(B(p) $\wedge$ p) turns out to be *insensitive*.[9]

To illustrate, I will analyze both Roxanne's case and the case of Moorean reasoning. If Roxanne's belief that *Tank is full at t_1* is sensitive, then the nearest possible world w where *B(Tank is full at t_1) $\wedge$ Tank is full at t_1* is false and Roxanne uses M_a is a world in which

[9] Nozick (1981: 236) argues that "knowledge also seems to be (always) closed under (known application of) adjudication: the inference from the premises p, q to the conjunctive conclusion p & q." Nozick is mistaken here. The inferences from B(p) and B(B(p)) to B(B(p) $\wedge$ p) does not transmit sensitivity and, therefore, also fails to transmit knowledge according to Nozick, if M_b is the belief forming method.

Note that if we formulate the sensitivity condition the way Luper-Foy (1984) and Williamson (2000) suggest, then S's belief turns out to be sensitive. In the nearest possible world w where B(p) $\wedge$ p is false, S does not believe that B(p) $\wedge$ p via M_b, since S does not believe that p in w. However, this paper aims at arguing that—in contrast to process reliabilism and safety—sensitivity faces a problem of instability, when it comes to bootstrapping as the belief-forming method. Hence, I regard the fact that subtle changes of the sensitivity condition lead to opposite results as further evidence for this instability. Moreover, neither Luper-Foy nor Williamson regards the reformulation of the sensitivity condition as a necessary condition for knowledge.

Roxanne does not believe that *Tank is full at* t_1 and, therefore, uses other beliefs about the tank, formed by looking at the gas gauge. (E.g. that the tank is empty at t_1.) In this case, Roxanne does not come to believe that *B(Tank is full at* t_1*)* ∧ *Tank is full at* t_1 and her belief of this conjunction turns out to be sensitive. If, for example, Roxanne believes that the tank is *empty* at t_1 by looking at the gas gauge she will come to believe *B(Tank is empty at* t_1*)* ∧ *Tank is empty at* t_1 instead.

However, the nearest possible world in which *B(Tank is full at* t_1*)* ∧ *Tank is full at* t_1 is false and Roxanne uses M_b is by definition a world in which Roxanne still believes that *B(Tank is full at* t_1*)* and that *Tank is full at* t_1 and, therefore, comes to believe that *B(Tank is full at* t_1*)* ∧ *Tank is full at* t_1 via M_b. As a result, Roxanne's belief turns out to be insensitive.

In the case of Moorean reasoning about COMPUTER, if B(COMPUTER) is an experience based sensitive belief, then M_a is the inference from my experience based beliefs to whether B(COMPUTER) ∧ COMPUTER is true. In this case, the nearest possible world in which B(COMPUTER) ∧ COMPUTER is false and where I use M_a in order to come to know whether B(COMPUTER) ∧ COMPUTER is true, is one where I do not believe that COMPUTER and where I use other experience based beliefs, e.g. B(EMPTY DESK). In this case, I do not come to believe that B(COMPUTER) ∧ COMPUTER anymore. But if the method is M_b, then it is the inference from my very beliefs that COMPUTER and that B(COMPUTER). In the nearest possible world where B(COMPUTER) ∧ COMPUTER is false and where I use M_b, I still believe that COMPUTER and that B(COMPUTER) and I still come to believe that B(COMPUTER) ∧ COMPUTER based on what I experience.

To sum up, if S comes to believe that B(p) ∧ p via inference from B(p) and B(B(p)), then there are two interpretations of bootstrapping and Mooreanism available. Given the wide method M_a, sensitivity transmits and B(B(p) ∧ p) is sensitive, if B(p) is sensitive. However, given the narrow method M_b, B(B(p) ∧ p) is insensitive.

Determining the correct method of belief-formation for the case of forming B(B(p) ∧ p)) via inference from B(p) and B(B(p)) is an instance of the *generality problem*. Conee and Feldmann (1998) and Conee (2012) present the generality problem as an objection to process reliabilism, which is the view that a belief is only justified or only constitutes knowledge if it is the result of a belief-forming process that reliably produces true beliefs. The generality problem arises because any particular belief is the result of a process *token*, i.e. of a process that occurs at a particular time and a particular place. But any such process token can be a token of many different process types with different values of reliability. The generality problem is then the problem of choosing the relevant process for determining the reliability of a particular process token.[10]

[10] For this description of the generality problem see Goldman (2011) who points out that this generality problem not only arises for process reliabilism but also for other externalist accounts of justification and knowledge.

Now the belief-forming token of inferring $B(p) \wedge p$ from p and B(p) is an instance of at least two different belief-forming types, method M_a and method M_b. According to M_a, S knows, but according to M_b, S does not know. This is an instance of the generality problem.

In order to determine whether sensitivity transmits from B(p) and B(B(p)) to $B(B(p) \wedge p)$ via inference, we have to decide whether M_a or M_b is the correct specification of bootstrapping and Mooreanism. Here, I do not see any clear case. Neither bootstrapping nor Moorean reasoning are clearly defined terms in literature. Rather, they are intuitively introduced. Broadly conceived, bootstrapping and Mooreanism are regarded as instances of epistemic circularity where a person acquires knowledge about the reliability or accuracy of a belief or a belief-forming source in an epistemically circular and, therefore, inappropriate way by drawing inferences from the belief or the source in question. M_a is the method of coming to believe that a belief is true via inference from the beliefs delivered from the same source, whereas M_b is the method of coming to belief whether a belief is true via inference from this very belief. Both formulations seem compatible with the intuitive understanding of bootstrapping and Mooreanism. However, M_a leads to the result that S knows, while M_b entails that S does not know; a highly problematic outcome.

One might object that the generality problem is a wide problem that equally concerns all externalist accounts of knowledge and that we cannot comment on particular instances of the generality problem without taking general accounts on the generality problem into account. I do not think that this objection is legitimate. Rather I think we can gain relevant insights about particular instances of the generality problem without committing ourselves to a particular general framework concerning the generality problem.

So far, we have elaborated a generality problem for higher-level beliefs if formulated as beliefs of one of the conjunctions $B(p) \wedge p$ or $B(p) \wedge \neg\neg p$. But what about beliefs that one does not falsely believe with the formal structure $\neg(B(p) \wedge \neg p)$ that Vogel originally suggested? It turns out that the same generality problem does not obtain, since the belief that $\neg(B(p) \wedge \neg p)$ is insensitive according to both the narrow and the wide reading.

The belief forming method for $B(\neg(B(p) \wedge \neg p))$ via bootstrapping or Mooreanism is the following deductive inference from belief (1) to belief (2):

(1) B(p)
(2) $B(\neg(B(p) \wedge \neg p))$

If this inference is specified as M_a, then it is the inference from the beliefs delivered by the same source that delivers B(p), whether $\neg(B(p) \wedge \neg p)$ is true. If it is specified as M_b, then it is the inference from the very belief B(p).

In the case of M_a, in the nearest possible world where $\neg(B(p) \wedge \neg p)$ is false and S uses M_a, S believes that p and reasonably still infers from her belief that p whether $\neg(B(p) \wedge \neg p)$. In w S still comes to believe $B\neg(B(p) \wedge \neg p)$ and $B\neg(B(p) \wedge \neg p)$ turns out to be *insensitive*. If the belief forming method is M_b, then it is *by definition* the inference from the very same belief B(p). Again, in the nearest possible world where $\neg(B(p) \wedge \neg p)$ is false and S uses M_b, S comes to believe that $\neg(B(p) \wedge \neg p)$ and $B\neg(B(p) \wedge \neg p)$ again turns out to be *insensitive*. Hence, no matter whether bootstrapping or Mooreanism are specified as M_a or M_b, the belief that $\neg(B(p) \wedge \neg p))$ turns out to be insensitive.[11]

In the case of $B\neg(B(p) \wedge \neg p))$, there is a tendency that the methods M_a and M_b collapse, but for $B(B(p) \wedge p)$ no such tendency applies. Whether S uses M_a or M_b, the nearest possible world where $\neg(B(p) \wedge \neg p))$ is false is in both cases one in which S still believes that p, and also uses her belief that p to infer whether $\neg(B(p) \wedge \neg p)$ is true. M_a and M_b collapse in this case. However, the nearest possible world in which $B(p) \wedge p$ is false and S uses M_a is one in which S does not believe that p anymore, and has to use other beliefs from the same source that delivered B(p). In this case, M_a and M_b do not collapse and the two methods deliver oppositional results with respect to the sensitivity of S's belief.

6. *Process reliabilism and basis-relative safety*

Next, I will investigate whether the particular generality problem for sensitivity and forming $B(B(p) \wedge p)$ via bootstrapping or Mooreanism also arises for the rival accounts of process reliabilism and of basis-relative safety. It will turn out that it does not.

6.1. *Process reliabilism*

Process reliabilism is the view that a belief is justified (or constitutes knowledge) only if it is the result of a belief-forming process that reliably produces true beliefs. The relevant process of coming to believe that $B(p) \wedge p$ is deductive inference from beliefs B(p) and B(B(p)).

We have seen that the belief forming method of deductive inference from B(p) and B(B(p)) can be specified in both a wide and narrow way.

[11] Salerno (2010: 74) argues that Vogel is mistaken in claiming that a belief that $\neg(B(p) \wedge \neg p)$ is sensitive. "Sometimes I know p and double-check my sources, thereby coming to know, additionally, that I do not believe falsely that p." Salerno is right that if S double-checks her sources by taking "independent sources" or an "independent method" into account, then S's belief that $\neg(B(p) \wedge \neg p)$ can easily be sensitive. However, double-checking one's sources is intuitively different from bootstrapping and Mooreanism, which are based on the idea that one comes to having beliefs about the truth of one's beliefs without taking additional evidence, other than those beliefs themselves, into account. Hence, Salerno misses the point with respect to bootstrapping and Mooreanism, because what we have in mind is that S comes to believe that $\neg(B(p) \wedge \neg p)$ via inference from B(p), i.e. without taking additional sources or methods into account, but in this case S's belief that $\neg(B(p) \wedge \neg p)$ turns out to be insensitive.

If O is the belief forming source of B(p), then the corresponding readings of the belief forming process for process reliabilism are as follows:

Process-type P_a

- P_a is the process of S's coming to believe that her beliefs delivered by O are true via inference from beliefs delivered by O, i.e. $B(B(p_1) \wedge p_1)$ formed via inference from S's beliefs delivered by O, $B(B(p_2) \wedge p_2)$ formed via inference from S's beliefs delivered by O…

Process-type P_b

- P_b is the process of S's coming to believe that her beliefs delivered by O are true via inference from those very beliefs, i.e. $B(B(p_1) \wedge p_1)$ formed via inference from $B(B(p_1))$ and $B(p_1)$, $B(B(p_2) \wedge p_2)$ formed via inference from $B(B(p_2))$ and $B(p_2)$… .[12]

In case of Mooreanism, p is an experience-based external world proposition like COMPUTER. The process of forming B(p) is then perception.

Following the two readings, Mooreanism can be understood as (a) S's coming to believe that her external world beliefs are true through inference from her external world beliefs in general or (b) S's coming to believe that her external world beliefs are true through inference from those very beliefs.

The crucial point is that in the case of process reliabilism, there is a *tendency* that P_a and P_b collapse, i.e. if S believes $B(B(p) \wedge p)$ via inference from the beliefs delivered by source O, then S believes $B(B(p) \wedge p)$ via inference from B(B(p)) and B(p).

If S is minimally rational in the sense that S believes that $a \wedge b$, only if S believes that a and believes that b, then: If S comes to believe that $B(p) \wedge p$ from beliefs delivered from the same source that delivers B(p), then S believes B(p) and B(B(p)) and infers $B(p) \wedge p$ from these very beliefs. Hence, P_a and P_b usually collapse in the case of process reliabilism, at least if S uses an inference that is intuitively an instance of bootstrapping.

The difference between process reliabilism and sensitivity-accounts is the following. Sensitivity: The nearest possible world where $B(p) \wedge p$ is false and S uses M_a is one where S does *not* believe B(p) and B(B(p)). The nearest possible world where $B(p) \wedge p$ is false and S uses M_b is one where S *does* believe B(p) and B(B(p)). Process Reliabilism: We take

[12] The process types P_a and P_b are defined relative to a particular person S and to beliefs delivered from a particular source O. However, it is also possible to define them relative to a *class* of persons, for example, or only to a particular belief B(p). If the process is restricted to a particular person and a particular belief, then process reliabilism tends to collapse with basis-relative safety. How to specify the belief-forming process adequately, given these alternatives, is an additional instance of the generality problem. However, the following results will be obtained for these alternative formulations of the process-types as well.

processes into account that produce $B(B(p) \wedge p)$. Hence, S believes that $B(p) \wedge p$ and if S is minimally rational, then S also that B(p) and that p, no matter whether we interpret the process as P_a or as P_b.

For determining whether the process of coming to believe $B(p) \wedge p$ *is* reliable we have to take Goldman's distinction between *conditionally* reliable processes and *outright* reliable processes into account. Goldman (1979: 340) argues that a "process is conditionally reliable when a sufficient proportion of its output beliefs are true *given that its input-beliefs are true*." He argues that we need the notion of 'conditional reliability' for processes like reasoning and memory, because the requirement of outright reliability is too strong. A reasoning process cannot be expected to produce true beliefs if it is applied to false premises. The same goes for memory. Outright reliable processes, in contrast, generally tend to produce beliefs that are true.

The deductive inference from believes that B(p)) and that p is part of a larger process of coming to believe that $B(p) \wedge p$, that also includes the belief forming processes of the believed premises p and B(p). The mere belief-forming process of inference from p and B(p) to $B(p) \wedge p$ is *conditionally* reliable since it is a deductive inference. The larger process, in contrast, that also takes the belief formation of beliefs that p and that B(p) into account can have outright reliability.

In case of Moorean reasoning, this larger process contains perception as the belief forming process for B(p), probably introspection as the one for B(B(p)) and a deductive inference. Whether this process is outright reliable depends on whether the involved processes of perception and introspection are outright reliable. However, if these two processes have sufficient outright reliability, then the overall process has outright reliability as well. Hence, outright reliability and, therefore, justification and knowledge, can transmit from B(p) and B(B(p)) to $B(B(p) \wedge p)$.

6.2. *Basis-relative safety*

Ernest Sosa (1999: 147) suggests replacing sensitivity by the alternative modal principle *safety*: A belief by S that p is safe if and only if S would believe that p only if it were so that p. Sosa (2007) later replaced his initial concept of safety by *basis-relative safety* which relativizes safety to a belief forming method.

> A belief that p is *basis-relative safe*, then, if and only if it has a basis that it would (likely) have only if true. By contrast, a belief that p is *basis-relative sensitive* if and only if it is based on a basis such that if it were false that p, then not easily would the believer believe that p on the same basis. More plausibly, then, what is required for knowledge is basis-relative safety, rather than outright safety. (Sosa, 2007: 26)

Since safety is a modal principle, we can also formulate it by using the notion of possible worlds. A belief B(p) formed via method M is basis-relative safe if and only if in most nearby possible worlds where S believes that p via M, B(p) is true.[13] Is B(B(p) $\wedge$ p) formed via inference from B(p) and B(B(p)) basis-relative safe? Suppose B(p) is delivered by source O. Again, there is a wide and a narrow reading of the inference from B(p) and B(B(p)) to B(B(p) $\wedge$ p).

Ms_a B(B(p) $\wedge$ p) is basis-relative safe iff in most nearby possible worlds where S believes that B(p) $\wedge$ p via inference from O, B(B(p) $\wedge$ p) is true.

Ms_b B(B(p) $\wedge$ p) is basis-relative safe iff in most nearby possible worlds where S believes that B(p) $\wedge$ p via inference from B(p) and B(B(p)), B(B(p) $\wedge$ p) is true.

As in the case of process reliabilism, there is a tendency that Ms_a and Ms_b collapse. In most nearby possible worlds where S believes that B(p) $\wedge$ p via inference from the beliefs delivered by O, S believes that B(p) $\wedge$ p via inference from the very beliefs B(p) and B(B(p)). The reason is, again, the following.

If S is minimally rational with respect to believing that a $\wedge$ b, only if S believes that a and that b, then: If S comes to believe that B(p) $\wedge$ p from beliefs delivered from the same source that delivers B(p), then S believes B(p) and B(B(p)) and infers from these very beliefs, if S performs an instance bootstrapping. Hence, there does not arise the same generality problem as in the case of sensitivity.

The difference between basis-relative safety and sensitivity is that in the case of basis-relative safety we only consider worlds where S believes that B(p) $\wedge$ p and, if minimally coherent, also believes that B(p) and believes that p. Hence, any world in which S uses Ms_a or Ms_b is one, in which S also believes that B(p) and believes that p and uses these two beliefs as premises. However, in the case of sensitivity, in the nearest possible world in which B(p) $\wedge$ p is false, S only believes that B(p) and that p in the case of M_b, but not in the case of M_a.

But is B(B(p) $\wedge$ p) basis-relative safe? This depends on the question whether B(p) and B(B(p)) are basis-relative safe. In the case of Mooreanism, B(p) is basis-relative safe if and only if in most nearby possible worlds, where S believes that p on the basis of S's experience, B(p) continues to be true. B(B(p)) is basis-relative safe if and only if in most nearby possible worlds, where S believes that B(p) on the basis of method M (for example introspection), B(B(p)) continues to be true.

In any possible world where p and B(p) are true, B(p) $\wedge$ p is also true. Therefore, in most nearby possible worlds where S believes that B(p) $\wedge$ p via inference from p and B(p), B(p) $\wedge$ p is true. Hence, B(B(p) $\wedge$ p) is basis-relative safe if B(p) and B(B(p)) are basis-relative safe. Basis-relative safety transmits from B(p) and B(B(p)) to B(B(p) $\wedge$ p).

[13] The following is Pritchard's (2007: 281) definition of a basis-relative safety principle: *S*'s belief is safe *iff* in most near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world the belief continues to be true.

7. *Conclusion*

We can formulate higher-level beliefs that our beliefs are true in different ways, either as $B(\neg(B(p) \wedge \neg p))$ or as $B(B(p) \wedge p)$. Beliefs of the first kind are insensitive. Beliefs of the second kind are sensitive according to a wide reading of bootstrapping but insensitive according to a narrow reading. This poses a particular generality problem for bootstrapping, one that does not arise for process reliabilism or basis-relative safety. It is a well-known fact that sensitivity accounts deliver a heterogonous picture of knowledge because they imply counterintuitive instances of closure failure. However, a heterogonous picture also results from specific instances of the generality problem.

Acknowledgements The research was funded by the Austrian Science Fund (FWF): J 3174-G15. I am indebted to Martina Fürst for helpful comments on this paper.

References

Baumann, P. 2009. "Was Moore a Moorean? On Moore and scepticism." *European Journal of Philosophy* 17 (2):181–200.

Becker, K. 2006. "Is counterfactual reliabilism compatible with higher-level knowledge?" *Dialectica* 60 (1):79–84.

Becker, K. 2007. *Epistemology modalized*. New York and London: Routledge.

Cohen, S. 2002. "Basic knowledge and the problem of easy knowledge." *Philosophy and Phenomenological Research* 65 (2):309–329.

________2005. "Why basic knowledge is easy knowledge." *Philosophy and Phenomenological Research* 70 (2):417–430.

Conee, E. 2012. "The specificity of the generality problem." *Philosophical Studies* 163(3): 751–762.

Conee, E. and R. Feldman. 1998. "The generality problem for reliabilism." *Philosophical Studies* 89 (1): 1–29.

Goldman, A. I. 1979. "What is justified belief?" *Justification and Knowledge*, George Pappas (ed), Dordrecht: Reidel: 1–23.

Goldman, A. I. 2011. "Reliabilism", *The Stanford Encyclopedia of Philosophy (Spring 2011 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2011/entries/reliabilism/>.

Kripke S. A. 2011. "Nozick on knowledge." *Philosophical troubles. Collected papers, volume I*. Oxford: Oxford University Press: 162–224.

Luper-Foy, S. 1984. "The epistemic predicament: knowledge, Nozickian tracking, and scepticism." *Australasian Journal of Philosophy* 62 (1):26–49.

Nozick, R. 1981. *Philosophical explanations*. Cambridge: Cambridge University Press.

Pritchard, D. 2007. "Anti-luck epistemology." *Synthese* 158 (3): 277–297.

Salerno, J. 2010. "Truth tracking and the problem of reflective knowledge." *Knowledge and Skepticism*. J. K. Campbell, M. O'Rourke and H. S. Silverstein (eds.). Cambridge/Mass.: MIT Press: 73–83.

Sosa, E. 1999. "How to defeat opposition to Moore." *Philosophical Perspectives* 13: 141–153.

Sosa, E. 2007. *A virtue epistemology: Apt belief and reflective knowledge, Volume I*, Oxford: Oxford University Press.

Vogel, J. 1987. "Tracking, closure and inductive knowledge." *The possibility of knowledge: Nozick and his critics.* S Luper-Foy (ed.), Totowa, NJ: Rowman and Littlefield: 197–215.

Vogel, J. 2000. "Reliabilism leveled." *Journal of Philosophy* 97 (11): 602–623.

Vogel, J. 2008. "Epistemic bootstrapping." *Journal of Philosophy* 105 (9): 518–539.

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