

The heterogeneity problem for sensitivity accounts

[Penultimate draft]

Published in *Episteme*

Abstract Offering a solution to the skeptical puzzle is a central aim of Nozick's sensitivity account of knowledge. It is well-known that this account faces serious problems. However, because of its simplicity and its explanatory power, the sensitivity principle has remained attractive and has been subject to numerous modifications, leading to a 'second wave' of sensitivity accounts. I will object to these accounts, arguing that sensitivity accounts of knowledge face two problems. First, they deliver a far too heterogeneous picture of higher-level beliefs about the truth or falsity of one's own beliefs. Second, this problem carries over to bootstrapping and Moorean reasoning. Some beliefs formed via bootstrapping or Moorean reasoning are insensitive, but some closely related beliefs in even stronger propositions are sensitive. These heterogeneous results regarding sensitivity do not fit with our intuitions about bootstrapping and Moorean reasoning. Thus, neither Nozick's sensitivity account of knowledge nor any of its modified versions can provide the basis for an argument that bootstrapping and Moorean reasoning are flawed or for an explanation why they seem to be flawed.

Key words

Sensitivity; higher-level knowledge; bootstrapping; Moorean reasoning; skepticism

1 Sensitivity

Overview: In section 1, I will introduce Nozick's original sensitivity account. In section 2, I will present insensitive and sensitive higher-level beliefs about the truth of one's own beliefs. In section 3, I will discuss bootstrapping and Moorean reasoning in more detail and in section 4, I will argue that we can construct insensitive and sensitive instances of bootstrapping and Moorean reasoning by repeating those processes that lead to insensitive and sensitive higher-level beliefs. In section 5, I will discuss the consequences for Nozick's sensitivity account and for various proposed modifications. The heterogeneity of the resulting picture affects each of the proposed sensitivity accounts in a different way, although the outcome is in some respect the same. Because of this heterogeneity problem, no sensitivity account for knowledge can adequately explain why bootstrapping and Moorean reasoning *are* flawed or why they *seem* flawed.

The notion of sensitivity has been introduced by Nozick (1981) who interprets knowledge modally. He argues that in cases of knowledge the belief tracks the truth. To a first

approximation Nozick suggests that 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 and (4) If p were true, S would believe that p . In accordance with the literature, I will understand condition (3) as the sensitivity-condition. Nozick (1981, 179) refines his first definition of knowledge in order to take methods of belief formation into account.

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 argues that sensitivity is a necessary condition for any kind of knowledge. Nozick's tracking account of knowledge has been heavily attacked for two main reasons. First, Nozick's view that sensitivity is necessary for any kind of knowledge has been criticized for being too strong (Vogel 1987, Sosa 1999, Williamson 2000). These authors present examples of insensitive beliefs that plausibly constitute knowledge. Second, the fact that Nozick's sensitivity account entails the failure of the highly plausible principle of knowledge closure is usually regarded as strong evidence against it.¹

However, because of its simplicity and its explanatory power, the sensitivity principle has remained surprisingly attractive and has been subject to numerous modifications, leading to a 'second wave' of sensitivity accounts.² These modifications attempt to handle cases of insensitive knowledge and/or the problem of the failure of knowledge closure, thereby providing alternative sensitivity-based solutions to the skeptical problem.

¹ See Kripke (2011). It is disputed whether Kripke's criticism of Nozick still holds if we take the belief-forming method into account, as Nozick's sensitivity account demands. See Adams and Clarke (2005) and Baumann (2012).

² DeRose (1995 and 2010) elaborates a sensitivity-based version of contextualism. Black (2002, 2008) favors a sensitivity-based *invariantism*. Roush (2005, 43) proposes a recursive sensitivity-based definition of knowledge according to which knowledge is closed under known entailment. For an overview of the contemporary debate see Becker and Black (2012). They talk about the "resilience of sensitivity".

2 The heterogeneity problem for higher-level beliefs

Vogel (2000) argues that if sensitivity is necessary for knowledge, then we cannot have any higher-level knowledge about the truth or falsity of our own beliefs. Becker (2006 and 2007) and Salerno (2010) object to Vogel, arguing that there are higher-level propositions that we can sensitively believe and, therefore, know according to sensitivity accounts of knowledge.

In this section, I will briefly present Vogel's argument and Becker's and Salerno's responses. I will conclude that the resulting picture of the sensitivity of higher-level beliefs is far too heterogeneous to provide a plausible picture of higher-level knowledge in terms of sensitivity.

2.1 Insensitive higher-level beliefs

Vogel (2000, 609-610) argues that sensitivity accounts of knowledge are flawed because they preclude us from having any higher-level knowledge about the truth of our beliefs. He presents the following inference for arguing that higher-level beliefs are insensitive.

- (1) You know Omar has a new pair of shoes.
- (2) You know that your belief that Omar has a new pair of shoes is true, or at least not false.

Vogel (2000, 610) claims that if sensitivity is necessary for knowledge then (2) is false i.e. "you do not know that your belief that Omar has new shoes is true, not false." He argues that this result is unacceptable and concludes that sensitivity is not necessary for knowledge.

Vogel presents the following argument for his claim that (2) is false if sensitivity is necessary for knowledge: He takes the belief that one does not falsely believe that p to be equivalent to:

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

Since the believed proposition $\neg(B(p) \wedge \neg p)$ is equivalent with the *disjunction* $\neg B(p) \vee p$, I will abbreviate it as d . In the nearest possible world w , where d is false, S believes that p . Vogel claims that if one believes that p , then one believes that one does not falsely believe that p , so if $B(p)$, then $B(d)$. Hence, in w both $B(d)$ and $\neg d$ hold. Therefore, $B(d)$ is insensitive. This is Vogel's argument.

Notably, $B(d)$ is insensitive, even if $B(p)$ is sensitive. The reason is that for evaluating the sensitivity of the two beliefs, we consider different possible worlds. In case of $B(p)$ we consider a world w_1 where just p is false. In case of $B(d)$ we consider w_2 where $\neg d$ is true

which is farther away from the actual world than w_1 . However, p entails d , since p is one of the disjuncts of the disjunction d . Hence, even if the belief in a stronger proposition p is sensitive, the belief in d is still insensitive.³

2.2 Sensitive higher-level beliefs

Beliefs of the form $B(d)$ are insensitive higher-level beliefs about the truth-value of one's own beliefs. Notably, there are some closely related higher-level beliefs that are sensitive. Take higher-level beliefs which have the following propositions as content.

- I believe that p and p

The underlying formal structure of this proposition is the following conjunction.

- $B(p) \wedge p$

Since this proposition is a *conjunction*, I will abbreviate it as c . The following propositions have a similar underlying structure:

- I believe that p and p is true
- My belief that p is true

In the following, I will use c to abbreviate any of the propositions 'I believe that p and p ', 'I believe that p and p is true' and 'My belief that p is true'.

Notably, beliefs in this conjunction are *sensitive* if $B(p)$ is sensitive; i.e. sensitivity carries over from $B(p)$ to $B(c)$.⁴ For determining whether $B(c)$ is sensitive, we have to look at the nearest possible world where c is false. If $B(p)$ is sensitive, then, by definition, any world where $B(p) \wedge \neg p$ is true is far off. Thus, any nearby world where c is false is such that $\neg B(p) \wedge \neg p$ or $\neg B(p) \wedge p$ is true. Moreover, we can assume that any world where S believes a conjunction without believing each of the conjuncts is also far off. Therefore, S does not believe c since she does not believe that p . Therefore, $B(c)$ is sensitive, if $B(p)$ is sensitive. Simply said, $B(\text{My belief that } p \text{ is true})$ is sensitive, if $B(p)$ is sensitive because in the nearest possible world where 'My belief that p is true' is false, I do not believe that p and, therefore, do not believe that I truly believe that p .

The reason why $B(c)$ is sensitive but $B(d)$ is insensitive is simply that in evaluating their sensitivity we take different possible worlds into account. In the nearest possible world where

³ Sosa (1999) argues that this is a counter-example against the claim that knowledge is sensitive true believing.

⁴ This point is also made by Becker (2006 and 2007) and Salerno (2010) and has probably been anticipated by Vogel (2000, 611, FN 17). For an alternative presentation of this problem see Melchior (2014a).

d is false, $B(p) \wedge \neg p$ is true, i.e. S still believes that p . In the case of $B(c)$, the nearest possible world is one where either $\neg B(p) \wedge \neg p$ or $\neg B(p) \wedge p$ is true, i.e. a world where S does not believe that p .

If the bi-valence principle is true, then $B(p) \wedge p$ is equivalent to $B(p) \wedge \neg\neg p$, which can be understood as a formalization of the following propositions:

- I believe that p and p is not false
- My belief that p is not false

Beliefs in these propositions are sensitive for the same reason as $B(c)$ is sensitive.

2.3 The belief forming method of higher-level beliefs

So far, we have not taken the belief forming method into account for determining the sensitivity of beliefs, as Nozick suggests. In this case, we achieve the same results. Take the following case of forming $B(d)$ by inferring it from the believed proposition p .

- (1) $B(p)$
- (2) $B(d)$ ⁵

If I form $B(d)$ by inferring it from the believed proposition p , then sensitivity does not carry over from $B(p)$ to $B(d)$, since in the nearest possible world where I falsely believe that p and where I use the inference from p to come to believe whether I falsely believe that p , I still come to believe that I am not falsely believing that p .⁶

Let's also have a closer look at processes of coming to believe that c on the basis of believing p . Take the following deductive inference from (1) and (2) to (3):

- (1) $B(p)$
- (2) $B(\text{I believe that } p)$
- (3) $B(\text{My belief that } p \text{ is true})$

⁵ This formulation of bootstrapping describes a way of coming to believe the proposition d via inference from the believed proposition p . It is not meant to be an inference from the proposition 'I believe that p ' to the proposition 'I believe that I do not falsely believe that p '.

⁶ Becker and Salerno achieve their results without taking the belief-forming method into account and therefore only for Nozick's first approximate definition of knowledge.

If (1) is sensitive, then (3) is also sensitive, if formed via inference from the propositions believed in (1) and (2). Furthermore, sensitivity carries over from $B(p)$ to $B(B(p) \wedge \neg\neg p)$. Thus, we acquire the same results even when taking plausible belief forming methods into account.

2.4 The heterogeneity-problem for higher-level knowledge

Vogel (2000) argues that if sensitivity is necessary for knowledge, then we cannot have any higher-level knowledge about the truth of our own beliefs. His line of argumentation is that $B(d)$ is insensitive and, therefore, we do not know c . Becker (2006 and 2007) and Salerno (2010) are right that Vogel is mistaken in this respect, for we can know c without knowing d , if knowledge is sensitive true belief.⁷ Moreover, we can also acquire sensitive higher-level beliefs that our beliefs are not false, if the underlying formal structure is equivalent to c . However, this reply to Vogel does not solve the problem of higher-level knowledge for sensitivity accounts but rather displaces it.

Some higher-level beliefs about the truth of single beliefs are sensitive and some are insensitive. Beliefs with the formal structure $B(c)$, which are beliefs of the conjunction $B(p) \wedge p$ turn out to be sensitive, but beliefs of the form $B(d)$, which are beliefs equivalent to the disjunction $\neg B(p) \vee p$ are insensitive. Since c entails d , the sensitively believed propositions are logically stronger than the insensitively believed ones. Sosa (1999) argues that the fact that $B(p)$ is sensitive although the belief in the disjunction $B(d)$ is insensitive provides evidence against the claim that knowledge is sensitive true believing. Since c is stronger than p , the fact that $B(c)$ is sensitive whereas $B(d)$ is insensitive provides even stronger evidence for Sosa's point. We want an account of knowledge that allows one to know the weaker propositions d if we know the stronger propositions c . Thus, the outcome for higher-level beliefs in terms of sensitivity is too heterogeneous to provide a useful explanation.⁸

⁷ In a footnote Vogel (2000, 611) calls his argumentation "clumsy and roundabout" that one fails to know that one's belief that p is true, if one fails to know that one's belief is not false. It has been shown that Vogel's argumentation is false. However, his conclusion that problems about knowing the one without knowing the other offer good reason to be suspicious about sensitivity accounts is exactly the conclusion of this paper.

⁸ I assumed that the proposition *I do not falsely believe that p* has the underlying formal structure $\neg(B(p) \wedge \neg p)$, whereas the formal structure of *My belief that p is not false* is $B(p) \wedge \neg\neg p$. It is disputable whether we can clearly assign these formal structures to these propositions. However, ambiguities about the underlying formal structures of the two propositions further increase the problems of sensitivity accounts.

3 Bootstrapping and Moorean reasoning

Vogel (2000 and 2008) introduced bootstrapping by means of the example of Roxanne who comes to know that the gas gauge of her car is reliable via 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)⁹

Similar cases can be constructed for knowledge about the reliability of any other measurement devices or, more generally, for any information-output-systems, such as newspapers.

Vogel argues that it is unacceptable that Roxanne can come to know via bootstrapping that her gas gauge is reliable, since bootstrapping would lead to knowledge about the reliability of a belief-forming source in much too easy a way. Vogel argues that the fact that reliabilism sanctions every step that Roxanne takes provides strong evidence against reliabilism. Cohen (2002 and 2005) argues more generally that bootstrapping does not only provide a problem for reliabilism but for any account of knowledge that allows one to gain knowledge from a source without having prior knowledge about the reliability of this source.¹⁰

⁹ It is important to be clear about how to interpret Vogel's formulation. Bootstrapping is a process of acquiring inferential knowledge about the reliability of a gas gauge. It is *not* a process of inferences between propositions about knowledge. Step (3), for example, is a potential way of coming to know the proposition 'Gauge reads "F" at t_1 & Tank is full at t_1 ' via an inference from the two known propositions 'Tank is full at t_1 ' and 'Gauge reads "F" at t_1 '. It is *not* an inference from the propositions 'Roxanne knows that Tank is full at t_1 ' and 'Roxanne knows that gauge reads "F" at t_1 ' to the proposition 'Roxanne knows that (Tank is full at t_1 and Gauge reads "F" at t_1)'.

¹⁰ Vogel (2000) argues that sensitivity accounts of knowledge are too narrow since they do not allow any higher-level knowledge that one's own beliefs are not false. However, when arguing against process reliabilism he simply assumes that bootstrapping is a flawed epistemic process. In this respect, Vogel does not establish any connection between higher-

Bootstrapping is not a clearly defined term in literature. It is rather understood as referring to Roxanne's case and to cases that are sufficiently similar to it. In the following, I will understand bootstrapping in a broader sense that also includes reasoning about the *accuracy* of a belief-forming source based on beliefs delivered from that very source.¹¹

One reason for the philosophical relevance of bootstrapping is its close connection to Mooreanism, understood as the claim that one can know that the skeptical hypotheses are false partly through inference from one's external world knowledge.¹² Moorean reasoning can be formulated in different ways. We can formulate it as follows as to demonstrate that it is a kind of bootstrapping.

Moorean reasoning

- (1) $K(p)$ Reliable process via perception
- (2) $K(I \text{ believe that } p)$ Introspection
- (3) $K(I \text{ believe that } p \text{ and } p)$ Logical inference
- (4) $K(\text{My belief that } p \text{ is true})$ Logical inference
- (5) Repeat
- (6) $K(\text{My perception is reliable})$ Induction

Moorean reasoning, understood this way, is a process of inference from one's perception based external world beliefs to the conclusion that one's perception is reliable.

It is far from clear whether Moore really endorsed this kind of argument.¹³ However, in the following, I will understand Moorean reasoning as an instance of bootstrapping, i.e. as inference from experience-based beliefs about the external world (and beliefs that we have these beliefs) to beliefs about the truth or reliability of experience-based beliefs. The method

level beliefs and bootstrapping. In this paper, I establish such a connection by showing how repetition of the inferences leading to higher-level beliefs leads to an instance of Moorean reasoning. In this respect, I diverge dialectically from Vogel.

¹¹ I think it is not useful to restrict bootstrapping to cases of *inductive* reasoning about the *reliability* of output delivering sources. It is widely acknowledged that the problem of easy knowledge as introduced by Cohen (2002) is closely related to the bootstrapping problem, although the problem of easy knowledge also applies to *deductive* inferences about the *truth* of one's own beliefs.

¹² For Vogel's reconstruction of Mooreanism that is analogous with bootstrapping see Vogel (2000, 618).

¹³ For a discussion of this question see Baumann (2009). Moreover, Mooreanism is often understood as the argument: Here is a hand. Here is another one. Therefore, the external world exists.

of Moorean reasoning can then be understood as consulting my experience in order to come to believe whether my experience-based beliefs are true or reliably formed.

Moorean reasoning is intuitively an inadequate response to the skeptical challenge but there exist various explanations as to why it is inadequate.¹⁴ According to sensitivity accounts of knowledge, bootstrapping cannot constitute knowledge because it cannot deliver a negative result, i.e. we would still believe, via bootstrapping, that a source is reliable if it were unreliable.¹⁵ Moreover, I come to believe that I am not a BIV by relying on my experiences and my experience based beliefs no matter whether I am a BIV or not. If I am not a BIV, then I reach the true conclusion that I am not a BIV because I base my reasoning on true external world beliefs; but if I am a BIV I reach the false conclusion that I am not a BIV because I am deceived and, hence, base my reasoning on false external world beliefs.

4 The heterogeneity problem for bootstrapping and Moorean reasoning

In this section, we will see that the heterogeneity problem for higher-level knowledge carries over to bootstrapping and Moorean reasoning. Some instances of bootstrapping and Moorean reasoning deliver insensitive beliefs, as expected, but some closely related processes deliver *sensitive* beliefs. Again, the resulting picture of bootstrapping and Moorean reasoning in terms of sensitivity is too heterogeneous to provide a plausible explanation for the (alleged) defectiveness of bootstrapping and Moorean reasoning.

¹⁴ Crispin Wright (2002 and 2004) claims that Moore's argument fails to be cogent, since warrant does not transmit from the premises to the conclusion. Furthermore, it is often implicitly assumed that Moorean reasoning is flawed because it cannot offer a persuasive argument to the skeptic. This view is criticized by Pryor (2004) who argues that Moorean arguments are "persuasively crippled", but in terms of their justificatory structure they are nothing to be ashamed of. One could also stipulate that bootstrapping and Moorean reasoning are flawed without offering any explanation at all. This seems an ad hoc and unconvincing strategy, though it seems to be Vogel's argumentation line. Vogel (2000, 615) rejects the view that insensitivity is the problem of bootstrapping, since he rejects insensitivity in general as a criterion for knowledge, but without offering any alternative explanation for the epistemic defectiveness of bootstrapping and Moorean reasoning. For Vogel's case against sensitivity see Vogel (1987). I will not handle these alternative explanations for why bootstrapping and Moorean reasoning are flawed or for they seem flawed here.

¹⁵ A similar strategy is suggested by Titelbaum (2010) and Douven and Kelp (2013), who argue that bootstrapping cannot constitute knowledge because it could never provide a negative result. Weisberg (2012) calls this the *no risk, no gain diagnosis*.

4.1 Insensitive Moorean reasoning

Take the following belief forming process via deductive inference.

- (1) $B(p_1)$
- (2) $B(\text{I do not falsely believe that } p_1)$
- (3) Repeat for $p_2 \dots p_n$
- (4) $B((\text{I do not falsely believe that } p_1) \text{ and } \dots \text{ and } (\text{I do not falsely believe that } p_n))$

Belief (4) is understood to have the following formal structure:

- $B(\neg(B(p_1) \wedge \neg p_1) \wedge \dots \wedge \neg(B(p_n) \wedge \neg p_n))$

In short:

- $B(d_1 \wedge \dots \wedge d_n)$

Belief (4) is insensitive. The nearest possible world w where (4) is false is one where one of its conjuncts d_i is false. But in w S still comes to believe that d_i via bootstrapping. Moreover, S also comes to believe that $(d_1 \wedge \dots \wedge d_n)$ via steps (1)–(3). Sensitivity does not carry over from step (1) to step (2) and neither to step (4).

In the case of Moorean reasoning, p is an experience-based external world proposition such as *There is a computer in front of me*, in short COMPUTER.¹⁶ Take the following example of Moorean reasoning.

- (1) $B(\text{COMPUTER})$
- (2) $B(\text{I do not falsely believe that COMPUTER})$
- (3) Repeat for DESK, GLASS OF WATER...
- (4) $B((\text{I do not falsely believe that COMPUTER}), (\text{I do not falsely believe that DESK}), \dots)$

In the nearest possible world where (4) is false, I still come to believe (4) via Moorean reasoning.

Beliefs are also insensitive, if an explanation for my potential deception such as a skeptical hypothesis is incorporated into the conclusion as in the following case.

¹⁶ All the following results regarding Moorean reasoning analogously apply to Roxanne's case of bootstrapping, though I will focus on the case of Moorean reasoning.

(4') B((I am not a BIV deceived in falsely believing that COMPUTER), (I am not a BIV deceived in falsely believing that DESK)...)

We can also form beliefs that *all* my beliefs from a certain source are true or reliably formed via bootstrapping and Moorean reasoning. It turns out that such beliefs are also *insensitive*. One can reach such a conclusion via the following deduction.

- (1) B(p_1)
- (2) B(I believe that p_1)
- (3) B(My belief that p_1 is true)
- (4) Repeat for $p_2 \dots p_n$
- (5) B($B(p_1) \dots B(p_n)$ are all my beliefs from source O)
- (6) B(All my beliefs from source O are true.)

In case of Moorean reasoning, such an inference can take the following form.

- (1) B(COMPUTER)
- (2) B(I believe that COMPUTER)
- (3) B(My belief that COMPUTER is true)
- (4) Repeat for DESK, GLASS OF WATER...
- (5) B($B(\text{COMPUTER}), B(\text{DESK}) \dots$ are all my beliefs based on my current experience)
- (6) B(All my beliefs based on my current experience are true.)

Beliefs of (6) formed via inference from (1)-(5) turn out to be *insensitive*.

We obtain the same results about insensitivity via *inductive* bootstrapping, i.e. by directly reaching the following conclusions (5') and (5'') from (1)-(4).

- (1) – (4)
- (5') B(All my beliefs from source O are true.)
- (5'') B(All my beliefs from source O are *reliably formed*.)

The inductive inference from (1)-(4) to (5'') is Vogel's version of bootstrapping. One might claim that sensitivity accounts of knowledge deliver the expected result for the bootstrapping

case that Vogel originally presented, since these beliefs are insensitive and, therefore, do not constitute knowledge. However, sensitivity accounts fail insofar as they deliver opposite results for very similar cases, as we will see.

4.2 Sensitive Moorean reasoning

Notably, there are also sensitive beliefs formed via bootstrapping or Moorean reasoning. Take the following inference that is a repetition of the inference leading to sensitive higher-level beliefs $B(c)$.

- (1) $B(p_1)$
- (2) $B(\text{I believe that } p_1)$
- (3) $B(\text{My belief that } p_1 \text{ is true})$
- (4) Repeat for $p_2 \dots p_n$
- (5) $B((\text{My belief that } p_1 \text{ is true}) \dots (\text{My belief that } p_n \text{ is true}))$

Belief (5) has the following formal structure:

- $B((B(p_1) \wedge p_1) \wedge \dots \wedge (B(p_n) \wedge p_n))$

In short:

- $B(c_1 \wedge \dots \wedge c_n)$

Belief (5) is sensitive, if each of the beliefs $B(p_1) \dots B(p_n)$ is sensitive. (5) is false, if at least one of its conjuncts c_i is false. We have seen that sensitivity carries over from $B(p_i)$ to $B(c_i)$ via inference from $B(p_i)$ and $B(B(p_i))$. But if I am minimally coherent, then I do not believe a conjunction, if I do not believe each of the conjuncts. So if I do not believe c_i , then I do not believe the conjunction $B(c_1 \wedge \dots \wedge c_n)$. Therefore, in the nearest possible world, where (5) is false, I do not believe it anymore and the belief is sensitive.

Take again a concrete example of Moorean reasoning, where beliefs of type (5) are sensitive.

- (1) $B(\text{COMPUTER})$
- (2) $B(\text{I believe that COMPUTER})$
- (3) $B(\text{My belief that COMPUTER is true})$
- (4) Repeat for $\text{DESK} \dots$
- (5) $B((\text{My belief that COMPUTER is true}), (\text{My belief that DESK is true}) \dots)$

Plausibly, a well-functioning perceptual apparatus delivers sensitive beliefs. Thus, our perceptual beliefs are typically sensitive and, therefore, our beliefs about these beliefs are also sensitive.

The types of beliefs that are involved in this inference are about the world, about what we believe about the world, and about the truth of our beliefs about the world. Together, they provide a comprehensive picture about us in the world, notably one, which can be sensitive throughout.

If the principle of bi-valence is true, the the proposition *my belief that p is true* is equivalent to *my belief that p is not false*, with the formal structure $B(p) \wedge \neg\neg p$. Accordingly, beliefs of the following structure are also sensitive when acquired via steps (1)-(4):

(5') $B((\text{My belief that } p_1 \text{ is not false) and...and (My belief that } p_n \text{ is not false)})$

(5'') $B((\text{COMPUTER and I am not a BIV deceived in falsely believing COMPUTER) and...})$

(5'') is remarkably close to a direct denial of the skeptical hypotheses of the form (I am not a BIV deceived in falsely believing that COMPUTER...), which is, as noted, insensitive.

4.3 Summary: The heterogeneity problem

The following table summarizes sensitive and insensitive higher-level beliefs and sensitive and insensitive beliefs acquired via bootstrapping and Moorean reasoning:

Insensitive beliefs:

$B(d)$

$B(\text{I am not falsely believing that } p)$

$B(\text{I am not a BIV deceived in falsely believing that } p)$

$B((\text{I do not falsely believe that } p_1) \text{ and... and (I do not falsely believe that } p_n))$

$B(\text{All my beliefs delivered by source O are true})$

Sensitive beliefs:

$B(c)$

$B(\text{My belief that } p \text{ is true})/B(\text{My belief that } p \text{ is not false})$

$B(p \text{ and I am not a BIV deceived in falsely believing that } p)$

$B((\text{My belief that } p_1 \text{ is not false) and...and (My belief that } p_n \text{ is not false)})$

These insensitive and sensitive beliefs are closely related. Moreover, some of the sensitively believed propositions are stronger than their insensitively believed counterparts. This outcome is a mess. Thus, we see how the heterogeneity problem for higher-level beliefs also affects bootstrapping and Moorean reasoning.

It is worth noting that I do not assume that all cases of higher-level beliefs and beliefs based on bootstrapping need to be treated the same in terms of justification or alleged justification. Rather, I am open to both accounts: first, that we need to treat them equally, and, second, that we do not need to. The point is that in the second case sensitivity does not offer the right criterion for distinguishing the good cases from the bad ones. Intuitively, instances of simple higher-level beliefs like $B(d)$ or $B(c)$ seem less problematic than instances of bootstrapping from $B(p)$ to $B(\text{I am not a BIV deceived in falsely believing that } p)$ or to $B(p \wedge \text{I am not a BIV deceived in falsely believing that } p)$.¹⁷ Thus, if we allow treating different instances of bootstrapping differently, then we want an account that is in accordance with these intuitions, but sensitivity accounts are not. They allow some complex instances of bootstrapping to constitute knowledge, whereas they prohibit some simple instances from constituting knowledge.

5 Why sensitivity accounts of knowledge can neither solve nor explain the skeptical puzzle

Bootstrapping and Moorean reasoning are intuitively flawed processes for knowledge acquisition. One strategy for explaining why they are flawed or why they seem flawed relies on the idea that they always deliver a positive result, regardless of whether the source in question is reliable or accurate. One way to spell out this intuition is in terms of sensitivity, i.e. by arguing that bootstrapping and Moorean reasoning are flawed or seem to be flawed, because they lead to insensitive beliefs. However, some instances of bootstrapping and Moorean reasoning lead to sensitive beliefs, while others do not. Therefore, this strategy must fail.

¹⁷ Critics of sensitivity accounts and their defenders share this view. Sosa (1999), for example, uses the simple inference from $B(p)$ to $B(d)$ as an illustrative counter-example against sensitivity accounts, since knowledge should transmit in these simple cases. DeRose (1995) who defends a sensitivity based contextualism assumes that knowing a “full-fledged” skeptical hypothesis is more demanding than knowing simple higher-level propositions like c or d .

I will, first, present a general classification of accounts that utilize sensitivity for explaining the (alleged) defectiveness of bootstrapping and Moorean reasoning. Second, I will discuss exemplarily in more detail how the heterogeneity problem affects the sensitivity accounts of Nozick, DeRose, Black and Roush.

We can choose between two general types of explanations for our persisting intuition that bootstrapping and Moorean reasoning are flawed:

- (1) Bootstrapping and Moorean reasoning *are* flawed.
- (2) Bootstrapping and Moorean reasoning are flawless, but they *seem* flawed.

Accordingly, we can utilize sensitivity for explaining our intuition about bootstrapping and Moorean reasoning in two different ways.

- (1) Bootstrapping and Moorean reasoning *are* flawed because they lead to insensitive beliefs.
- (2) Bootstrapping and Moorean reasoning are flawless, but ‘sensitivity’ is part of an explanation for why they *seem* flawed.

In case (1), we can further distinguish between two views, first, that bootstrapping and Moorean reasoning are flawed because sensitivity is necessary for *any kind* of knowledge and, second, that it is necessary for knowledge that the skeptical hypotheses are false.

Those who defend the first view, as Nozick does, are confronted with numerous counterexamples of insensitive knowledge outside the skeptical context. However, the results obtained here show that the weaker claim is also not defensible, since some instances of bootstrapping and Moorean reasoning lead to *sensitive* beliefs.¹⁸

If we choose option (2), we can find two strategies for explaining why bootstrapping and Moorean reasoning only seem to be flawed in the literature.

- (2.1) Bootstrapping and Moorean reasoning are flawless, but they seem flawed because they lead to insensitive beliefs. However, we mistakenly believe that sensitivity is necessary for (any kind of) knowledge.
- (2.2) Bootstrapping and Moorean reasoning are flawless, but they seem flawed because we mistakenly think that they lead to insensitive beliefs.

According to (2.1) we err in believing that sensitivity is necessary for knowledge. According to (2.2) we are mistaken about the insensitivity of bootstrapping and Moorean reasoning.

¹⁸ Of course, we could narrow the definitions of bootstrapping and Moorean reasoning by adding as a further condition that the resulting beliefs are insensitive. However, this definition would no longer be in accordance with our intuitive understanding of bootstrapping and Moorean reasoning.

If we opt for (2.1), then we can choose between different explanations for why we mistakenly believe that sensitivity is necessary for knowledge. Here are two possibilities.

(2.1.1) We confuse two kinds of knowledge, one that demands sensitivity and one that does not.¹⁹

(2.1.2) We mistakenly believe that sensitivity is necessary for knowledge. Period.

Most of these alternatives have been defended in the literature. (2.1.1) is featured in DeRose's (1995 and 2010) sensitivity-based contextualism. (2.1.2) seems ad hoc and, to the best of my knowledge, it is not defended in literature. (2.2) is defended by Black (2002 and 2008a).

In what follows, I will investigate the sensitivity accounts of Nozick, DeRose, Black and Roush in more detail. We will see that they have to accept either the heterogeneity problem as outlined here or a variant of it. Therefore, none of them can utilize sensitivity for explaining why bootstrapping and Moorean reasoning are flawed or why they seem flawed.

5.1 Nozick's sensitivity account

Nozick argues that sensitivity is a necessary condition for knowledge and that knowledge is not closed under known entailment. The *simple* Nozickian view is that any anti-skeptical hypothesis is insensitively believed and therefore not known. This view is mistaken. If sensitivity is necessary and – in addition to Nozick's conditions (1), (2) and (4) – sufficient for knowledge, then there are some anti-skeptical hypotheses which we know and some which we do not know. Nozick (1981) was well aware of the fact that his account implies that we know the conjunction $\neg sh \wedge p$, though we do not know $\neg sh$. Thus, Nozick was already explicitly holding a *differentiated* view with respect to this conjunction, a fact that is often neglected in literature.

According to this differentiated view, skepticism is a more isolated phenomenon than according to the simple view, since it is only some specific anti-skeptical hypotheses that we cannot know. This can be seen as an advantage of the differentiated view. However, the simple view has it that knowledge closure does not hold between external world propositions and any kind of anti-skeptical hypotheses, which is already highly controversial. According to the differentiated view, knowledge closure fails between strong and weak anti-skeptical

¹⁹ One can also argue that we confuse knowledge with something else, the first does not demand sensitivity, but the second does. This view can be supported by an account suggested by Weatherson (manuscript) that knowledge does not require sensitivity, but that *testing* a measuring device does.

hypotheses. This is an even more problematic instance of closure failure. Astonishingly, Nozick is willing to bite this bullet.

We can find a more detailed explanation for Nozick's conception of closure-failure, if we consider his often ignored account of *inferential* knowledge. Nozick proposes:

S knows (via inference from p) that q iff

- (1) S knows that p
- (2) q is true, and S infers q from p (thereby, being led to believe that q)
- (3) if q were false, S wouldn't believe that p (or S wouldn't infer q from p)
- (4) if q were true, S would believe that p (and would infer q from p if she were to infer either q or $\neg q$ from it).²⁰

Nozick argues that I cannot acquire knowledge that $\neg sh$ from my external world belief that p , because condition (3) is violated. If I were a BIV deceived in falsely believing that p , I would still believe that p and would infer from p that $\neg sh$. Hence, Nozick not only claims that one can know that p without knowing that $\neg sh$, he also provides an explanation for why we do not have this *inferential* knowledge via Moorean reasoning.

Notably, the results concerning insensitive and sensitive beliefs remain the same, if we consider Nozick's concept of inferential knowledge. Take the example of $B(d)$ and $B(c)$. Beliefs of type $B(d)$ acquired via inference from $B(p)$ cannot constitute inferential knowledge, because if d were false, S would still believe that p . However, beliefs of type $B(c)$ acquired via bootstrapping are sensitive if $B(p)$ is sensitive because if c were false, S would not believe that p .

A particularly counter-intuitive result of Nozick's account of inferential knowledge is that knowledge does not transmit via conjunction elimination. Thus, we can know the conjunction $p \wedge \neg sh$ without knowing the conjunct $\neg sh$, even if we form the belief in the conjunct via inference from the conjunction. However, the view that knowledge does not transmit via conjunction elimination seems very counter-intuitive.²¹ Thus, given Nozick's account of inferential knowledge, the heterogeneity problem remains the same. Nozick is not ignorant

²⁰ For a detailed discussion of Nozick's account of inferential knowledge and for this formulation see Baumann (2012).

²¹ See, for example, Hawthorne (2005) who regards knowledge by conjunction elimination as "incredibly plausible".

about the consequences of his account. Rather, he offers an additional explanation for why this is so. However, his explanations are highly implausible.²²

5.2 DeRose's indirect sensitivity account

DeRose's (1995 and 2010) *indirect* sensitivity account has it that the insensitivity of certain beliefs *explains* why they fail to be knowledge. DeRose (2010, 163) argues that we have "some at least fairly general—though perhaps not exceptionless—tendency to judge that insensitive beliefs are not knowledge". Accordingly, the insensitivity of our beliefs in anti-skeptical hypotheses can explain our intuitions that these beliefs fail to constitute knowledge, even if sensitivity is not a necessary condition for knowledge in general. Moreover, DeRose (1995) rejects conjunctions of the form "S knows that she has hands but S does not know that she is not a handless BIV". His sensitivity based contextualist solution to the skeptical problem takes the following form: In ordinary contexts, persons have external world knowledge and knowledge that the skeptical hypotheses are false. However, when it is *asserted* that S knows that p (or that she does not know that p), the standards for knowledge tend to be raised up to a level that S's belief has to be sensitive for constituting knowledge. By asserting that S knows that she is not a BIV we establish such high standards and the

²² Nozick's account of knowing via a method faces a technical problem when it comes to *one-sided methods* that can recommend believing that p , but cannot recommend believing that $\neg p$. Intuitively, we want to allow that one-sided methods can generate knowledge, but according to Luper-Foy (1984), they necessarily violate the sensitivity condition and therefore, cannot generate knowledge. Luper-Foy paraphrases Nozick's sensitivity condition as follows:

(3) $(\neg p \wedge (S \text{ believes that } p \text{ via } M \text{ or } S \text{ believes that } \neg p \text{ via } M)) \rightarrow (S \text{ does not believe that } p \text{ via } M)$

He argues that if M is one-sided and never recommends believing that $\neg p$, then the antecedent of (3) requires that S believes that p via M . But in this case the consequent of (3) and therefore, (3) turns out to be false. Hence, if M is a one-sided method, then Nozick's sensitivity condition can never be met and S cannot know via M , which is counter-intuitive. In order to solve the problem of one-sided methods, Luper-Foy (1984) suggests replacing Nozick's condition (3) with the following condition.

(3*) $\neg p \rightarrow (S \text{ does not believe that } p \text{ via } M)$

Notably, the results remain the same if we apply this modified sensitivity account to bootstrapping and Moorean reasoning. Hence, the heterogeneity problem remains stable over variations of sensitivity accounts. For this reformulation of sensitivity, see also Williamson (2000, 154). Luper-Foy and Williamson raise problems for this modified version of sensitivity and finally reject it. For a critical discussion of this modified sensitivity account, see also Zalabardo (2012).

assertion turns out to be false. In these more demanding contexts, a person has neither external world knowledge nor knowledge that she is not a BIV.

Let's take belief-forming method into account, as sensitivity accounts suggest. We have seen that some anti-skeptical beliefs formed via Moorean reasoning are insensitive, as DeRose assumes, but that some are *sensitive*. However, Nozick's claim that we know the sensitively believed anti-skeptical hypotheses, but do not know the insensitively believed anti-skeptical hypotheses, is not in line with our intuitions about skepticism. In contrast, we tend either to accept that we know insensitively believed anti-skeptical hypotheses *and* sensitively believed ones or *neither*. Moreover, we maintain these intuitions even if we become aware of the fact that some beliefs of anti-skeptical hypotheses are sensitive and that some are insensitive. Hence, our intuitions about whether we know anti-skeptical hypotheses cannot be based on the sensitivity or insensitivity of our beliefs of anti-skeptical hypotheses.

This is a serious problem for DeRose, who thinks that his contextualist account also contains an explanation for why we raise the standards for knowledge. However, his sensitivity-based explanation is simply inadequate. DeRose can still defend a contextualist solution to the skeptical problem, but without being able to explain in terms of sensitivity why we shift the standards. Hence, he is mistaken that his indirect sensitivity account is 'on the right track'.²³

5.3 Black's Moorean invariantism

Black (2002 and 2008) defends a Moorean *invariantist* response to the skeptical argument according to which the standards for knowledge are always comparatively low and, therefore, we always have external world knowledge and knowledge that we are not BIVs. Black criticizes Nozick's (1981, 184) specification of belief forming methods that any method that is experientially the same, the same "from the inside", counts as the same method. According to Nozick's account, a normal person and a BIV share the same method. Black objects that our method of coming to believe that we are not BIVs is partly perception, but that perception is in part characterized "from the outside", i.e. as a method based on using our sense apparatus. Thus, normal persons and BIVs do not share the same method, since BIVs, *ex hypothesi*, do not use any sense apparatus.²⁴

²³ For a more detailed criticism of DeRose's indirect sensitivity account that does not focus on bootstrapping and Moorean reasoning see Melchior (2014b).

²⁴ For a discussion of how to individuate belief-forming methods correctly in the context of sensitivity accounts see Becker (2012).

Black argues that given this reformulation of belief-forming methods, our beliefs that we are not BIVs turn out to be *sensitive*. S's belief via method M that p is sensitive iff in the nearest possible world where p is false and S uses M, S does not come to believe that p via M. Since BIVs do not share our method of perception, there simply is no nearest possible world where I am a BIV and where I use the same method that I use in the actual world. Hence, my belief that I am not a BIV is sensitive. Since Black defends a sensitivity account of knowledge, he concludes that we can know that we are not BIVs.

One advantage of Black's account is that beliefs in the conjunction $\neg sh \wedge p$ as well as beliefs in the conjunct $\neg sh$ turn out to be sensitive. Thus, Black offers a uniform anti-skeptical solution that allows knowing both. In this respect, he is better off than Nozick and DeRose, who both have to treat beliefs in $\neg sh$ and beliefs in $\neg sh \wedge p$ differently.

Black (2002) concedes that his solution to the skeptical problem only applies to skeptical hypotheses that rely on subjects like BIVs whose belief forming methods are essentially different from our own. However, for many of the discussed cases there is a nearest possible world where the proposition in question is false but where the same belief forming method is still available. Take the inference from $B(p)$ to $B(d)$ and suppose that p is the content of a perception based belief like 'Mark is on the other side of the street'. In this case, there are possible worlds w , where d is false i.e. where I falsely believe that Mark is on the other side of the street, and where I use the same method of perception and inference as in the actual world. In w , I take, for example, someone else to be Mark. Hence, my belief $B(d)$ is still insensitive. The same is true for the following beliefs formed via bootstrapping:

B(All my beliefs formed by looking at the gas gauge are true.)

B(All my beliefs formed by looking at the gas gauge are reliably formed.)

In both cases, we could still use the same method 'from the outside' of consulting the gas gauge plus bootstrapping, if these conclusions were false. Therefore, these beliefs still turn out to be insensitive.

If we apply Black's sensitivity account to all instances of bootstrapping, then we know that we are not BIVs via Moorean reasoning without knowing d and without acquiring knowledge via alternative instances of bootstrapping. This seems counter-intuitive. Thus, Black's

sensitivity account might offer a partial solution to the skeptical problem, as he argues, but the heterogeneity problem is only displaced.²⁵

5.4 Roush's recursive sensitivity account

Roush (2005) offers a tracking account of knowledge that is inspired by Nozick's account, but that avoids the undesired consequence of closure-failure. Roush (2005, 43) argues, as a first approximation, that S knows that p if and only if S's belief fulfills Nozick's condition (1)-(4) or " p is true, S believes p , and there is a q not equivalent to p such that q implies p , S knows that q implies p , and S knows that q ". According to Roush's account, S knows that she is not a BIV deceived in falsely believing that p , if S knows that p and knows that p implies that she is not a BIV deceived in falsely believing that p .

How does Roush explain our *intuition* that we do not know that we are not BIVs? Roush (2005, 54f) argues that the fact that we cannot find conclusive reasons for believing that we know, even if our beliefs constitute knowledge, explains our experience of helplessness when confronted with the skeptical hypothesis and our frustration about our inability to refute the skeptic. Since sensitivity is not part of an explanation for why Moorean reasoning *seems* flawed according to Roush, her view is partly in accordance with the main theses of this paper.

However, Roush's account faces two problems concerning bootstrapping. Roush (2005, 120) claims that bootstrapping obviously does not constitute knowledge. She claims that her recursive sensitivity account can explain why: First, the beliefs in the conclusion are insensitive and, second, the last step of inference of bootstrapping is not deductive.

However, we have seen that we can construe instances of bootstrapping that violate each of these conditions. First, there are beliefs based on bootstrapping that are *sensitive*. Second, we can also construe *deductive* instances of bootstrapping about the reliability of a source O. Hence, Roush cannot argue that we fail to know in these two cases because the resulting belief is insensitive and not all inferences involved are deductive. Without providing further arguments, she is committed to accepting that we do know in these cases of bootstrapping, although we do not know in Vogel's original example. Again, the heterogeneity problem is only displaced.

²⁵ Black (2008b) offers a solution to the problem of easy knowledge that is based on a particular understanding of the principle of knowledge closure.

6 Conclusion

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 sensitivity. Thus, sensitivity accounts of knowledge cannot adequately capture higher-level knowledge. Moreover, this problem carries over to bootstrapping and Moorean reasoning. Some instances of bootstrapping and Moorean reasoning deliver insensitive beliefs, but some closely related processes deliver sensitive beliefs. Therefore, neither Nozick's sensitivity account nor one of the modified sensitivity accounts of knowledge can explain why bootstrapping and Moorean reasoning are flawed or why they seem flawed. Moreover, I do not see how *any* sensitivity-based account of knowledge could offer such an explanation, given the heterogeneity problem.

Acknowledgements

The research was funded by the Austrian Science Fund (FWF): J 3174-G15. I am indebted to Ernest Sosa and the members of his epistemology group at Rutgers University for their helpful comments, to Martina Fürst, Lisa Miracchi and Danilo Šuster for their remarks on earlier drafts of this paper and to Peter Klein for insightful and encouraging discussions.

Guido Melchior is currently Assistant Professor of Philosophy at the University of Graz. After receiving his PhD at the University of Graz, he was Visiting Scholar at the Department of Philosophy at Rutgers University. His primary research interest is in epistemology, including modal knowledge accounts and skepticism.

References

- Adams, F. and Clarke, M. (2005). Resurrecting the tracking theories. *Australasian Journal of Philosophy* 83(2), 207-221.
- Baumann, P. (2009). Was Moore a Moorean? On Moore and scepticism. *European Journal of Philosophy* 17(2), 181-200.
- Baumann, P. (2012). Nozick's defense of closure. *The Sensitivity Principle in Epistemology*. K. Becker and T. Black (eds.) Cambridge: Cambridge University Press, 11-27.
- 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.
- Becker, K. (2012). Methods and how to individuate them. *The Sensitivity Principle in Epistemology*. K. Becker and T. Black (eds.) Cambridge: Cambridge University Press, 81-97.
- Becker, K. and Black, T. (2012). *The Sensitivity Principle in Epistemology*. Cambridge: Cambridge University Press.
- Black, T. (2002). A Moorean response to brain-in-a-vat skepticism. *Australasian Journal of Philosophy* 80(2), 148-163.
- Black, T. (2008). Defending a sensitive neo-Moorean invariantism. *New Waves in Epistemology*. V. F. Hendricks and D. Pritchard (eds.) Basingstoke: Palgrave Macmillan, 8-27.
- Black, T. (2008a). Solving the problem of easy knowledge. *Philosophical Quarterly* 58(233), 597-617.
- Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research* 65(2), 309-329.
- Cohen, S. (2005). Why basic knowledge is easy knowledge. *Philosophy and Phenomenological Research* 70(2), 417-430.
- DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review* 104(1), 1-52.
- DeRose, K. (2010). Insensitivity is back, baby! *Philosophical Perspectives*, 24, 161-187.
- Douven, I. and Kelp, C. (2013). Proper bootstrapping. *Synthese* 190(1), 171-185.
- Hawthorne, J. (2005). The case for closure. *Contemporary Debates in Epistemology*, Mathias Steup and Ernest Sosa, (eds.), 26-42.
- 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.
- Melchior, G. (2014a). A generality problem for bootstrapping and sensitivity. *Croatian Journal of Philosophy* 14(40), 31-47.
- Melchior, G. (2014b). Skepticism: the hard problem for indirect sensitivity accounts. *Erkenntnis* 79(1), 45-54.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge: Cambridge University Press.
- Pryor, J. (2004). What's wrong with Moore's argument?. *Philosophical Issues* 14(1), 349-378.
- Roush, S. (2005). *Tracking Truth. Knowledge, Evidence, and Science*. Oxford: Oxford University Press.
- 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.
- Titelbaum, M. (2010). Tell me you love me: bootstrapping, externalism, and no-lose epistemology. *Philosophical Studies* 149(1), 119-134.
- 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.
- Weatherson, B. (manuscript). Easy knowledge and other epistemic virtues. <http://brian.weatherson.org/ekoev.pdf>
- Weisberg, J. (2012). The bootstrapping problem. *Philosophy Compass* 7/9, 597-610.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wright, C. (2002). (Anti-)sceptics simple and subtle: Moore and McDowell. *Philosophy and Phenomenological Research*, 65(2), 330-48.
- Wright, C. (2004). Warrant for nothing (and foundations for free)?. *Aristotelian Society Supplementary Volume* 78(1), 167-212.
- Zalabardo, J. L. (2012). *Scepticism and Reliable Belief*. Oxford: Oxford University Press.