

Easy knowledge, closure failure or skepticism: a trilemma

Guido Melchior

[Penultimate draft]

Published in *Metaphilosophy*

Abstract This paper aims to provide a structural analysis of the problems related to the easy knowledge problem. The easy knowledge problem is well known. If we accept that we can have basic knowledge via a source without having any prior knowledge about the reliability or accuracy of this source, then we can acquire knowledge about the reliability or accuracy of this source too easily via information delivered by this source. However, rejecting any kind of basic knowledge leads into an infinite regress and, plausibly, to skepticism. I will argue that the third alternative, accepting basic knowledge but rejecting easy knowledge entails closure failure. This is obviously the case for deductive bootstrapping, but, notably, the problem also arises for inductive bootstrapping. Hence, the set of problems related to the easy knowledge problem has the structure of a trilemma. We are forced to accept easy knowledge, closure failure, or skepticism.

Keywords: Bootstrapping, easy knowledge problem, knowledge closure, Mooreanism, skepticism

Introduction

Typical externalist knowledge accounts and many internalist knowledge accounts accept the view that basic knowledge is possible, i.e. that a subject S can acquire knowledge via a source without having prior knowledge about the reliability of the source. Basic knowledge immediately leads to the easy knowledge problem viz. that S can acquire knowledge about the reliability or accuracy of a source via deductive and inductive reasoning from information delivered by this source. However, these reasoning processes of bootstrapping are counterintuitive methods for acquiring knowledge. In this paper, I explore the options concerning the easy knowledge problem in a systematic way. In section 1, I will review the concepts of bootstrapping and easy knowledge. In section 2, I will distinguish different instances of bootstrapping and show that if we accept basic knowledge, then for each of these

instances we face the dilemma between accepting easy knowledge or closure failure. In section 3, I will show that rejecting basic knowledge entails skepticism.

1 Bootstrapping and easy 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 it. He 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)

Vogel assumes that bootstrapping is obviously an epistemically flawed reasoning process. He uses bootstrapping cases as *reductio* arguments against process reliabilism. Since process reliabilism sanctions every step of bootstrapping, process reliabilism has to be false according to Vogel.

Cohen (2002) argues more generally that problems of bootstrapping and easy knowledge not only arise for process reliabilism but for any account of knowledge that allows one to have knowledge via a source without having prior knowledge about the reliability of this source. Cohen introduces the following principle KR.

KR: A potential knowledge source K can yield knowledge for S, only if S knows K is reliable.
(Cohen 2002, 309)

Cohen (2002 and 2005) argues that any theory that denies KR implies the undesired result that one can acquire knowledge about the reliability of our knowledge sources much too easily by basing it on basic knowledge. Cohen (2002) calls this "The Problem of Easy Knowledge". However, such suspicious inferences are not only restricted to reasoning about the *reliability* of a source. Cohen presents the following inference:

- (1) The table looks red
- (2) The table is red

- (3) If the table is red, then it is not white with red lights shining on it
(4) The table is not white with red lights shining on it.
(Cohen 2005, 418)

Cohen argues that such inferences from (1)-(3) cannot constitute knowledge of (4). The notions of basic knowledge and easy knowledge are based on Cohen's principle KR. Accordingly, basic knowledge is knowledge via a source without having any prior knowledge about the reliability of this source and easy knowledge is knowledge about the reliability of a source via inferences from output delivered by this source. The notion of bootstrapping, in contrast, is rather informally introduced by Vogel as Roxanne's reasoning about her gas gauge and processes of reasoning that are sufficiently similar to this case, i.e. processes of drawing conclusions about the reliability of a source via inference from output delivered by this source.

In this paper, I will elaborate connections between knowledge about the *accuracy* and knowledge about the *reliability* of a potential knowledge source. For this reason, I define bootstrapping more broadly than usually suggested as inferences that also include propositions about the accuracy of an output-delivering source. By slightly diverging from the common terminology, I define 'basic knowledge', 'bootstrapping' and 'easy knowledge' as follows:

Basic knowledge

S has basic knowledge that p via source O, iff S knows that p via O without having prior knowledge about the reliability or accuracy of O.

Bootstrapping

Bootstrapping is any process of reasoning about the reliability or accuracy of a potential knowledge-source via inferences (at least partly) from output delivered by this very source.

Easy knowledge

S has easy knowledge that p iff S has knowledge via bootstrapping that p.

2 Instances of bootstrapping

Cohen's example of reasoning about the color of a table is a *deductive* inference; Vogel's bootstrapping case is an *inductive* inference from multiple premises. Generally speaking, we can distinguish inductive and deductive instances of bootstrapping and single-premise and multi-premise bootstrapping. In this section, I will present single-premise deductive bootstrapping, multi-premise deductive bootstrapping, and multi-premise inductive

bootstrapping in more detail. Furthermore, I will argue that if we accept basic knowledge, then in each of these cases we are facing the dilemma of being forced to accept easy knowledge or to reject knowledge-closure. This is obvious for deductive bootstrapping, but notably this problem also arises for *inductive* bootstrapping. In order to make this point, I will first discuss the principle of knowledge-closure in more detail.

2.1 Knowledge-closure

Knowledge-closure or *closure* is the principle that a person, who knows that p and knows that p entails q , also knows that q . Hence, knowing that p and knowing that p entails q is sufficient for knowing that q . If closure holds, then knowledge is said to be closed under known entailment.

This *orthodox* version of knowledge-closure is often regarded as too strong and, thus, many philosophers have developed alternative accounts which try to replace it by a weaker and more plausible principle. There are two general strategies of weakening closure: First, by strengthening its antecedent and, second, by weakening its consequent.¹ The first strategy is more popular than the second, perhaps because knowledge is the central epistemic notion and thus there is more interest in the sufficient conditions for knowledge than in the sufficient conditions for weaker notions. Since I focus on knowledge in this paper, I will only consider variants of closure with strengthened antecedents.

Here is one problem for the orthodox version of knowledge-closure. The consequent of closure is $K(q)$. Knowing that q implies believing q , but neither knowing that p nor knowing that p entails q nor knowing both propositions implies believing q , since it is possible that a person S knows that p and knows that p entails q but fails to believe q , because S simply fails to put 2 and 2 together. Considering the belief problem, one can weaken closure by assuming that if S knows that p and knows that p entails q and believes q , then S knows that q . Let's call this modified version of closure *proper belief closure*.

A second problem for knowledge closure arises as follows: One might have the intuition that a person who believes q for implausible reasons fails to know that q , even if she knows that p and knows that p entails q .² This objection can be met by incorporating a condition of belief *acquisition* into the antecedent of closure. Williamson (2000) claims that intuitive closure is the principle that knowing p_1, \dots, p_n , competently deducing q and thereby coming to believe q

¹ See David and Warfield (2008).

² David and Warfield (2008) call the first problem the *belief problem* and the second problem the *warrant problem*. For a brief discussion of the belief problem, see also Blome-Tillmann (2006).

is, in general, a way of coming to know that q . I will call versions of closure that incorporate an acquisition condition *proper basing closure*. This version of closure states that S knows that q , if S knows that p and knows that p entails q and believes q based on deducing it from p and ‘ p entails q ’.³

Knowledge-closure via inference from a single premise is usually regarded as highly plausible, but *multi-premise* closure faces additional problems. Suppose, S knows propositions p_1, \dots, p_n and knows that these propositions entail the conjunction $p_1 \wedge \dots \wedge p_n$. In this case, it is disputable that S knows the conjunction, since the justification of the conjunction can be weaker than the justification of each of the conjuncts. Hence, one can accept single-premise closure, but reject multi-premise closure. This issue will become relevant later, when we discuss single-premise bootstrapping and multi-premise bootstrapping.

2.2 Deductive single-premise bootstrapping

One type of deductive bootstrapping involves inferences with the following structure:

- (1) p (via source O)
- (2) O does not falsely deliver that p (via inference from (1))

Such inferences are instances of deductive closure, if the underlying formal structure is as follows:

- (1) p
- (2) $\sim((O \text{ delivers } p) \wedge \sim p)$

(2) can be regarded as a proposition about the accuracy of the deliverance of O . It is equivalent with the disjunction $\sim(O \text{ delivers } p) \vee p$. Since p is one of the disjuncts, it entails this disjunction. This entailment relation also holds if a potential explanation for why O falsely delivers p is added, as in the following cases:

- (2') $\sim((O \text{ delivers } p \text{ because of } x) \wedge \sim p)$
- (2'') $\sim(O \text{ falsely delivers } p \text{ because of } x)$

³ Versions of proper basing closure are most prominently defended by Hawthorne (2004 and 2005), who adds the further premise that S also has to retain her knowledge that p . David and Warfield (2008) present a further variant of proper basing closure that requires that believing q must solely be based on deduction from p and (p entails q).

In the first case, O delivers p for a certain reason, which is not necessarily a reason for why O falsely delivers p, in the second case x is also a reason for the falsity of the deliverance. In any of these two cases, the proposition is entailed by p.

Cohen's example of reasoning about the white table not being illuminated by red light can be interpreted as an instance of this kind of deductive bootstrapping. Here are two further examples:

Example 1:

- (1) The ambient temperature is 73°. (by looking at thermometer T₁)
- (2) T₁ is not falsely reporting that the ambient temperature is 73°. (via inference from (1))
- (2') T₁ is not manipulated by the landlord to (falsely) report that the ambient temperature is 73°. (via inference from (1))

Example 2:

- (1) There is a computer in front of me (by having an experience as of a computer in front of me)
- (2) I am not falsely experiencing as of a computer in front of me. (via inference from (1))
- (2') I am not a brain in a vat (BIV) deceived in falsely experiencing as of a computer in front of me. (via inference from (1))⁴

If we accept the possibility of basic knowledge of (1), then we can either accept or reject the possibility that these instances of bootstrapping lead to easy knowledge. If we reject the possibility of easy knowledge of (2) or (2') via inference from (1), then we are committed to accepting closure failure for the following reason: p entails $\sim((O \text{ delivers } p) \wedge \sim p)$ and S can easily know this entailment relation. However, it is possible that S does not know $\sim((O \text{ delivers } p) \wedge \sim p)$ via any other belief-forming source than O. Hence, if basic knowledge is possible but easy knowledge is not, the following conjunction can turn out to be true:

- $K(p)$ and $K(p \text{ entails } \sim((O \text{ delivers } p) \wedge \sim p))$ and $\sim K\sim((O \text{ delivers } p) \wedge \sim p)$

This conjunction contradicts the orthodox formulation of knowledge-closure. Moreover, rejecting easy knowledge also implies that S still does not know that $\sim((O \text{ delivers } p) \wedge \sim p)$, even if S also *believes* $\sim((O \text{ delivers } p) \wedge \sim p)$ or if this belief is *properly based* on S's beliefs that p and that $\sim((O \text{ delivers } p) \wedge \sim p)$. Hence, the weaker versions of proper belief closure and proper basing closure also turn out to be false. We obtain the same results if we incorporate an

⁴ One can also argue that the underlying formal structure here is the one of 2'.

explanation for why O falsely delivers p into the conclusion as in the case of $\sim((O \text{ delivers } p \text{ because of } x) \wedge \sim p)$ or $\sim(O \text{ falsely delivers } p \text{ because of } x)$.

Klein (2004, 177ff) points out that if propositions of the form ‘The table is not white but illuminated with red light’ are negations of conjunctions with the formal structure $\sim((O \text{ delivers } p \text{ because of } x) \wedge \sim p)$, then they are equivalent with disjunctions with the structure $\sim(O \text{ delivers } p \text{ because of } x) \vee p$. He argues, contra Cohen (2002), that such disjunctions do not give us any actual knowledge about the truth or accuracy of a belief delivering source. They just state that p or something else is the case but they do not make any claim about the way p is delivered. Adopting this line of argumentation, inferences from p to $\sim((O \text{ delivers } p \text{ because of } x) \wedge \sim p)$ do not lead to easy knowledge, because the conclusion only *seems* to be a proposition about the accuracy of a belief-forming source, though it is actually not.

Klein might be right if the conclusions of the discussed inferences are formulated the way he suggests. However, we can easily formulate inferences that clearly fulfill the criteria for easy knowledge. Take the following inferences from premises P₁ and P₂ to one of the following conclusions.

Premises:

P₁: p (via source O)

P₂: O delivers p (via observing O)

Conclusions:

C₁: (O delivers p) \wedge (p is true)

C₂: (O delivers p) \wedge (p is not false)

C₃: (O delivers p) \wedge (p is not false because of x)⁵

Each of the inferences from P₁ and P₂ to any of the conclusions C₁-C₃ is a deductive inference. This is obviously true for conclusions C₁-C₂. It is also true for the inference to C₃, since C₃ is weaker than C₂.

These inferences are based on two premises, P₁ and P₂. However, only premise P₁ is formed via source O. I understand single-premise bootstrapping as reasoning about the accuracy or reliability of a source O based on a single premise formed via O whereas multi-premise bootstrapping is based on several premises formed via O. Thus, these reasoning processes are still instance of *single*-premise bootstrapping. Here is an example:

⁵ One could draw a distinction between ‘p is true’ and ‘p is the case’ just as one can draw a distinction between ‘p is the case’ and ‘ $\sim p$ is not the case’. Accordingly, one can also formulate ‘(O delivers p) \wedge p is the case’ and ‘(O delivers p) \wedge $\sim p$ is not the case’ as a further conclusion. However, for the sake of simplicity, I will treat these conclusions as equivalent.

- P₁: The ambient temperature is 73°. (by looking at thermometer T₁)
- P₂: T₁ delivers that the ambient temperature is 73°. (by looking at T₁)
- C₁: T₁ delivers that the ambient temperature is 73° and this deliverance is true.
- C₂: T₁ delivers that the ambient temperature is 73° and this deliverance is not false.
- C₃: T₁ delivers that the ambient temperature is 73° and the landlord has not manipulated T₁ to generate this false information.

Similar examples can be formulated for experiential knowledge about my surroundings. If we accept basic knowledge, then we can, again, accept or reject easy knowledge via bootstrapping from P₁ and P₂ to C₁-C₃. If easy knowledge is rejected and if S does not acquire knowledge of C₁-C₃ via any other source, then S does not know C₁-C₃, even if S knows P₁ and P₂. This is also the case if S knows that P₁ and P₂ entail any of the propositions C₁-C₃. Hence, the following knowledge ascriptions can turn out to be true for P₁, P₂ and C₁.

- K(p) and K(O delivers p) and K((p and (O delivers p))) entail (O delivers p ∧ p) and ~K(O delivers p ∧ p)

The same is true for the conclusions C₂ and C₃. These knowledge ascriptions are instances of closure failure. In the example above, S knows P₁ about the ambient temperature and P₂ about the deliverance of thermometer T₁ but does not know any of the propositions C₁-C₃ about the accuracy of T₁. Moreover, these cases also violate weaker versions of knowledge-closure. They violate proper belief closure if S also *believes* that ((O delivers p) ∧ p) and proper basing closure, if S *forms* the belief that ((O delivers p) ∧ p) via inference from P₁ and P₂ and (P₁ and P₂ entail ((O delivers p) ∧ p)).

2.3 Deductive multi-premise bootstrapping

The instances of bootstrapping discussed so far concern deductive inferences from a *single* proposition p delivered by O. However, one can easily construe instances of deductive *multi*-premise bootstrapping. Take the following deductive inference:

- P₁: (p₁)
- P₂: (O delivers p₁)
- P₃: (p₁ ∧ (O delivers p₁))
- Repeat for p₂...p_n
- Conclusion: ((p₁ ∧ (O delivers p₁) ∧ ... ∧ (p_n ∧ (O delivers p_n)))

The conclusion is about the truth or accuracy of the output delivered by O via inference from multiple deliverances of O and, therefore, an instance of multi-premise bootstrapping. It is a conjunction of conclusions of type C₁. The example can easily be modified by using conclusions that are conjunctions of type C₂ and C₃. Here is a further case of deductive multi-premise bootstrapping:

P₁: (p₁)
P₂: (O delivers p₁)
P₃: (p₁ ∧ (O delivers p₁))
Repeat for p₂...p_i
P_n: p₁...p_i are all the deliverances by O
Conclusion: All deliverances by O are accurate/true/not inaccurate/not false

The instance of closure that fails in this context is multi-premise closure, i.e. the view that if a person knows multiple propositions p₁...p_n and knows that they entail p_i, then S also knows p_i. However, multi-premise closure can be disputed even if single-premise closure is accepted. In this respect, the failure of multi-premise closure is less problematic. However, these instances of deductive multi-premise bootstrapping turn into deductive single-premise bootstrapping if S concludes from the conjunction of premises P₁-P_n. Hence, these instances of closure failure are not less problematic.

If we accept basic knowledge but reject easy knowledge, then, again, orthodox knowledge-closure fails, if S knows the relevant entailment relations between premises and the conclusion but does not know the conclusion via another source than bootstrapping. More restricted versions of knowledge-closure fail if S fulfills those further conditions that these versions demand.

2.4 Inductive bootstrapping

Vogel (2002) originally introduces bootstrapping as a process of drawing *inductive* inferences from output delivered by a source to conclusions about the *reliability* of this source. However, we can also use inductive bootstrapping for reaching conclusions about the *truth* of the deliverances of O. Take the following examples:

Premises:
P₁: (p₁)
P₂: (O delivers p₁)
P₃: (p₁ ∧ (O delivers p₁))
Repeat for p₂...p_n
Conclusions:

C₁: O is a reliable output-delivering source.

C₂: Each output delivered by O is true.

The inferences from P₁-P_n to C₁ and to C₂ are inductive. If we accept basic knowledge, we can, again, choose between accepting or rejecting easy knowledge via inductive bootstrapping. In the case of deductive bootstrapping, it is obvious that rejecting easy knowledge implies closure failure. Notably, we face the same problem with respect to *inductive* bootstrapping. If we reject easy knowledge, then S cannot know whether an output-delivering source O₁ (e.g. a thermometer) is accurate (or not inaccurate, reliable, not unreliable) via inference from outputs delivered by O₁. However, S can use a second output-delivering source O₂ (a second thermometer) for coming to know whether O₁ is accurate (not inaccurate, reliable, not unreliable). But how can S acquire knowledge about the reliability or accuracy of O₁ *and* O₂? Our informal introduction to bootstrapping and easy knowledge does not make any explicit claims about when reasoning about two or more output-delivering sources is an instance of bootstrapping. However, it is plausible to assume that any inference from outputs of any of the sources in question counts as bootstrapping.

There is a *reductio* argument in order to support the view that those like Vogel who argue that bootstrapping cannot constitute knowledge are forced to accept this view about several output-delivering sources. Suppose S can acquire knowledge about the reliability of O₁ *and* O₂ via inference from an output delivered by O₁ *or* O₂. In this case Roxanne can acquire knowledge about the tank at period T₁ via inference from the output the gauge delivered at T₂ and vice versa. Roxanne can hereby acquire knowledge about the overall reliability of her tank by looking at the tank gauge, given that we perceive looking at the tank as a class of sources characterized by periods. However, this is precisely what Vogel wants to deny. Thus, he is committed to accept the concept of bootstrapping for conjoined sources O₁ *and* O₂ proposed here.

Accordingly, any inference from an output delivered by O₁ or O₂ to a conclusion about the reliability of O₁ *and* O₂ counts as bootstrapping. However, if there is no further source O₃ available to S in order to know whether O₁ and O₂ are accurate, then S cannot know whether O₁ *and* O₂ are accurate. In this case, S violates knowledge-closure as follows:

- (1) S knows that O₁ is reliable/accurate. (via O₂)
- (2) S knows the relevant entailment relations.
- (3) S does not know that O₁ *and* O₂ are not (equally) unreliable/inaccurate.

(1) entails (3) for the following reason: (1) has the formal structure $K(O_1 \text{ is } r/a)$ and (3) has the formal structure $\sim K\sim((O_1 \text{ is } \sim r/\sim a \text{ and } O_2 \text{ is } \sim r/\sim a))$. Thus, (1) is an instance of $K(p)$ and (3) is an instance $\sim K\sim(\sim p \wedge \sim q)$ which is equivalent to $\sim K(p \vee q)$. Furthermore, S knows (1) and knows that (1) entails (3) but if easy knowledge is rejected, then S still does not know (3). This is a clear instance of closure failure. In this respect, inductive bootstrapping also leads to closure failure.⁶ Take the following example:

S has two thermometers T_1 and T_2 for determining the ambient temperature.
 S knows that T_1 is accurate by using T_2 .
 S knows the relevant entailment relations.
 S does not know that T_1 and T_2 are not both (equally) inaccurate.

If there are any doubts about the conditions for bootstrapping for two or more sources, then we can define a source O_n as the source that conjoins O_1 and O_2 . Suppose that if O_1 is reliable or accurate, then O_n is not totally unreliable or totally inaccurate. In this case knowledge-closure is violated as follows:

(1') S knows that O_1 is reliable/accurate.
 (2') S knows the relevant entailment relations.
 (3') S does not know that O_n is not totally unreliable/inaccurate.

Moreover, given certain assumptions we can construe a further type of closure violation. Suppose that S knows that O_1 is reliable or accurate by using the output delivered by O_2 . Can S know that O_2 is reliable or accurate by using the output delivered by O_1 ? There might be something fishy about this procedure but I will not enter the discussion on whether this is a valid method of knowledge acquisition. However, *if* it is accepted, then it leads to a further type of closure violation. Suppose that S knows that O_1 is reliable or accurate by using O_2 and that O_2 is reliable or accurate by using O_1 . Furthermore, suppose that there is no further source O_3 available to S for determining whether O_1 and O_2 are reliable or accurate. In this case, knowledge-closure is violated as follows:

(1) S knows that O_1 is reliable/accurate. (via O_2)
 (2) S knows that O_2 is reliable/accurate. (via O_1)
 (3) S knows the relevant entailment relations.

⁶ Closure-failure does not follow from inductive bootstrapping itself, but it follows directly from the knowledge that we can base on inductive bootstrapping and the knowledge that we cannot.

(4) S does not know that O_1 and O_2 are reliable/accurate.⁷

To sum up: Accepting basic knowledge and rejecting easy knowledge obviously implies closure-failure for *deductive* bootstrapping but it also implies closure failure for *inductive* bootstrapping.

2.5 Solution types

If we accept basic knowledge, then we face the dilemma of being forced to accept easy knowledge or closure failure. Here we can distinguish at least four solution types.

Type-1 solutions

Reject all instances of easy knowledge.

Consequence: All bootstrapping-based instances of closure failure must be accepted.

Type-2 solutions

Accept all instances of easy knowledge.

Consequence: No instance of closure failure must be accepted.

Type-3 solutions

Accept some instances of easy knowledge and reject other instances.

Consequence: Some instances of closure failure must be accepted and others not.

Type-4 solutions

Reject easy knowledge for some *forms* of knowledge but accept it for other forms.

Consequence: Closure-failure must be accepted for some forms of knowledge, but not for others.

One can also formulate mixed versions of type 3 and type 4 by holding that *some* forms of basic knowledge allow for *some* instances of easy knowledge.⁸ Most of these solutions have been proposed in the literature in one way or another.

Type-1 solutions

One can support type-1 solutions by providing a further *explanation* for why we mistakenly regard knowledge-closure as a valid principle. However, I defined bootstrapping in a very

⁷ Again, we can also define O_n as the source that delivers the output from O_1 or O_2 and construe instances of closure violation by replacing (4) by (4'): S does not know that O_n is reliable/accurate.

⁸ Type 4 solutions are not meant to be a subset of type 3 solutions. Type-3 solutions distinguish between different types of bootstrapping but only take one kind of knowledge into account. Type-4 solutions, in contrast, consider different kinds of knowledge.

broad sense that also includes simple deductive inferences that are usually regarded as undisputed.⁹ I do not know of any author who rejects *all* instances of easy knowledge in the broad sense I have in mind.

Type-2 solutions

Type-2 solutions can keep the principle of knowledge-closure but for the price of accepting every instance of easy knowledge. However, some instances of bootstrapping seem highly implausible. Thus, for being satisfactory type-2 solutions not only have to defend easy knowledge in general, they also have to explain why it *seems* defective to us. Here, different explanations are possible. Markie (2005, 415) argues that bootstrapping gives evidence for the belief that a belief-forming source is reliable but it is of limited value, because it “is not, in particular, the basis for a nonquestion-begging reply to anyone who challenges the reliability.” Cohen (2005) argues, contra Markie, that the reason why bootstrapping is flawed is not that it is not cogent for someone else but that it is not cogent for *him*. This indicates that any satisfactory defense of easy knowledge has to offer a complex explanation for why bootstrapping seems defective for us that might include missing persuasiveness in some context, but other aspects in other contexts. Bergmann (2004 and 2006) chooses a similar strategy as Markie does for defending easy knowledge. He distinguishes between questioned source situations or contexts and unquestioned source situations or contexts. In the first case, the person *is or should be* seriously questioning or doubting the trustworthiness or reliability of a source. In the second case, the person neither is nor should be seriously doubting or questioning. Bergmann argues that bootstrapping is malignant in the first case but benign in the second. Moreover, Bergmann utilizes questioned source contexts for explaining our intuition that bootstrapping is malignant.

Type-3 solutions

Opting for a type-3 solution means to accept some instances of easy knowledge and to reject others. Type-3 solutions allow for the outcome that some instances of bootstrapping are benign but that others are malignant. Such an outcome seems to be in accordance with our intuitions and, in this respect, attractive. However, in order to support such a solution, one has

⁹ Sosa (1999) and Vogel (2000) argue that assuming that one can know that p without knowing that $\sim(B(p) \wedge p)$ is counter-intuitive. Therefore, any knowledge-account, such as sensitivity accounts, that prohibits this instance of easy knowledge are flawed.

to offer a plausible criterion for distinguishing benign instances of bootstrapping that lead to easy knowledge from malignant ones that do not. Here different approaches are possible.

- (a) Distinguishing different propositions
- (b) Distinguishing between direct and complex bootstrapping
- (c) Treating different belief forming sources differently¹⁰

(a) Distinguishing different propositions

One way of distinguishing different forms of bootstrapping is by means of the involved propositions. One can, for example, argue that deducing *My belief that p is true* from *p* and *I believe that p* is a way of coming to know that I truly believe that *p*, but that this is not a way of coming to know *My belief that p is not false* or at least not a way of coming to know *I am not a BIV deceived in falsely believing that p*. Although this outcome might be somewhat plausible, it seems arbitrary to distinguish between logically equivalent propositions such as *My belief that p is true* and *My belief that p is not false* or to allow knowledge of *My belief that p is true* but to reject knowledge of a weaker proposition *I am not a BIV deceived in falsely believing that p*. Vogel (2000 and 2008) seems to endorse a type-3 solution, since he rejects easy knowledge via inductive bootstrapping about the reliability of the source but explicitly allows higher-level knowledge that $\sim(B(p) \wedge \sim p)$ via deductive inference from *p*. However, there is a tension between Vogel's (1990) views that there are no convincing instances of closure failure and his view that inductive bootstrapping does not lead to knowledge, since rejecting inductive bootstrapping also leads to closure failure as we have seen.

(b) Distinguishing direct and complex bootstrapping

One could also distinguish benign from malignant bootstrapping by relying on the distinction between direct and complex bootstrapping. Accordingly, one could argue that direct inductive bootstrapping about the reliability of a source that is exclusively based on output delivered by this source is malignant and does not lead to easy knowledge but that complex bootstrapping involving *additional* background knowledge from other sources is benign and can lead to easy knowledge. According to this account, knowledge-closure holds for complex bootstrapping but not for direct bootstrapping.

¹⁰ A further possibility seems to be to argue that deductive bootstrapping can lead to easy knowledge, since knowledge-closure would fail otherwise, but that inductive bootstrapping cannot lead to easy knowledge. However, limiting easy knowledge to deductive bootstrapping is not a viable strategy to avoid closure failure, as I have argued.

However, this account has the undesired consequence that there are propositions about our overall reliability that we can never know. Our overall reliability concerns all potential knowledge-sources. Therefore, any form of bootstrapping to the conclusion that we are generally reliable is a form of direct bootstrapping, which does not lead to knowledge. This seems implausible given that we want to have knowledge about our overall reliability. Moreover, S can know of each potential knowledge source that it is reliable without knowing that she is generally reliable, according to this account, which is a further implausible instance of closure failure. Obviously, the structure of this problem is very similar to the problem of closure failure for inductive bootstrapping already discussed.

(c) Treating different belief forming sources differently

One can also distinguish between benign instances of bootstrapping that allow easy knowledge and malignant instances that do not by treating different belief-forming sources differently.¹¹ One could, for example, argue that perception allows easy knowledge about the reliability or accuracy of our sense apparatus but that knowledge via technical instruments does not allow for knowledge about the reliability or accuracy of these instruments.

Sosa, for example, defends the view that perceptual knowledge deserves a privileged status. Sosa (2010, 138) draws a distinction between perceptual justification and instrumental justification and argues “that our senses enjoy a kind of default rational justification denied to (ordinary) instruments. That is to say, we are default justified in accepting the deliverances of our senses, but we need a rational basis for accepting the deliverances of our instruments.” Also Pryor (2000 and 2004) argues that we can have basic perceptual knowledge (or at least basic perceptual justification) but admits that we might treat different belief-forming sources differently, i.e. that belief forming sources other than perception might not deliver basic knowledge.¹² This line of argumentation naturally provides the basis for the view that perception allows for basic knowledge and for easy knowledge whereas other potential knowledge sources do not.

However, is it reasonable to treat different knowledge sources differently? There is a problem, since these accounts treat different belief-forming sources *prima facie* differently, i.e. they treat them differently without applying a general criterion such as reliability. But how can we, for example, *argue* that perception delivers basic knowledge and easy knowledge but that

¹¹ For a similar typology see Vogel (2008, 525) who distinguishes between the views that easy knowledge is prohibited for all sources, for some sources, and for no sources.

¹² See Pryor (2004, 355).

instrumental devices do not, without taking their actual reliability into account? I think that such arguments are hard to find. In this respect, any such account seems ad hoc. Presumably, such accounts can deliver the desired results about the extent of our knowledge. In this respect, these accounts are in accordance with our intuitions. However, the problem is simply that it is hard to find arguments for such accounts.

Type-4 solutions

Type-4 solutions reject easy knowledge for some forms of knowledge but accept it for other forms. A strong version of a type-4 solutions introduces two *kinds* of knowledge, one that is closed under known entailment and, therefore, allows easy knowledge, and one for which knowledge-closure does not hold Cohen (2002), for example, adapts Sosa's (2007) distinction between animal knowledge and reflective knowledge and suggests that one can have basic animal knowledge, but that this kind of knowledge is not closed under known entailment, since it only involves tracking the truth. Furthermore, one can argue that reflective knowledge *is* closed under known entailment, hereby accepting any instance of easy knowledge for reflective knowledge, but rejecting easy knowledge for animal knowledge. In this case, bootstrapping can lead to the more reflective form of knowledge, but not to the less reflective form.¹³

A weaker version introduces two kinds of truly talking about knowledge, one that allows truly talking about easy knowledge and one that does not. The second version is an instance of contextualism about knowledge. Contextualists cannot claim that easy knowledge is possible in one context but not in the other unless they give up knowledge-closure in at least one context. However, contextualists typically endorse knowledge-closure and, therefore, any instance of easy knowledge within one context. DeRose (1995), for example, calls the conjunction of knowing that p, knowing that p entails q and not knowing that q 'abominable'. DeRose (1995) defines full-fledged skeptical hypotheses that not only contain a hypothesis that one falsely believes, but also an explanation for why one falsely believes. He argues that only when confronted with a full-fledged skeptical hypothesis we raise the standards up to level where sensitivity is necessary and where the skeptic is right. Accordingly, in context low we can have basic knowledge and easy knowledge, in context high we can have neither. Thus, his account allows for treating different propositions differently. However, his account is not

¹³ See Weisberg (2012) for a criticism of Cohen's account.

in accordance with our intuitions, since some full-fledged skeptical hypotheses can be sensitively believed.¹⁴

3 Rejecting basic knowledge entails skepticism

Accepting basic knowledge but rejecting easy knowledge entails numerous instances of closure failure. Since easy knowledge is knowledge via inference from basic knowledge, we can avoid easy knowledge by rejecting basic knowledge. In this section, I will stress the point that rejecting any kind of basic knowledge has the undesired consequence of leading to an infinite regress, which implies global skepticism according to finitist conceptions of knowledge and justification.¹⁵ We can formulate the assumption that no kind of basic knowledge is possible as follows:

No basic knowledge (NB)

S can only know that p via belief-forming source O if S has prior knowledge that O is reliable.

It becomes obvious that NB leads to an infinite regress, if we make explicit how S can only have *prior* knowledge that O is reliable. S's knowledge that O is reliable is prior to S's knowledge via O only if S knows that O is reliable via a knowledge-forming source O' and O and O' are not identical. Accordingly, the infinite regress takes the following form:

S only knows via O₁ that p if S has prior knowledge that O₁ is reliable.
S only knows that O₁ is reliable if S knows this via a further source O₂.
S only knows via O₂ that O₁ is reliable if S has prior knowledge that O₂ is reliable.
S only knows that O₂ is reliable if S knows this via a further source O₃.
...

It is assumed that none of the belief forming sources O₁, O₂... are identical. Here is an example:

S acquires knowledge about the current temperature in the room via thermometer T₁ only if S has prior knowledge that T₁ is reliable.
S knows that T₁ is reliable if S knows this via thermometer T₂.
S acquires knowledge via T₂ that T₁ is reliable only if S has prior knowledge that T₂ is reliable.
S knows that T₂ is reliable if S knows this via thermometer T₃.

¹⁴ See Melchior (2014). For a more a general criticism of sensitivity based analyses of the bootstrapping problem see Melchior (2015).

¹⁵ This point has already been stressed in the literature. See Cohen (2002).

There is wide agreement that due to their cognitive limitations human beings cannot acquire knowledge based on infinite regresses of reasoning.¹⁷ Hence, the claim that we cannot have basic knowledge leads to an infinite regress, which traditionally implies the general skeptical claim that we cannot know any proposition p .

4 Conclusion

An analysis of the problems related to the easy knowledge problem provides the following broader picture. Accepting basic knowledge commits one to accepting either easy knowledge or knowledge-closure. Moreover, rejecting basic knowledge leads to skepticism, unless one accepts infinitism. Hence, finite conceptions of knowledge face the trilemma of being forced to accept easy knowledge, closure failure, or skepticism. Agrippa's trilemma is the claim that the only three alternatives to skepticism are foundationalism, coherentism, and infinitism, but that each of these alternatives is unacceptable. There is wide agreement that coherentism and infinitism are not appropriate solutions to Agrippa's trilemma. However, if one rejects coherentism and infinitism, then one faces this new trilemma. Thus, one might think of this new trilemma as supporting coherentism or infinitism. I cannot enter this discussion here, but it is doubtful that coherentism, at least, will be better off.¹⁸

University of Graz
Department of Philosophy
Heinrichstraße 26/V
8010 Graz, Austria
Email: guido.melchior@uni-graz.at

¹⁶ In this example, S only uses belief-forming sources of the same type, but S can also use belief-forming sources of different types. For example, S can know that T_1 is reliable by reading the testing reports of the thermometer producing company.

¹⁷ This view is rejected by infinitism, which is the claim that infinite regresses are not necessarily vicious. See Klein (1999 and 2007).

¹⁸ For example, Lehrer's (2000) coherentism is not better off than foundationalism concerning bootstrapping. This version of coherentism does not accept basic knowledge, but it accepts knowledge about the reliability or trustworthiness of source without knowledge from an additional source. See Melchior (2012).

References

- Bergmann, Michael. 2004. "Epistemic Circularity: Malignant and Benign." *Philosophy and Phenomenological Research* 69(3):709-727.
- Bergmann, Michael. 2006. *Justification without Awareness*. Oxford: Oxford University Press.
- Blome-Tillman, Michael. 2006. "A Closer Look at Closure Skepticism." *Proceedings of the Aristotelian Society* 106:383-392.
- Cohen, Stewart. 2002. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research* 65(2):309-329.
- Cohen, Stewart. 2005. "Why Basic Knowledge is Easy Knowledge." *Philosophy and Phenomenological Research* 70(2):417-430.
- David, Marian and Ted Warfield. 2008. "Knowledge-Closure and Skepticism." In *Epistemology: New Essays*, edited by Quentin Smith, 137-187. Oxford: Oxford University Press,
- DeRose, Keith. 1995. "Solving the Skeptical Problem." *The Philosophical Review* 104(1):1-52.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. New York and Oxford: Oxford University Press.
- Hawthorne, John. 2005. "The Case for Closure." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa, 26-42. Malden, Oxford: Blackwell Publishing.
- Klein, Peter. 1999. "Human Knowledge and the Infinite Regress of Reasons." *Philosophical Perspectives* 13:297-325.
- Klein, Peter. 2004. "Closure Matters: Academic Skepticism and Easy Knowledge." *Philosophical Issues* 14(1):165-184.
- Klein, Peter. 2007. "Human Knowledge and the Infinite Progress of Reasoning." *Philosophical Studies*, 134(1):1-17.
- Lehrer, Keith. 2000. *Theory of Knowledge*, 2nd edition. Boulder: Westview.
- Markie, Peter. 2005. "Easy Knowledge." *Philosophy and Phenomenological Research* 70(2):406-416.
- Melchior, Guido. 2012. "Skepticism: Lehrer versus Mooreanism." *Philosophical Studies* 161(1):47-58.
- Melchior, Guido. 2014. "Skepticism: The Hard Problem for Indirect Sensitivity Accounts." *Erkenntnis* 79(1):45-54.

- Melchior, Guido. 2015. "The Heterogeneity Problem for Sensitivity Accounts." *Episteme* 12(4):479-496.
- Pryor, James. 2000. "The Skeptic and the Dogmatist." *Noûs* 34(4):517–549.
- Pryor, James. 2004. "What's Wrong with Moore's Argument?." *Philosophical Issues* 14(1):349–378.
- Sosa, Ernest. 1999. "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13:141-153.
- Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford: Oxford University Press.
- Sosa, Ernest. 2010. *Knowing Full Well*. Princeton: Princeton University Press.
- Vogel, Jonathan. 1990. "Are There Counterexamples to the Closure Principle?" In *Doubting: Contemporary Perspectives on Skepticism*, edited by Michael D. Roth and Glenn Ross, 13-29. Dordrecht: Kluwer.
- Vogel, Jonathan. 2000. "Reliabilism Leveled." *Journal of Philosophy* 97(11):602-623.
- Vogel, Jonathan. 2008. "Epistemic Bootstrapping." *Journal of Philosophy* 105(9):518-539.
- Weisberg, Jonathan. 2012. "The Bootstrapping Problem." *Philosophy Compass* 7/9:597-610.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.