

HIGHER-ORDER DEFEAT WITHOUT EPISTEMIC DILEMMAS

Mattias Skipper

Penultimate draft, forthcoming in *Logos & Episteme*.

Abstract. Many epistemologists have endorsed a version of the view that rational belief is sensitive to higher-order defeat. That is to say, even a fully rational belief state can be defeated by (sufficiently strong) misleading higher-order evidence, which indicates that the belief state is irrational. In a recent paper, however, Maria Lasonen-Aarnio (2014) calls this view into doubt. Her argument proceeds in two stages. First, she argues that higher-order defeat calls for a *two-tiered* theory of epistemic rationality. Secondly, she argues that there seems to be no satisfactory way of avoiding *epistemic dilemmas* within a two-tiered framework. Hence, she concludes that the prospects look dim for making sense of higher-order defeat within a broader theoretical picture of epistemic rationality. Here I aim to resist both parts of Lasonen-Aarnio's challenge. First, I outline a way of accommodating higher-order defeat within a *single-tiered* framework, by amending epistemic rules with appropriate provisos for different kinds of higher-order defeat. Secondly, I argue that those who nevertheless prefer to accommodate higher-order defeat within a two-tiered framework can do so without admitting to the possibility of epistemic dilemmas, since epistemic rules are not always accompanied by 'oughts' in a two-tiered framework. The considerations put forth thus indirectly vindicate the view that rational belief is sensitive to higher-order defeat.

1. Introduction

Many epistemologists have endorsed a version of the view that rational belief is sensitive to higher-order defeat. That is to say, even a fully rational belief state can be defeated by (sufficiently strong) misleading higher-order evidence, which indicates that the belief state is irrational. Here is a putative example:¹

Self-Enhancement Bias: John rationally believes that he is better than most people at driving. However, when reading today's newspaper, John learns about the well-documented *self-enhancement bias*: the widespread tendency to overrate oneself on a wide range of qualities and abilities, including intelligence, driving skills, and so on. As it happens, John is one of the few people who does not suffer from the bias.

¹ Similar cases of misleading higher-order evidence can be found in, e.g., Christensen (2010; forthcoming), Dorst (forthcoming a,b), Horowitz (2014), Skipper (forthcoming a,b), Whiting (2017; forthcoming), and Worsnip (2018; forthcoming).

Upon learning about the self-enhancement bias, how, if at all, should John revise his opinion about his own driving skills? In order to avoid getting sidetracked by issues concerning the (un)reliability of current journalism and experimental psychology, let us simply assume that John, after having read the newspaper, has good reason to think that the self-enhancement bias is indeed a widespread and pervasive phenomenon. Given this, it seems plausible to say that John should give up his belief (or at least become less confident) that he is better than most people at driving. After all, he has strong—albeit misleading—reason to think that his belief is the result of an irrational bias, and it seems plausible to say that one should give up beliefs that one has strong reason to think are irrational.

In recent years, a number of attempts have been made at vindicating the intuition that agents like John should revise their opinions, even if the higher-order evidence at hand is misleading. Perhaps the most prominent suggestion is due to David Christensen (2007; 2010; 2011; forthcoming), who points out that someone who disregards a body of misleading higher-order evidence seems to fall prey to an objectionable kind of dogmatism or question-begging reasoning: if John does not revise his opinion in response to the evidence about the self-enhancement bias, he must, it seems, take the evidence to be misleading. But the evidence is obviously only misleading if John does not suffer from the self-enhancement bias. So if John does not revise his opinion in response to the evidence about the self-enhancement bias, he must treat himself as an exception to the rule. However, the thought goes, John thereby seems to beg the question the against the studies on the self-enhancement bias. Contraposing: since John *shouldn't* engage in this kind of question-begging reasoning, he *should* revise his opinion about his own driving skills in response to the evidence about the self-enhancement bias.²

Despite the intuitive pull of this diagnosis, Maria Lasonen-Aarnio (2014) has recently challenged the view that rational belief is sensitive to higher-order defeat. Her argument proceeds in two stages. First, she argues that higher-order defeat calls for a *two-tiered* theory of epistemic rationality. Secondly, she argues that there seems to be no satisfactory way of avoiding *epistemic dilemmas* within a two-tiered framework. Hence, she concludes that the

² For other considerations in favor of the view that rational belief is sensitive to higher-order defeat, see also Schoenfield (2016) and Elga (2007). In effect, Elga argues that *peer disagreement* must act as a higher-order defeater, since one could otherwise become justified in considering an epistemic *peer* to be an epistemic *inferior* merely by discovering that he or she disagrees with oneself. While many have found Elga's view basically correct, those who don't typically still accept that peer disagreement can have defeating force—see, e.g., Kelly (2010) and Lackey (2008). To my knowledge, Titelbaum (2015) is the only current proponent of the view that disagreement can never have defeating force.

prospects look dim for making sense of higher-order defeat within a broader theoretical picture of epistemic rationality. If she is right, we are faced with the counterintuitive result that John should *not* revise his opinion about his own driving skills upon learning about the self-enhancement bias (and, more generally, that that a rational belief can remain rational even in light of strong evidence to the contrary).

Here I aim to resist both parts of Lasonen-Aarnio's challenge. After laying out Lasonen-Aarnio's argument in more detail (§2), I outline a way of accommodating higher-order defeat within a *single-tiered* framework, by amending epistemic rules with appropriate provisos for different types of higher-order defeat (§3). I then argue that those who nevertheless prefer to accommodate higher-order defeat within a two-tiered framework can do so without admitting to the possibility of epistemic dilemmas, since epistemic rules are not always accompanied by 'oughts' in a two-tiered framework (§4). The considerations put forth thus indirectly vindicate the view that rational belief is sensitive to higher-order defeat.

2. Lasonen-Aarnio on Higher-Order Defeat

Lasonen-Aarnio frames her discussion within a *rule-based* picture of epistemic rationality on which rational epistemic agents are characterized by following correct epistemic rules. We can think of epistemic rules as abstract functions from epistemic situations to sets of rationally permitted belief states. For example, one epistemic rule might enjoin me to believe that p whenever I have reliable testimony that p , whereas another epistemic rule might enjoin me to believe that p whenever it looks to me as if p . Needless to say, this picture gives rise to a number of tricky questions: how are epistemic rules to be formulated? What does it take for an epistemic rule to be 'correct'? What does it mean to 'follow' an epistemic rule in the relevant sense? I shall largely dodge such foundational issues in what follows.³ For present purposes, we can simply grant that a broadly rule-based framework provides a fruitful way of thinking about epistemic rationality.

In line with Lasonen-Aarnio, let us say that a rule-based theory of epistemic rationality is 'single-tiered' given that following correct epistemic rules is both necessary and sufficient for epistemic rationality:

Single-Tiered Framework: An agent's belief state S is epistemically rational if and only if S is the result of following correct epistemic rules.

³ See Boghossian (2008) for an illuminating discussion of various foundational issues concerning epistemic rule-following.

Suppose we aim to accommodate higher-order defeat within such a single-tiered framework. We must then ensure that whenever an agent in a rational belief state S receives sufficiently strong evidence that S is irrational, S is defeated by that evidence. That is, as Lasonen-Aarnio (2014, pp. 321-22) observes, epistemic rules must satisfy the following condition in order to be correct:

Rule Condition: If an agent has sufficiently strong overall evidence that her belief state S is not the result of following correct epistemic rules, then S is not the result of following correct epistemic rules.

This condition ensures that whenever an agent in a rational belief state receives a higher-order defeater, the belief state is no longer rational for the agent. Hence, by requiring that epistemic rules must satisfy the Rule Condition to be correct, we have a systematic way of making sense of higher-order defeat within a single-tiered framework.

However, according to Lasonen-Aarnio (pp. 322-24), we shouldn't expect that every correct epistemic rule satisfies the Rule Condition. Consider, for instance, a rule that enjoins me to believe that q if I believe that p & $(p \rightarrow q)$. This rule doesn't satisfy the Rule Condition: it's perfectly possible for my belief state to be the result of correctly applying modus ponens, even if I have strong reason to think that modus ponens is invalid or that I'm unable to apply modus ponens correctly. Yet, Lasonen-Aarnio submits, this shouldn't lead us to conclude that modus ponens is incorrect. Rather, we should maintain that modus ponens is correct, but that one can get misleading evidence to the contrary. Later, in §3, I will suggest that we might reasonably hold that modus ponens is an incorrect epistemic rule in virtue of violating the Rule Condition, but maintain that there is nonetheless something importantly right about modus ponens, which isn't diminished by the possibility of acquiring misleading evidence to the contrary. For now, however, let us proceed on the assumption that epistemic rules need not satisfy the Rule Condition to be correct. This means that we must look beyond a single-tiered framework to make sense of higher-order defeat.

According to two-tiered theories, as Lasonen-Aarnio understands them, following correct epistemic rules is necessary, but not sufficient, for epistemic rationality:

Two-Tiered Framework: An agent's belief state S is epistemically rational if and only if:

- (i) S is the result of following correct epistemic rules; and
- (ii) The agent does not have strong overall evidence that S is not the result of following correct epistemic rules.

If we adopt a two-tiered version of the rule-based framework, we can say that a belief state S is irrational despite resulting from following correct epistemic rules. For example, my belief

state might be the result of following modus ponens, yet be irrational because I have strong reasons to think that modus ponens is invalid or that I'm unable to correctly apply modus ponens. And John's belief that he is better than average at driving might be the result of following correct inductive rules, yet be irrational because he has strong reasons to think that his belief is the result of an irrational bias. As such, two-tiered theories give us the resources to make sense of higher-order defeat without assuming that epistemic rules must satisfy the Rule Condition to be correct.

But according to Lasonen-Aarnio (2014, pp. 328-30), two-tiered theories have a serious drawback, namely that they generate *epistemic dilemmas*: situations where an agent ought to adopt incompatible doxastic attitudes (for instance, believing and suspending judgment about the same proposition). Her reasoning goes as follows: let R be a correct epistemic rule, which enjoins me to believe that p in my current epistemic situation. Assuming that R violates the Rule Condition, it must be possible for me to receive a higher-order defeater in light of which I should give up my belief in p , although R still enjoins me to believe that p . Hence, there must exist another correct epistemic rule R' , which enjoins me to suspend judgment about p upon having received the higher-order defeater. As a result, the rules R and R' end up giving me conflicting recommendations: R enjoins me to believe that p , whereas R' enjoins me to suspend judgment about p . This leads Lasonen-Aarnio to conclude that “[i]f correct rules are accompanied by oughts, it looks like [I] ought to believe that p , and [I] ought to suspend judgment in p ” (Lasonen-Aarnio 2014, p. 329). In other words, it looks like I face an epistemic dilemma.⁴

These are the bare bones of Lasonen-Aarnio's two-piece challenge to the proponent of higher-order defeat: (i) higher-order defeat seems to call out for a two-tiered theory of epistemic rationality, and (ii) two-tiered theories seem to generate epistemic dilemmas. In her subsequent discussion of this challenge, Lasonen-Aarnio considers and rejects various strategies for meeting the challenge, either by embracing the possibility of epistemic dilemmas, or by trying to specify a version of the rule-based framework (either single-tiered or two-tiered) that doesn't generate epistemic dilemmas after all. I will return to this part of Lasonen-Aarnio's discussion in due course, but for now it suffices to have the basic challenge on the table.

⁴ Lasonen-Aarnio (2014, pp. 329-30) also argues that cases of conflicting recommendations by correct epistemic rules can arise even if epistemic rules are formulated in terms of what one is permitted to believe, rather than in terms of what one is required to believe. But for present purposes, we can simply grant that epistemic rules specify what is required rather than what is permitted.

3. Higher-Order Defeat in a Single-Tiered Framework

This section concerns the first part of Lasonen-Aarnio's challenge. I want to defend two claims, a negative and a positive. Negatively, I will argue that Lasonen-Aarnio's case against the Rule Condition is ill-founded. Positively, I will use the Rule Condition to outline and motivate a way of accommodating higher-order defeat within a single-tiered framework.

First, the negative point. Recall that Lasonen-Aarnio's case against the Rule Condition was based on the observation that there are seemingly correct epistemic rules, such as modus ponens, that nevertheless fail to satisfy the Rule Condition. I think there is something right and something wrong about this observation, which can be brought out by distinguishing two different conceptions of what it means for an epistemic rule to be 'correct'. On one understanding, an epistemic rule is correct to the extent that it is truth-conducive or truth-preserving, where truth-conduciveness is a matter of generating true beliefs, and truth-preservation is a matter of preserving true beliefs.⁵ On this understanding, modus ponens is clearly correct in virtue of being deductively valid. But various non-deductive (e.g. inductive and abductive) rules are presumably also correct in this sense, despite falling short of perfect truth-preservation. On another understanding, an epistemic rule is correct to the extent that it is rational for agents to follow it across all—or at least a sufficiently wide range of—epistemic situations. On this understanding, it is much less clear that modus ponens is correct, precisely because one can have strong reason think that modus ponens is invalid or that one is unable to apply modus ponens correctly. Hence, whether modus ponens is a correct epistemic rule may well depend on the notion of correctness we have in mind.

In the present context, since we are considering what it's rational for agents to believe in different epistemic situations, we presumably want to interpret the Rule Condition as a condition on what it takes for an epistemic rule to be correct in the latter sense. Accordingly, when we say that modus ponens is incorrect in virtue of violating the Rule Condition, we are not saying that modus ponens is anything less than perfectly truth-preserving. Rather, we are saying that it isn't always rational to follow modus ponens, since one can have strong reason to think that modus ponens is invalid or that one is unable to apply modus ponens correctly. And this doesn't strike me as an implausible thing to say, especially if we consider logical rules of a more complex nature than modus ponens.

Someone might worry that we end up with a circular account of epistemic rationality, if we spell out what it means for an epistemic rule to be correct in terms of whether it is rational to follow the rule across a sufficiently wide range of circumstances. But there seems to be a natural way of avoiding such circularity by holding that it's rational for an agent to follow a

⁵ This distinction is equivalent (or at least very similar in spirit) to Goldman's (1979) distinction between 'conditional' and 'unconditional' reliability.

given epistemic rule provided that the agent has (sufficiently strong) overall reason to think that doing so is (sufficiently) conducive to the formation of true beliefs. This allows us to avoid circularity while being faithful to the initial motivation for saying that it can be irrational for an agent to follow an epistemic rule, even if the rule is in fact perfectly truth-conducive or truth-preserving.

Another potential worry about the distinction between a rule's being truth-conducive (or truth-preserving) and its being rational to follow is that it obscures the connection between rationality and truth. If rational belief is supposed to aim at truth, how could it be irrational to follow a rule that is perfectly truth-conducive or truth-preserving? I think we can answer this worry by drawing an analogy between propositions and rules. Just as it can be irrational to believe a true proposition (if one has strong reason to think that the proposition is false), so it can be irrational to follow a truth-conducive or truth-preserving rule (if one has strong reason to think that following the rule isn't conducive to the formation of true beliefs). And just as we can nevertheless aim to believe only what is true, we can nevertheless aim to follow only rules that are truth-conducive or truth-preserving. As such, the distinction between a rule's being truth-conducive (or truth-preserving) and its being rational to follow seems compatible with there being an important normative connection between the two. Obviously, the connection will be less direct than saying that a rule is rational to follow if and only if it is truth-conducive (or truth-preserving). But this seems no less plausible than denying that a proposition is rational to believe if and only if it is true. Hence, the claim that epistemic rules must satisfy the Rule Condition to be correct (in the relevant sense of 'correct') strikes me as quite plausible, or at least not as implausible as Lasonen-Aarnio seems to think.

Next, the positive point. Even if the above considerations are basically correct, it remains to be seen whether there is a satisfactory way of ensuring that correct epistemic rules satisfy the Rule Condition. In passing Lasonen-Aarnio considers the possibility of amending epistemic rules with provisos for different higher-order defeaters. However, she takes it to be "at least a *prima facie* challenge" to show how an epistemic rule could have provisos "for *all possible* higher-order defeaters built into it" (Lasonen-Aarnio 2014, p. 324). Here I want to suggest a way of meeting this challenge. As a first step, consider the following qualification of modus ponens:

Modus Ponens (provisional): Follow modus ponens, unless:

- (i) Some correct epistemic rule enjoins you to believe that modus ponens is incorrect;
or
- (ii) Some correct epistemic rule enjoins you to believe that you are likely to misapply modus ponens.

This rule features two provisos, which correspond to the two central types of higher-order defeat discussed by Lasonen-Aarnio: (i) higher-order defeat due to evidence that a given epistemic rule is incorrect, and (ii) higher-order defeat due to evidence that one is incapable of applying an epistemic rule correctly. Perhaps some higher-order defeaters do not fall neatly within these two categories. If so, we might have to amend modus ponens with additional provisos. But the provisos included here cover at least the most commonly discussed cases of higher-order defeat in the recent literature.

If all correct epistemic rules come with provisos akin to those in Modus Ponens (provisional), we have the resources to accommodate higher-order defeat within a single-tiered framework. Consider a structurally similar case to Self-Enhancement Bias, adapted from Christensen (2016, p. 401):

Modus Ponens on Drugs: Suzy rationally believes that p & $(p \rightarrow q)$, and she forms a belief in q as a result of correctly applying modus ponens. But she is now told from a reliable source that she has been given a reason-distorting drug that subtly, but significantly, impairs her ability to perform even simple deductive inferences. As a matter of fact, Suzy has not been drugged.

Assuming that this is a case of higher-order defeat, Suzy is rationally permitted to believe that q *before* being told that she has been drugged, but she isn't rationally permitted to believe that q *after* being so told. And the proposed single-tiered account yields precisely this result, since the second proviso in Modus Ponens (provisional) is met *after*, but not *before*, Suzy is told about the drug. (The same would obviously hold if Suzy's higher-order defeater had instead consisted of evidence indicating that modus ponens is invalid.)

However, while epistemic rules like Modus Ponens (provisional) allow us to make sense of *ordinary* cases of higher-order defeat—like Modus Ponens on Drugs and Self-Enhancement Bias—they do not allow us to make sense of *all possible* cases of higher-order defeat. For as Lasonen-Aarnio (2014, p. 323) points out, Suzy could in principle get strong reason to think that Modus Ponens (provisional) is incorrect or that she is unable to apply it correctly. Hence, if we want to accommodate such non-standard cases of higher-order defeat, we need to amend modus ponens with a second layer of provisos:

Modus Ponens (doubly provisional): Follow Modus Ponens (provisional), unless:

- (i) Some correct epistemic rule enjoins you to believe that Modus Ponens (provisional) is incorrect; or

- (ii) Some correct epistemic rule enjoins you to believe that you are likely to misapply Modus Ponens (provisional).

Continuing this path, we run into an infinite regress: a third layer of provisos is needed to accommodate a case where Suzy gets strong reason to think that Modus Ponens (doubly provisional) is incorrect or that she is unable to apply it correctly, and so on. Thus, if we want to accommodate all possible cases of higher-order defeat, it looks like correct epistemic rules must feature infinite layers of provisos for different kinds of higher-order defeat.

What should we make of the resulting view? By way of comparison, consider a different single-tiered view, which Lasonen-Aarnio (2014, §4) discusses under the label ‘The Über-Rule View’. On this view, a belief state is epistemically rational if and only if it is the result of following an overarching epistemic rule—the Über-rule—which determines, for any given epistemic situation, which belief states are rationally permitted in that situation. Accordingly, higher-order defeat is to be made sense of by reverse-engineering the content of the Über-rule in such a way that it delivers the desired verdicts in cases like Self-Enhancement Bias.

While Lasonen-Aarnio raises a number of different worries about The Über-Rule View, perhaps the most serious problem is that an extensionally adequate Über-rule must presumably be a highly complex construct of gerrymandering, which can hardly be expressed in “any finite, informative way” (Lasonen-Aarnio 2014, p. 332). Consequently, the worry goes, the Über-rule can hardly offer the kind of epistemic guidance that we should expect correct epistemic rules to be able to offer (Lasonen-Aarnio 2014, p. 333).

Let us straightaway grant this criticism of the Über-rule view. Does the single-tiered view outlined above fall prey to a similar criticism? It seems not. Despite the addition of infinite layers of provisos, the view is nevertheless able to offer genuine guidance. Consider Modus Ponens on Drugs: we previously saw that Suzy’s initial belief that q is deemed rational, since none of the provisos in Modus Ponens (provisional) are initially met. But when Suzy receives the higher-order defeater, she may no longer believe that q , because the second proviso in Modus Ponens (provisional) is now met. Hence, it looks like Suzy has all the guidance she needs to decide which doxastic attitude she should adopt towards q before and after receiving the higher-order defeater. And the same obviously goes for other ordinary cases of higher-order defeat such as Self-Enhancement Bias.

What about non-standard cases of higher-order defeat? Suppose Suzy gets a second higher-order defeater, which consists of evidence that Modus Ponens (provisional) is incorrect. Here things become a bit more complicated. On the proposed view, Suzy should resort to Modus Ponens (doubly provisional), because one of the provisos in Modus Ponens

(provisional) is met. But which doxastic attitude does this rule enjoin Suzy to adopt towards q ? It seems to depend on the exact content of the second higher-order defeater. Suppose the second higher-order defeater indicates that Modus Ponens (provisional) is incorrect, without indicating *why* it is incorrect. We can then imagine how Suzy might entertain different possible explanations of why Modus Ponens (provisional) is incorrect. One possible explanation is that it is always reasonable to follow modus ponens, contrary to what Modus Ponens (provisional) says. If so, Suzy is presumably enjoined to believe that q , since this is what modus ponens enjoins. But another possible explanation is that Modus Ponens (provisional) is incorrect in virtue of failing to take into account every type of situation in which it is unreasonable to follow modus ponens. If so, Suzy is presumably enjoined to suspend judgment about q , since her current situation still counts as one in which it is unreasonable to follow modus ponens. Thus, it looks like Suzy faces a problem of underdetermination about which doxastic attitude she is enjoined to adopt towards q .

Is this sort of underdetermination a vice or a virtue of the proposed view? Someone who thinks that there must always be a determinate answer to questions about which doxastic attitudes we should adopt towards different propositions might want to conclude that the proposed single-tiered theory fails to offer enough guidance in situations like Suzy's. However, it isn't clear to me that there *is* a determinate answer to the question of which doxastic attitude Suzy should adopt towards q . Suzy's situation might just be a genuine case of underdetermination concerning which doxastic attitude she should adopt. If so, we should consider the underdetermination property of the proposed view as a feature, not a bug.

In any case, regardless of how we settle the issue about underdetermination, the proposed view at least seems well placed to offer genuine guidance in ordinary cases of higher-order defeat like Modus Ponens on Drugs and Self-Enhancement Bias. As such, I see no reason to be pessimistic about the prospects for accommodating higher-order defeat within a single-tiered framework.

4. Higher-Order Defeat in a Two-Tiered Framework

This section concerns the second part of Lasonen-Aarnio's challenge. My aim is to argue that even if the single-tiered view outlined in the previous section cannot be made to work, we can still make sense of higher-order defeat within a two-tiered framework without committing ourselves to the possibility of epistemic dilemmas.

Recall Lasonen-Aarnio's basic reasons for thinking that two-tiered theories tend to generate epistemic dilemmas: (i) correct epistemic rules can give conflicting

recommendations in a two-tiered framework, and (ii) correct epistemic rules are “accompanied by oughts” (Lasonen-Aarnio 2014, p. 329). My strategy will be to grant the former claim, but deny the latter: correct epistemic rules may well deliver conflicting recommendations in a two-tiered framework, but we should not expect such conflicting recommendations to constitute genuine epistemic dilemmas.

Let us begin with a simple observation: epistemic rules play a very different role in a two-tiered framework than in a single-tiered framework. Suppose a rule R enjoins an agent to believe that p in a given situation. On a single-tiered view, R 's recommendation will invariably constitute a rational requirement on part of the agent to believe that p , since single-tiered theories work with a straightforward ‘one-to-one’ correspondence between requirements of rationality and recommendations of correct epistemic rules. Accordingly, conflicting recommendations by correct epistemic rules always generate epistemic dilemmas within a single-tiered framework. By contrast, in a two-tiered framework, R 's recommendation need not constitute a rational requirement to believe that p , since two-tiered theories do not feature a one-to-one correspondence between requirements of rationality and recommendations of correct epistemic rules. Accordingly, conflicting recommendations by correct epistemic rules need not generate epistemic dilemmas within a two-tiered framework.

This observation already suffices to mitigate, at least to some degree, the second part of Lasonen-Aarnio's challenge. Suppose I'm enjoined to believe that p by a correct epistemic rule R . Assuming, as we do, that R need not satisfy the Rule Condition, it must be possible for me to acquire a higher-order defeater in light of which I should give up my belief that p , although R still enjoins me to believe that p . Hence, there must exist another correct epistemic rule R' , which enjoins me to suspend judgment about p upon having received the higher-order defeater. As a result, the rules R and R' end up giving me conflicting recommendations: R enjoins me to believe that p , whereas R' enjoins me to suspend judgment about p . This is what led Lasonen-Aarnio to conclude that two-tiered theories tend to generate epistemic dilemmas. However, since R and R' make their conflicting recommendations in a situation where I have received a higher-order defeater, proponents of the two-tiered view will presumably want to say that R 's recommendation no longer constitutes a rational requirement to believe that p , in which case I do not face an epistemic dilemma. After all, it was the ability to accommodate such cases of higher-order defeat that motivated the move to a two-tiered framework to begin with. Hence, we should not expect two-tiered theories to generate epistemic dilemmas in cases of higher-order defeat.

Obviously, this does not yet show that there is a satisfying way of filling in the details of a two-tiered theory, which ensures that conflicting recommendations by correct epistemic

rules never constitute epistemic dilemmas. In particular, it does not fend off Lasonen-Aarnio's criticism of the specific two-tiered view that she calls 'The Hierarchy View' (Lasonen-Aarnio 2014, §6). On this view, a belief state is epistemically rational if and only if it is the result of following a correct epistemic rule, which isn't overridden by a correct epistemic rule higher up the hierarchy. The hierarchy of rules is then determined by a meta-rule, which assigns an ordering of epistemic rules to each epistemic situation. Hence, given the right meta-rule, The Hierarchy View might help avoid epistemic dilemmas in cases of higher-order defeat: if R' overrides R , I will be required to suspend judgment about p upon having acquired the higher-order defeater, but I will not be required to believe that p .

However, just as Lasonen-Aarnio deems The Über-Rule View inapt to accommodate higher-order defeat within a single-tiered framework, she deems The Hierarchy View inapt to accommodate higher-order defeat within a two-tiered framework. Here the main worry concerns what happens if one gets strong reason to think that the meta-rule is incorrect or that one is unable to apply it correctly. According to Lasonen-Aarnio, if we want to accommodate such cases of higher-order defeat, we must introduce a set of meta-rules together with a 'meta-meta-rule', which gives us a ranking of the meta-rules, for each epistemic situation. But this strategy quickly leads to an infinite regress due to the further possibility of acquiring reasons to think that the meta-meta-rule is incorrect or that one is unable to apply it correctly. Hence, according to Lasonen-Aarnio, we face the same worry that led us to reject The Über-Rule View: if The Hierarchy View must feature infinitely many orders of meta-rules, it becomes unclear how the view can offer genuine guidance (Lasonen-Aarnio 2014, pp. 340-41).

This guidance-worry seems to have less force against The Hierarchy View than against The Über-Rule View. Of course, there is a sense in which neither view is "finitely expressible," but The Hierarchy View nevertheless seems informative in a way that The Über-Rule View is not. In fact, The Hierarchy View seems informative in much the same way as the single-tiered view proposed in the previous section. Consider Modus Ponens on Drugs: on The Hierarchy View, Suzy's belief that q is initially rational, since it is the result of following a correct epistemic rule, call it R , which is not overridden by any other correct epistemic rule. But when receiving the higher-order defeater, Suzy is no longer rationally permitted to believe that q , since R 's recommendation is now overridden by another rule, call it R' , which recommends her to suspend judgment about p . Hence, it looks like Suzy has all the guidance she needs to determine which doxastic attitude to adopt towards q before and after receiving the higher-order defeater. Obviously, The Hierarchy View allows for the possibility that the recommendation by R' is itself overridden at some later stage, in which case it might become

less clear what The Hierarchy View recommends (for reasons explained in the previous section). But at the very least, the Hierarchy seems able to offer genuine guidance in ordinary cases of higher-order defeat like Modus Ponens on Drugs and Self-Enhancement Bias. Hence, I see no reason to think that the kind of infinity featuring in The Hierarchy View is detrimental to the view. But even if it turns out that The Hierarchy View cannot ultimately be made to work, it would be premature to conclude that no two-tiered theory can be made to work, unless there are general grounds for expecting that two-tiered theories tend to generate epistemic dilemmas in cases of higher-order defeat. And I take the foregoing considerations to show that such general grounds are lacking.

References

- Boghossian, P. (2008): "Epistemic Rules." In: *The Journal of Philosophy* 105(9), pp. 472-500.
- Christensen, D. (2007): "Epistemology of Disagreement: The Good News." In: *Philosophical Review* 116, pp. 187-217.
- (2010): "Higher-Order Evidence." In: *Philosophy and Phenomenological Research* 8, pp. 185-215.
- (2011): "Disagreement, Question-Begging, and Epistemic Self-Criticism." In: *Philosophers' Imprint* 11.
- (2016): "Disagreement, Drugs, Etc.: From Accuracy to Akrasia." In: *Episteme* 13, pp. 397-422.
- (forthcoming): "Formulating Independence." In M. Skipper and A. Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*, Oxford University Press.
- Dorst, K. (forthcoming a): "Higher-Order Uncertainty." In M. Skipper and A. Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*, Oxford University Press.
- (forthcoming b): "Evidence: A Guide for the Uncertain." Forthcoming in *Philosophy and Phenomenological Research*.
- Elga, A. (2007): "Reflection and Disagreement." *Noûs* 41, pp. 478-502.
- Goldman, A. (1979): "What if Justified Belief?" In: *Epistemology. An Anthology*, E. Sosa and J. Kim (Eds.). Blackwell, pp. 340-53.
- Horowitz, S. (2014): "Epistemic Akrasia." In: *Noûs* 48, pp. 718-44.
- Kelly, T. (2010): "Peer Disagreement and Higher Order Evidence." In: *Social Epistemology: Essential Readings*, A. Goldman and D. Whitcomb (eds.), Oxford: Oxford University Press, pp. 183-217.

- Lackey, J. (2008): "What Should We Do When We Disagree?" In: *Oxford Studies in Epistemology* 3, T. Gendler and J. Hawthorne (eds.), Oxford: Oxford University Press, pp. 274–93.
- Lasonen-Aarnio, M. (2014): "Higher-Order Evidence and the Limits of Defeat." In: *Philosophy and Phenomenological Research* 88, pp. 314-45.
- Schoenfield, M. (2016): "An Accuracy Based Approach to Higher Order Evidence." In: *Philosophy and Phenomenological Research* 93, pp. 1-26.
- Skipper, M. (forthcoming a): "Reconciling Enkrasia and Higher-Order Defeat." Forthcoming in *Erkenntnis*.
- (forthcoming b): "Higher-Order Defeat and the Impossibility of Self-Misleading Evidence." In M. Skipper and A. Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*, Oxford University Press.
- Titelbaum, M. (2015): "Rationality's Fixed Point (Or: In Defense of Right Reason)." In: *Oxford Studies in Epistemology* 5, T. Gendler and J. Hawthorne (eds.), pp. 253-94.
- Whiting, D. (2017): "Against Second-Order Reasons." In: *Noûs* 51(2), pp. 398-420.
- (forthcoming): "Whither Higher-Order Evidence?" In M. Skipper and A. Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*, Oxford University Press.
- Worsnip, A. (2018): "The Conflict of Evidence and Coherence." In: *Philosophy and Phenomenological Research* 96(1), pp. 3-44.
- (forthcoming): "How Your Total Evidence Can Mislead About Itself." In M. Skipper and A. Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*, Oxford University Press.