

Higher-Order Defeat is Object-Independent

Forthcoming in *Pacific Philosophical Quarterly*

Penultimate draft. Please cite published version.

Abstract: Higher-order defeat occurs when one loses justification for one's beliefs as a result of receiving evidence that those beliefs resulted from a cognitive malfunction. Several philosophers have identified features of higher-order defeat that distinguish it from other familiar types of defeat. If higher-order defeat has these features, these are data an account of rational belief must capture. In this paper, I identify a new distinguishing feature of higher-order defeat, and I argue that both on its own, and in conjunction with the other distinguishing features, it favors an account of higher-order defeat, and hence rational belief, partly grounded in non-evidential, "state-given reasons" for belief.

"What is out of common is usually a guide rather than a hindrance."

--Sherlock Holmes, *A Study in Scarlet*

1. Introduction

Higher-order evidence is evidence about the relation between one's evidence and one's beliefs.

Higher-order defeat occurs when one loses justification for one's beliefs as a result of acquiring higher-order evidence that these beliefs resulted from some cognitive malfunction, flaw in reasoning, or other flawed process rather than from a proper assessment of the evidence.

Several philosophers have identified features of higher-order defeat that arguably distinguish it from more familiar types of defeat.¹ David Christensen rightly contends that, "[A]ny satisfactory epistemology will need to address the apparent peculiarities with which [higher-order evidence] presents us. We will not fully understand rational belief until we understand higher-order evidence."² Higher-order evidence and higher-order defeat are, indeed, peculiar,³ and these peculiarities are best seen as a guide to helping us understand the normative underpinnings of rational belief. This paper identifies a new distinguishing feature

¹ See, for instance, Feldman (2005), Christensen (2010), Schechter (2013), and Lasonen-Aarnio (2014).

² Christensen (2010a: 213).

³ Or, "out of common" as Holmes says.

of higher-order defeat, and thereby describes a new observation that theories of rational belief must capture.

It is too early to develop a complete theory of higher-order defeat because, I think, we are still in the data-gathering phase on this issue. Nevertheless, in the penultimate section, I tentatively suggest that the feature I identify favors a non-evidentialist account of higher-order defeat over an evidentialist one (i.e., a theory of higher-order defeat may need to appeal to more than variation in evidential support to explain this feature), and that this non-evidentialist account may also explain other peculiar features of higher-order defeat.

In §2, I (i) describe the two most familiar types of defeat, (ii) provide examples of higher-order defeat (HOD, hereafter), (iii) explain what others have thought distinguishes HOD from these familiar types of defeat, and (iv) motivate the project of studying distinctive features of defeat. In §3, I argue that HOD is “object-independent.” To preview, this means that whether one should revise one’s belief in response to receiving HOE is independent of the object of that belief. In §4, I argue that the object-independence of HOD favors a hybrid account of the normative underpinnings of rational belief in terms of both evidential, “object-given,” reasons and non-evidential, “state-given,” reasons. In §5, I conclude.

2. Background and Motivation

When acquiring new information results in a subject losing rationality or justification for one of her beliefs, that new information *defeats* the rationality or justification of that belief. I take this new information to be *evidence*, and I will use ‘justification’ and ‘rationality’ (and their

cognates) interchangeably.⁴ Although defeat is a relation between evidence and justification on my usage, I will sometimes speak of beliefs themselves being defeated.

Beliefs and defeat are both of two sorts, categorical and graded. I discuss both sorts of belief. If acquiring evidence makes a previously justified belief unjustified, the defeat is *categorical*. If the belief ends up being *less*, though still, justified, the defeat is *graded*. I focus on categorical defeat.

There are two familiar sorts of defeat: rebutting and undercutting. A standard example of rebutting defeat involves a subject believing a universally quantified proposition on the basis of a strong inductive generalization, and then encountering a counterexample to the universal proposition:

Rebutting: Bob observes a large number of swans, noticing that they are all white.

Let E_1 be his relevant evidence. On the basis of E_1 , Bob believes that all swans are white. Then Bob learns E_2 , that there is a black swan.

There are two central features of cases of rebutting defeat. First, the original evidence, E_1 , supports belief or high confidence in some proposition. Second, the defeating evidence, E_2 , supports the *negation* of this proposition. By hypothesis, Bob's original evidence supports his belief that all swans are white. Then he acquires evidence that there is a black swan. Bob's belief is no longer rational because, in this case, not only does his evidence no longer support this belief, it in fact supports its negation.

Undercutting defeat differs from rebutting defeat over the second feature. As John Pollock explains, undercutting defeat "attacks the connection between the evidence and the

⁴ It is plausible that one's views on the nature of evidence and evidence possession may influence one's views on defeat, and vice versa. Thanks to an anonymous referee.

conclusion, rather than attacking the conclusion itself.”⁵ Here is a common example of undercutting defeat:

Undercutting: Coco sees a wall that appears red to her. Let E_1 be her relevant evidence. On the basis of E_1 , Coco believes that the wall is red. Then she learns E_2 , that there is a red light shining on the wall.

E_2 is not evidence that the wall is not red. Intuitively, however, it does defeat Coco’s justification for believing that the wall is red. Roughly, E_2 is evidence that, in this context, E_1 is not evidence for the proposition that the wall is red.⁶ Like Rebutting, in Undercutting the subject’s original evidence supports belief in some proposition. Unlike Rebutting, in Undercutting the defeating evidence is not evidence that this proposition is false.

Although undercutting defeaters are not, *qua* undercutters, evidence that the defeated belief is false, this doesn’t mean that they *cannot* rebut this belief. With the proper background evidence in place, evidence that a wall is being illuminated by a red light *can* be evidence that the wall is not red. If Coco knows for certain that the wall is red only if it’s not being illuminated by a red light, then learning E_2 will not only undercut but also rebut her justification. When identifying the unique contributions made by one type of defeater in order to distinguish it from other types of defeaters, we need to ignore cases with this kind of rich background evidence.

Several philosophers argue that HOD is distinct from both rebutting defeat and undercutting defeat. Call this the *defeat distinctness hypothesis*. These philosophers have identified differences between HOD and these other types of defeat. Before reviewing these claims, we should consider examples of HOD. Here is a common case:

⁵ Pollock (1986: 39).

⁶ Alternatively, E_2 is evidence that makes E_1 less weighty. See Schroeder (2011) for a proposal like this.

Hypoxia: Andy is piloting an airplane and after doing some calculations, she becomes highly confident that she has enough fuel to fly to Hawaii on the basis of the following evidence, E_1 :

- A full tank contains 20,000 miles worth of fuel.
- The tank is $\frac{3}{4}$ full.
- Hawaii was 16,000 miles from Andy's point of departure.
- Andy has flown 5,000 miles toward Hawaii.

Then air traffic control provides her with a credible warning that as a result of a drop in her oxygen supply any reasoning she's done in the last five minutes has probably been distorted without her knowing it.⁷ Call this information HOE_1 .

Andy's reasoning is perfectly reasonable, we can suppose, as E_1 entails that she has enough fuel.⁸ Intuitively, though, learning HOE_1 makes maintaining her high confidence irrational.⁹

Consider another case:

Disagreement: Dale and Donna reasonably consider each other to be epistemic peers on meteorological issues, and they share the same (somewhat complicated) evidence, E_2 , concerning whether it's going to rain tomorrow. After Dale becomes confident that it will not rain, he learns, HOE_2 , that Donna believes that it will. Independent of their disagreement, Dale has no reason to think that Donna is not assessing her evidence properly.

⁷ This can actually happen. The condition is called hypoxia.

⁸ I assume that if she has more fuel than she needs, then she has enough. The first two propositions entail that she has 15,000 miles worth of fuel and the second two propositions entail that she needs 11,000 miles worth of fuel to make it to Hawaii. So, this evidence entails that she has more fuel than she needs, so it entails that she has enough.

⁹ For the same verdict, see Elga (2013), Christensen (2010b), and Schechter (2013). This example is originally due to Elga.

We can suppose that, in fact, Dale responded reasonably to his evidence. Nevertheless, intuitively, learning HOE_2 makes maintaining his confidence irrational.¹⁰

These are standard cases of HOD. On the basis of relatively complex reasoning, a subject adopts a reasonable belief. Then she acquires strong evidence that her belief is due to some sort of cognitive malfunction or flaw in reasoning. This HOE is also typically evidence that, from the subject's perspective, if this cognitive malfunction did influence her beliefs, she would not know it. In other words, in typical cases of HOD, the subject acquires evidence both that she is in the bad case and that she cannot tell from her perspective whether she is in the bad case. Acquiring this evidence makes maintaining her belief irrational. So, acquiring the HOE defeats the belief's rationality.

I'll now list three features of HOD identified by others in support of the defeat distinctness hypothesis. All three features distinguish HOD from rebutting defeat, and all but the first feature distinguish it from undercutting defeat.

First, unlike rebutting defeat, HOD does not typically provide subjects with evidence that the defeated proposition is false. This claim isn't that HOD *never* provides evidence of this, just as the corresponding claim isn't that undercutting defeat *never* does either. Like undercutting defeat, HOD *typically* defeats without rebutting.

Second, HOD has a *retrospective aspect*: it provides subjects with evidence that their beliefs were *originally* irrational.¹¹ This is not how rebutting or undercutting defeat work. When Bob learns that there is a black swan or when Coco learns that a red light is shining on the wall, this new evidence doesn't cast doubt on the rationality of their original beliefs.

¹⁰ Few philosophers working on disagreement dispute this claim, but see Kelly (2005) and Titelbaum (2015).

¹¹ See Christensen (2010: 185) and Lasonen-Aarnio (2014: 317). The phrase "retrospective aspect" is Lasonen-Aarnio's. Here is an alternative (rough) characterization of this retrospective aspect, covering cases in which the subject has not yet formed the relevant belief: HOD provides subjects with evidence that if they had formed beliefs of a certain kind, in a certain way, then those beliefs would have been irrational. Thanks to an anonymous referee for suggesting that I address this kind of case.

However, the HOE that the subjects in Hypoxia and Disagreement acquire (misleadingly) indicates that their original beliefs were irrational.

The third point can be separated into two claims. Depending on one's view of evidential support, however, it may turn out that the claims are equivalent. Joshua Schechter notes that conditional probabilities differ between undercutting cases and (some) HOD cases. Compare the conditional probabilities in Undercutting and Hypoxia. The probability that a wall that looks red is red is presumably greater than the probability that a wall that looks red is red given that the wall is illuminated by red lights. Whereas the probability that Andy has enough fuel given her original evidence is no greater than the probability that she has enough fuel conditional on her original evidence and that she is suffering from hypoxia.¹² After all, her evidence entails that she has enough fuel.

Christensen, going a step further, argues that HOE defeats the rationality of one's attitudes *without changing whether one's evidence supports* those attitudes.¹³ This is true, he argues, both in cases (like Andy's) where the evidence entails the believed proposition and also in cases in which the evidence merely supports, without entailing, the believed proposition.¹⁴ The claim that HOD leaves evidential support intact when the subject's evidence entails the defeated proposition is most plausible in a framework that measures evidential support probabilistically. In this framework, anything entailed by one's evidence will be assigned

¹² Schechter (2013: fn. 41).

¹³ In conversation, I've met with resistance to this interpretation. Some claim that although Christensen thinks that the first-order evidential support remains intact, he wouldn't say that the subject's *total* evidence still supports the defeated belief as strongly as it did before. This resistance is misguided; this is *exactly* what he says (2010a: 195): In entailment cases of misleading HOE, "it is not obvious why my *total evidence*, after I [acquire the HOE], does not support my original conclusion *just as strongly as it did beforehand*. ... there is a clear sense in which the facts which are not in doubt—the parameters of the puzzle—leave no room for anything other than my original answer. Or, to put it another way, the undoubted facts support my answer in the strongest possible way—they entail my answer—, and this kind of connection cannot be affected by adding more evidence" [emphasis added]. Hedging aside, in this passage Christensen clearly argues for the view I'm attributing to him in the text. Lasonen-Aarnio (2014: 318) also attributes this view to him.

¹⁴ Christensen (2010a: 197).

probability 1 because one's evidence is typically assigned probability 1, and a proposition that has probability 1 conditional on one's evidence has full evidential support. Moreover, it will continue to receive full support, no matter how much evidence one acquires, as long as one possesses the entailing evidence. Since Andy's evidence entails that she has enough fuel, unless she loses evidence, after she acquires the HOE it will still entail and, hence, fully support this proposition.¹⁵ Even if Andy's evidence only provides strong, but non-conclusive support for her belief, her evidence might still support this belief after she gets the hypoxia warning. Since the evidential relation between Andy's evidence and whether she has enough fuel doesn't depend on any claims about her physiology, adding evidence of hypoxia doesn't change whether her evidence supports the claim that she has enough fuel.

These are the main differences between HOD and other types of defeat that have so far been identified.¹⁶ I do not mean to endorse these claims. Rather than assessing their merits, however, I will argue for a new distinguishing feature, and this argument will not presuppose these claims. But I will return to them in §4 to show that they lend additional, independent support to the suggestion I make in that section. Before I turn to my main argument, it's worth motivating this project.

Most philosophers accept that theories of justification must account for defeat. Often, however, philosophers address defeat simply by adding "No-Defeaters" clauses to their principles. As Schechter explains, "One of the lessons of contemporary epistemology is

¹⁵ Kelly (2013: 47) also notes this probabilistic interpretation of Christensen's argument.

¹⁶ This list is not exhaustive, but other alleged differences are not as compelling. Feldman (2005), for instance, argues that whereas undercutting defeat attacks the connection between the evidence and the conclusion *in a particular case* while leaving the general connection between that evidence and conclusion intact, HOE attacks the *general* connection. Christensen (2010a: 194) convincingly counters this suggestion. Kelly (2005: 187-8) argues that HOE is *agent-relative* in the sense that its epistemic effects depend on *who* you are. On reflection, though, it's not obvious that this feature distinguishes it from other types of defeating evidence. The features discussed in the text, on the other hand, are *prima facie* plausible as distinguishing features of HOD. Thanks to an anonymous referee for motivating me to reconsider agent-relativity.

that epistemic principles should generally include such clauses.”¹⁷ If they *do* make specific proposals about how their principles interact with defeat, these proposals typically suffer from one of two problems: either they fail to account for all and only cases of defeat, or they conflict with the theory’s motivation.

For instance, Alvin Goldman admirably ends his seminal defense of process reliabilism by attempting to account for defeat in his reliabilist framework.¹⁸ Three such attempts, Goldman realizes, either betray the motivations for reliabilism (by not being sufficiently reductive or by being “too internalist”—my phrase not his), or get the wrong results (by implying justification where there should be defeat).¹⁹ And the principle Goldman ends up endorsing is terribly vague.²⁰ If theories of justification ought to account for defeat, and it turned out that process reliabilism could not do so, this would be a major cost of the view. Whether or not reliabilism can in fact handle defeat, my point is that without considering the details, we would never know whether it can.

Philosophers working on HOD have carefully attended to the normative implications of defeat. Let’s consider three examples.²¹ First, recall, Christensen argues that HOE defeats while leaving evidential support intact; that is, acquiring HOE can change

¹⁷ Schechter (2013: 437).

¹⁸ Goldman (1979/1992).

¹⁹ See Goldman (1979/1992: 122) for these three attempts (numbered 7-9). The first version contains the term ‘justified,’ so will not work as a base clause for his recursive account of justification. The second version, Goldman admits, gets the intuitively wrong result. The third version appeals to “meta-beliefs,” which has problematic implications regarding unreflective creatures.

²⁰ Goldman’s final principle: “If S’s belief in *p* at *t* results from a reliable cognitive process, and there is no reliable or conditionally reliable process available to *S* which, had it been used by *S* in addition to the process actually used, would have resulted in *S*’s not believing *p* at *t*, then *S*’s belief in *p* at *t* is justified.” Goldman (1979/1992: 123) recognizes that it’s vague, but defends its vagueness by appealing to the vagueness of our concept of justification. I find this move unsuccessful, but I cannot argue for this here. See Beddor (forthcoming) on process reliabilism’s trouble with defeat.

²¹ These examples do not exhaust the interesting (potential) implications of HOE. For instance, Schoenfield (forthcoming a) suggests that higher-order considerations may be relevant to the evaluation of transitions of thought, rather than the evaluation of belief states. Relatedly, in her (forthcoming b) she argues that considering HOE reveals that there are distinct notions of rationality, corresponding to third-person evaluation and to first-person deliberation.

whether one's beliefs are justified without changing whether one's total evidence supports those beliefs. This directly contradicts *evidentialism*, the view that justification supervenes on evidential support. Thus, if Christensen's arguments are sound, capturing HOD requires abandoning evidentialism.

Second, Christensen also argues that HOE is *rationaly toxic*: once one has it one is often forced into epistemic imperfection because one is bound to violate some epistemic ideal. He argues for this conclusion on a number of grounds,²² but just suppose that his claim that HOD leaves evidential support intact is true. Sometimes, then, agents will have HOE that recommends believing something other than what their evidence supports. Christensen argues that, in spite of this, believing what your evidence supports is an epistemic ideal. But so is respecting one's fallibility by accounting for HOE in one's beliefs. Thus, in a case of HOD, a person who continues to believe what her evidence supports violates the "respect your fallibility" ideal; but, if she believes what her HOE indicates, then she violates the "believe what your evidence supports" ideal. Either way, she violates an ideal. And, Christensen argues, violating these epistemic ideals forces her into epistemic imperfection. Whether this argument is sound depends on whether HOD has the features Christensen attributes to it. Thus, a theory of HOD may have implications concerning whether all epistemic norms are always jointly satisfiable.

Finally, Maria Lasonen-Aarnio argues for several metaepistemological conclusions regarding HOD.²³ Those who take HOD seriously, she thinks, work from two assumptions. First, that for any epistemic rule one can acquire evidence that an attitude that is the output of that rule is actually flawed. And, second, that one should revise one's attitude when one acquires such evidence. She concludes that "there is no non-paradoxical notion of

²² See his (2010a, 2013). At (2013: 92), he makes a logically weaker case for this view.

²³ Lasonen-Aarnio (2014).

justification or epistemic rationality that can accommodate these ideas.”²⁴ She endorses the rationality of epistemic akrasia: one can rationally believe some proposition, while rationally believing that one epistemically should not believe that proposition. This is a surprising conclusion that many are likely to find unacceptable.²⁵ Again, simply appealing to a No-Defeaters clause would conceal the challenges Lasonen-Aarnio raises.

To determine whether standard theories of justification can accommodate defeat in general, and HOD in particular, we need to examine the features of defeat. Although Christensen’s and Lasonen-Aarnio’s assessments of HOE’s implications differ, they converge on the conclusion that an investigation of HOE leads to revisionary results. As I mentioned at the outset, I believe that reaching conclusions about the implications of HOE would be premature. However, the fact that conflicting, *prima facie* cases have been made in favor of this revisionary verdict should make it clear that we need to examine HOD carefully. In the next section, I continue that examination.

3. Object-Independence

In this section, I identify a feature of HOD – what I call *object-independence* – that distinguishes HOD from rebutting and undercutting defeat. From the outset I want to be clear that if it turned out that some cases of HOD don’t have this feature, I wouldn’t take my argument to be undermined. I intend to argue that the *standard* cases of HOD have this feature. In some cases undercutters attack the conclusion, rather than merely attacking the connection between the evidence and the conclusion, but these cases are non-standard, and they don’t show that undercutters should not be distinguished from rebutters in the usual way. Likewise, even if there are cases of object-dependent HOD, this wouldn’t show that HOD

²⁴ Lasonen-Aarnio (2014: 342). Cf. Schoenfield (forthcoming b).

²⁵ For relatively early discussions of the problems associated with epistemic akrasia see Elga (2005) and Feldman (2005). For recent, more detailed articulations of its problems see Greco (2013) and Horowitz (2014).

should not be distinguished from undercutters and rebutters in this way. These cases, I suspect, would be non-standard. Anyway, I will argue that the cases of HOD discussed in this paper have this feature, and that the cases of rebutting and undercutting defeat do not. I believe this argument can be generalized.

To begin to see the object-independence of HOD, consider the following example:

Protective Pessimism: Ron is a Republican, and Don is a Democrat. By Ron's estimate, it's likely, but not *terribly* likely that a Republican will win the next presidential election. By Don's estimate, it's likely, but not *terribly* likely that a Democrat will win.²⁶ When it comes to politics, they both suffer from protective pessimism: they tend to have lower confidence than their evidence supports in political outcomes they find attractive. However, both Ron and Don are completely unaware that they suffer from this affliction.²⁷

Suppose Ron and Don were to learn of their protective pessimism. They would thereby learn that their respective political beliefs likely resulted from a cognitive malfunction or flaw in reasoning, and this would make maintaining their attitudes irrational. If we assume that both of their beliefs were originally justified—perhaps this time they both overcame their pessimism—as I will,²⁸ then this is a case of HOD.²⁹

Intuitively, in response to learning that they are protectively pessimistic, Ron and Don should each become more confident that their preferred candidate will win. Yet, the contents of their beliefs are clearly different: Ron believes *that a Republican will win*, and Don believes *that a Democrat will win*. Thus, whether someone should revise his belief in response

²⁶ For the sake of concreteness, imagine that each has only a 60% credence that his candidate will win.

²⁷ This case is inspired by Christensen's (2010b) Ava/Brayden cases.

²⁸ This assumption is inessential. I'll discuss this below in the text.

²⁹ Don't let the fact that the resulting attitude should be *greater confidence* rather than *less confidence* throw you. Maintaining their previously rational beliefs is irrational as a result of acquiring HOE, so this is a case of HOD.

to learning that he is protectively pessimistic appears to be independent of the content of that belief. I submit that this is a general feature of HOD: whether a subject should revise a belief in response to acquiring some HOE that indicates that that belief resulted from some cognitive malfunction or flaw in reasoning is independent of the object of that belief. In this sense, HOD is *object-independent*.

To say that HOD is object-independent is just to say that *whether* a subject should revise her belief in response to HOE is independent of the content of that belief. If we judge that someone should revise her belief in response to some HOE, then we will make this judgment regardless of whether she believes P or not-P. This does *not* imply anything about what the resulting revision should be. It's worth emphasizing two points. First, we need to distinguish two questions: (1) Should I revise my belief? (2) What should I (now) believe? Second, in cases of HOD, I claim, the answer to the first question is independent of the content of the subject's belief, whereas the answer to the second question may not be. In Protective Pessimism, Ron and Don should both revise. In particular, they should both become more confident. More confident in *what*? The answer to that question depends on how their confidence is distributed; it depends on the objects of their beliefs. Ron should become more confident that a Republican will win, and Don should become more confident that a Democrat will win. *How much more* confident? That depends on lots of things. For instance, when discussing disagreement, Thomas Kelly argues that although both parties to the disagreement should reduce their confidence, the person who originally responded correctly to her evidence need only reduce a little compared to the person who originally bungled the evidence.³⁰ Likewise, how much more confident Ron and Don should become

³⁰ Kelly (2010).

may depend on what their original evidence supports. Or it may not. Saying that HOD is object-independent doesn't bear directly on this issue.

Juxtaposing cases like Ron's and Don's helps bring out the object-independence of HOD. To see that this isn't unique to this case, compare two versions of Hypoxia:

Hypoxia (Original): ...Andy is highly confident that she has enough fuel...Then she acquires *HOE_f*.

Hypoxia (New): ...Andy is highly confident that she *doesn't* have enough fuel... Then she acquires *HOE_f*.

The difference between these cases is the content of Andy's beliefs. In Original, she believes that she has enough fuel; in New, she believes the opposite. Acquiring *HOE_f* seems to call for revision in both cases, in spite of Andy's being confident in one proposition in one case and confident in the negation of that proposition in the other case. It's not that she should revise because of what she believes. Rather, she should revise because, *whatever she believes*, that belief probably resulted from a cognitive malfunction.

If it helps, imagine the following conversation between Andy and Sandy, the air traffic controller who gave Andy the hypoxia warning, and who happens to be an epistemologist.

Sandy: Andy, given our readings, you're probably suffering from hypoxia!

Andy: Well, even if I am, I don't think it matters. I managed to arrive at the belief my evidence supports. I believe...

Sandy: [Cuts Andy off.] Look, I don't care what you believe. You're probably suffering from hypoxia, which means you think you're reasoning as well as usual, but you're not! Whatever you believe about your fuel, you need to reduce your confidence, breathe in some oxygen, and we'll go from there.

If, like Sandy, all we knew about Andy's belief was that she arrived at high confidence about whether she has enough fuel, without knowing which proposition earned her confidence, we would still claim that she should revise that belief. The object of her belief plays no role in the judgment that she should revise, and the contribution of *HOE*, is the same in both cases.³¹

Neither rebutting defeat, nor undercutting defeat is object-independent. If all we knew about Bob and Coco was that they learned, respectively, that there is a black swan and that there is a red light shining on the wall, we would not know whether they should revise their respective beliefs. To know this we would have to know what they already believe. I'll focus on Undercutting, but the same points apply to Rebutting:

Undercutting (Original): ...Coco believes that the wall is red...Then Coco learns that the wall is being illuminated by red lights.

Undercutting (New): ...Coco *doesn't* believe that the wall is red [alternatively: Coco believes that the wall is *not* red]...Then Coco learns that the wall is being illuminated by red lights.

Learning that the wall is being illuminated by red lights seems to call for revision in Undercutting (Original), but not Undercutting (New). In the latter case, the new information may not bear at all on her beliefs or it may provide some confirmation. Either way, acquiring this information calls for revision in Undercutting (Original), but not in Undercutting (New). And the only relevant difference between the cases is the content of Coco's beliefs.

Again, if it helps, imagine a conversation between Coco and her epistemologist friend Dodo.

³¹ We see basically the same thing in Disagreement, except because it's an impure case of HOD, to factor out the rebutting effect and to keep the HOE the same across both cases, we have to specify the HOE as saying something like, "my peer believes the opposite of what I believe."

Coco: Dodo, I was previously confident about whether the wall is red. I'm wondering, since I just learned that it's being illuminated by red lights, should I revise my belief?

Dodo: Well, that depends. If you were pretty confident that the wall is red, then you should definitely revise your belief. If you were already confident that the wall is not red, then there's no need to revise.

Dodo's response seems reasonable. We cannot know whether Coco should revise her confidence without knowing what she already believes.

It is unsurprising that rebutting and undercutting defeat are object-dependent when we consider their standard characterizations. Rebutting defeaters "attack the conclusion," while undercutting defeaters "attack the connection between the evidence and the conclusion." A rebutting defeater for a proposition is evidence that this proposition is false; typically, the same piece of evidence will not indicate (for people with similar evidence) that P and not-P are both false. And, regarding undercutters, there is no reason to suppose that a single piece of evidence will attack both the connection between evidence and a conclusion and the connection between evidence and the negation of that conclusion.

When we reflect on examples of HOD, the object-independence of HOD, if not unsurprising, at least makes sense. When Ron and Don learn that they are protectively pessimistic, they learn that they tend to be underconfident on certain matters. When Dale discovers that Donna disagrees with him, from his perspective, he learns that "my peer believes the opposite of me." When Andy learns that her belief probably resulted from hypoxia, she learns that her belief probably resulted from a distorting influence. At bottom, each of these subjects learns the same thing: whatever I believe regarding some issue, *that belief* is flawed. The defeating evidence is not tied to the content of any particular belief.

Someone may object by presenting the following dilemma. Horn 1: the pairs of cases don't both involve defeat. Horn 2: the evidence differs between the cases. The dilemma can be motivated by two thoughts. First, defeat requires defeated beliefs to be justified initially. Second, a single body of evidence can justify only one of a pair of inconsistent beliefs. Given these two thoughts, either the evidence is the same in both cases and both beliefs are not initially justified, or both are justified and the evidence differs. If Horn 1 is true, then we can't compare the cases to learn lessons about HOD (because they are not both cases of defeat). If Horn 2 is true, then there's another difference between the cases besides the object of the subjects' beliefs. Either way, the argument doesn't show that HOD is object-independent.

This objection appeals to the denial of *permissivism*. According to permissivism, two people with the same evidence can have justified beliefs in opposing propositions. But this denial is controversial, and if my view relied on permissivism, that wouldn't be a huge cost.³² However, the dilemma doesn't essentially rely on this denial. The dilemma could be run by pointing to particular cases, claiming that *those* cases are impermissible. This will be more or less plausible depending on the case. Protective Pessimism is a paradigmatic permissive case (according to permissivists), but Hypoxia is not plausibly permissive as Andy's evidence entails her belief.

Nevertheless, this objection misses the point for two reasons.³³ First, the argument for the object-independence of HOD doesn't depend on there being *no* differences between the pairs of cases aside from the content of the subjects' beliefs. Admittedly, the case for object-independence is *stronger* if that is the only difference. But even if there are other

³² One of the seminal (recent) works on this issue is White (2005). For critical discussion of that paper, see Meacham (2014). See, also, Schoenfield (2014) for a defense of permissivism.

³³ I mention it because I anticipate that those working on the issues discussed in this paper might be inclined to think in these terms.

differences, if they don't explain why HOE appears to be object-independent, then these differences are irrelevant in this context.

Second, and more importantly, my argument identifies a difference between the defeating effect of HOE, on the one hand, and that of rebutting and undercutting evidence, on the other. Regardless of whether Horn 1 or Horn 2 is true, the argument appealing to the pairs of cases does this. In both cases, acquiring the hypoxia warning provides Andy with a reason to revise her belief. The corresponding claim is not true of Rebutting and Undercutting. That is the observation driving my argument, and it remains regardless of how we fill out the cases.

This point can also be made by “counting mistakes.”³⁴ Suppose that in Hypoxia (New) Andy's original evidence is the same as her evidence in Hypoxia (Original). In both cases, then, Andy's evidence supports high confidence in the proposition that she has enough fuel. Now consider only Hypoxia (New). Suppose that after she gets the hypoxia warning, she remains highly confident in her belief. In this scenario, Andy has made two mistakes. She didn't believe what her original evidence supported, and now she fails to respond appropriately to the HOE. The latter mistake is the same mistake Andy would make by failing to revise in Hypoxia (Original). In the analogous setup, in Undercutting (New) Coco makes only one mistake. Even if she believes that the wall is not red when her original evidence supports the proposition that the wall is red, maintaining her belief after learning that the wall is being illuminated does not constitute an additional mistake.

4. Upshot

In this section, I speculate on what implications the object-independence of HOD has for a theory of rational belief. I argue that both on its own, and in conjunction with the other

³⁴ See Schechter (2013: 433) for a similar argumentative strategy.

distinguishing features described in §2, the object-independence of HOD favors an account of HOD grounded in non-evidential, “state-given” reasons for belief. If this is correct, then an account that grounds HOD in state-given reasons has the potential to explain its peculiarities. This explanatory payoff is what makes this account significant.³⁵

Metaethicists distinguish between two sorts of reasons: state-given reasons and object-given reasons.³⁶ Using this distinction (to be explained presently), we can state two hypotheses about HOD:

SG-Hypothesis: HOD defeats by providing subjects with state-given epistemic reasons against their beliefs.

OG-Hypothesis: HOD defeats by providing subjects with object-given epistemic reasons against their beliefs.

I will suggest that the object-independence of HOD favors the SG-Hypothesis over the OG-Hypothesis.

Consider the proposition that God exists; call it *G*. The Teleological argument provides one reason to believe *G*: roughly, God’s existence is the best explanation of the apparent design we observe. This argument (suppose) gives us reason – evidence – to believe that *G is true*. Another reason to believe *G* is provided by Pascal’s Wager: roughly, there is extreme disvalue associated with not believing *G*. Pascal’s Wager doesn’t purport to provide *evidence* that *G* is true, but it is meant to provide reason to *believe* *G*. These are different kinds of reasons for belief. We can refer to the first kind of reason as an *object-given reason* and the

³⁵ The following argument provides only a sketch, and along the way, I flag claims that need additional investigation.

³⁶ See, e.g., Parfit (2011) and Schroeder (2012).

second kind as a *state-given reason*.³⁷ This distinction, I will assume, is jointly exhaustive: if a reason isn't object-given, then it is state-given (and vice versa).

Although my main focus in this section is metaepistemological, it's worth briefly mentioning some metaethical implications of the SG-Hypothesis. Derek Parfit is skeptical of the significance of state-given reasons because, he contends, state-given reasons are redundant given object-given reasons; since we of course need to appeal to object-given reasons, it's not worth discussing state-given reasons.³⁸ Mark Schroeder defends the significance of state-given reasons against the prevailing orthodoxy that these reasons are the "wrong kind of reason," by arguing, *inter alia*, that we have state-given epistemic reasons *against* believing and disbelieving, in favor of *withholding judgment* – an attitude that has no object, and hence, that we cannot have object-given reasons to adopt.³⁹ Schroeder's argument for epistemic state-given reasons, however, is limited to reasons to withhold. If capturing the object-independence of HOD does indeed require state-given reasons, as I argue, then that is reason to reject Parfit's skepticism and to extend our judgment of the significance of state-given reasons beyond the limits of Schroeder's argument. It also provides *independent* reason to reject the prevailing orthodoxy that state-given reasons are the wrong kind of reasons for belief, since the reasons provided by HOD are clearly of the right kind. Thus, if the SG-Hypothesis is true, it has several metaethical ramifications in addition to its epistemological implications.

Return now to the epistemological domain. As our focus here is on *defeat*, we are interested in reasons to *revise* one's beliefs. Reasons to revise a belief must be grounded in

³⁷ It may be that some state-given reasons are also object-given reasons. For the value/disvalue associated with the state that grounds a particular state-given reason may be located in the object of that state. In what follows, I focus on *mere* state-given reasons. I will leave the "mere" implicit.

³⁸ Parfit (2011: 51). See also Appendix A in that volume. The "it's not worth" claiming or discussing locution applied to state-given reasons is Parfit's.

³⁹ Schroeder (2012).

(apparent) facts about the disvalue of maintaining that belief; if there were nothing wrong with holding that belief, there would be no reason to revise it. Object-given reasons to revise a belief locate the disvalue of that state in its content. In the case of object-given *epistemic* reasons, that disvalue is linked to the (probable) falsity of the belief's content.⁴⁰ State-given reasons to revise a belief do not locate the disvalue of that state in its content; there may be something about holding that belief, independent of its content, that makes maintaining it bad, and this badness gives rise to a state-given reason against maintaining it. Thus, whereas object-given reasons to revise a belief do depend on the belief's content, state-given reasons to revise do not.

Object-given epistemic reasons, then, are the same as evidence, and the OG-Hypothesis is a standard evidentialist picture of HOD. State-given reasons, on the other hand, are not evidence, and hence, the SG-Hypothesis is an anti-evidentialist picture. Several prominent philosophers take HOE to be “just more evidence,” and, accordingly take HOD to be amenable to an evidentialist picture like the OG-Hypothesis.⁴¹ An argument for the SG-Hypothesis, therefore, is an argument against this prominent approach to HOD.

Here is my main argument. Higher-order defeaters provide subjects with reasons to revise their beliefs. Since HOD is object-independent, it provides subjects with reasons to revise their beliefs that do not depend on the content of those beliefs. This is why the object-independence of HOD favors the SG-Hypothesis over the OG-Hypothesis: the reasons provided by HOD cannot be object-given reasons if HOD is object-independent, so they must be state-given reasons.

⁴⁰ In particular, the reason to revise may be that the object of the belief is *false* or that it is *less likely to be true*, and hence, *more likely to be false* than previously thought.

⁴¹ See, for instance, Feldman (2009) and Kelly (2010). In addition, Lasonen-Aarnio (2014) admits that she is “very sympathetic to the thought that higher-order evidence should be treated as just more evidence,” though she disagrees with Feldman and Kelly about the implications this thought has for an account of HOD.

This argument raises at least three questions:

1. The reasons provided by HOD are clearly *epistemic*, but Pascal's Wager provides a *pragmatic* reason. If this is the paradigmatic case of a state-given reason, does this mean there is no precedent for epistemic state-given reasons, making the SG-Hypothesis is a non-starter?
2. If the SG-Hypothesis is a "non-evidentialist" picture of HOD, what role does higher-order *evidence* play?
3. If state-given reasons depend on content-independent disvalue, what is the disvalue that gives rise to the state-given reasons that ground HOD?

In what follows, I address each question in turn. Then I conclude the section by replying to a more general worry, arguing that the other distinguishing features of HOD described in §2 also favor the SG-Hypothesis.

Regarding Question 1, the first thing to note is that there may be examples of state-given *epistemic* reasons unrelated to HOD. Consider: while you're waiting for a medical test result, your expectation that you'll soon receive further evidence provides you with an epistemic reason to suspend judgment, even though this fact is not evidence either way regarding the result.⁴² Or another: the fact that P entails Q is an epistemic reason to avoid the state of believing P and disbelieving Q, but this sort of reason isn't clearly evidence either.⁴³ Or finally: the fact that two hypotheses are on a par with respect to all theoretical virtues except that one is simpler than the other provides a reason to believe the simpler one. But, again, this is not obviously evidence since simplicity is non-truth-conducive. Moreover, the (epistemic) value of adopting belief in a simpler hypothesis is, one might argue, that it increases coherence of one's total belief state. So, this reason is arguably state-given rather than object-given because the value associated with adopting such a belief derives not from

⁴² This example is from Schroeder (2012).

⁴³ This example is from Sosa & Sylvan (forthcoming).

the object of the belief itself, but from a feature of the total belief state.⁴⁴ These seem to be examples of state-given epistemic reasons for/against certain mental states.⁴⁵ If that is correct, then there is a precedent for admitting state-given epistemic reasons into one's normative framework.

Setting this precedent would certainly bolster the case for the SG-Hypothesis, and if it turns out to be a contender, this would be a step to take in future research. However, even if no precedent could be set, this wouldn't doom the SG-Hypothesis. HOD differs from other kinds of defeat. Depending on how different HOD is, it may be completely unsurprising that there is no precedent. Recall the revisionary results suggested by Christensen's and Lasonen-Aarnio's arguments, that HOE is rationally toxic and that there is no non-paradoxical notion of justification that can accommodate HOD. No one thinks rebutting or undercutting defeat have implications like this. It may turn out that the name of the HOD game is "unprecedented."

Question 2: What is the relationship between HOE and HOD on the SG-Hypothesis, a "non-evidentialist" picture? Answer: HOE provides *evidence of the epistemic disvalue* of maintaining a belief, and having evidence of this disvalue *gives rise to a state-given epistemic reason* against maintaining this belief. In general, when a subject has evidence that there is some value or disvalue associated with being in some state *S* – besides having a true/false object – this gives her a state-given reason to adopt or give up *S*. For instance,

⁴⁴ I'm grateful to an anonymous referee for reminding me of this potential precedent. Lycan (1988), for instance, defends a version of coherentism according to which what justifies a new belief is that it increases explanatory coherence among one's set of beliefs, and whether one hypothesis is a better explanation than another depends on its possession of theoretical virtues. Lycan admits that some of these virtues are not truth-conducive.

⁴⁵ Whether these examples succeed may depend on complex issues about the nature of belief and suspension of judgment, issues that I cannot address here. For instance, if suspension of judgment has an object, then perhaps the first example is an example of an object-given reason; or, if believing *P* and disbelieving *Q* is just to believe the conjunction *P* & ~*Q*, then perhaps the second example is too. As I argue in the next paragraph in the text, if they don't happen to be examples of state-given reasons, the main argument of this section would only be slightly weakened. Thanks to an anonymous referee for urging me to address this issue.

suppose Ethel is suffering from cancer, and her doctor informs her of studies showing that patients who believe they will recover tend to recover more frequently than those who don't. Ethel has evidence (provided by the doctor's testimony) that there is value associated with believing that she will survive. And this gives rise to a (pragmatic) state-given reason to believe she will survive. On the SG-Hypothesis, the same applies to HOE. When Andy gets the hypoxia warning, she acquires HOE that there is some epistemic disvalue associated with her belief. This gives rise to a state-given reason against holding this belief. Thus, although HOD, like all defeat, can ultimately be traced to evidence, what *explains* HOD is not that the subject's evidence no longer *supports* the relevant belief—the evidence may still support this belief. Rather, it's that the evidence indicates that there is content-independent disvalue associated with maintaining the belief.

Question 3: What is the alleged, object-independent epistemic disvalue to which HOE points that gives rise to the state-given reasons that ground HOD? I don't have a full account to offer here, but there are at least three *prima facie* plausible answers that come to mind: When Andy gets the hypoxia warning, she simultaneously gains evidence (i) that her belief is not supported by her evidence, (ii) that her belief is epistemically irrational, and (iii) that it does not constitute knowledge. And there is epistemic disvalue associated with each of (i)-(iii), disvalue that needn't depend on the belief's object, and that may be enough to warrant revision.⁴⁶

The SG-Hypothesis differs in important respects from other accounts of HOD. Its appeal to non-evidential, state-given epistemic reasons makes it anti-evidentialist, unlike

⁴⁶ That HOE provides evidence of (i) is a central premise in the debate over the rationality of epistemic akrasia (see fn. 24 above for references). Christensen (2010b) also makes this suggestion. See the references in fn. 10 for (ii). Since (i) and (ii) are plausibly necessary conditions for (iii), the arguments favoring (i) and (ii) favor (iii).

Richard Feldman's and Thomas Kelly's accounts.⁴⁷ It also has an advantage over *level-bridging accounts* of HOD. On these accounts, the rationality of a subject's lower-order belief gets defeated solely as a result of the subject having a higher-order belief that something is epistemically wrong with the lower-order belief. In Andy's case, for example, HOD requires two steps: first, after acquiring the hypoxia warning, she should become confident that something is wrong with her first-order belief; second, as a result of adopting this confidence, and *only* as a result of this, she should revise her first-order belief. If the HOE fails to make her confident that something is wrong with her first-order belief, then it does not affect the rationality of that belief.⁴⁸ This, however, strikes me as the wrong result. The intuition that Andy should revise her first-order belief once she receives the HOE is present without knowing whether she changes her higher-order beliefs. Even if Andy thinks nothing is wrong with her first-order belief, maintaining this first-order belief is irrational because her evidence indicates that something *is* epistemically wrong with this belief.⁴⁹ The SG-Hypothesis gets the right result here. As explained in response to Question 2, HOD occurs when a subject has a (sufficiently strong) state-given reason against her belief, and such a reason arises when she has (sufficiently strong) HOE that there is something epistemically wrong with that belief. So, on the SG-Hypothesis once Andy acquires the HOE, the rationality of her first-order belief is defeated regardless of whether she changes her mind about the rationality of that belief.

The SG-Hypothesis also avoids the potentially problematic implications associated with the views of Lasonen-Aarnio and Christensen (§2). In providing an account of HOD, it

⁴⁷ Feldman (2009) and Kelly (2010).

⁴⁸ Three notes. First, this obviously does not imply that level-bridging accounts cannot criticize the resulting belief-state. Second, Christensen (2013: 91) seems to endorse a level-bridging account. Third, in articulating this account, I'm echoing David Lewis's (1986: 85) discussion of the relation between evidence and beliefs about chance and outcomes.

⁴⁹ Cf. Schechter (2013: fn. 35).

avoids the claim, made by Lasonen-Aarnio and others⁵⁰, that accommodating HOE leads to rational epistemic akrasia. It can also avoid claiming, as Christensen does, that accommodating HOE forces one into epistemic imperfection. Recall, Christensen argues for this claim by showing that once one has HOE one is often bound to violate an epistemic ideal.⁵¹ But, from the perspective of the SG-Hypothesis – a perspective that takes epistemic *reasons* seriously – the step from violating ideals to epistemic imperfection is questionable. In the case of action, for instance, not lying may be an ideal and avoiding causing harm may be another ideal, but when these ideals conflict, no imperfection must result. This is because the ideals should not be seen as exceptionless rules, but as being or grounding different contributory reasons, which should be weighed against each other to yield an overall recommendation.⁵² Likewise, on the SG-Hypothesis, it may be true that there are conflicting epistemic ideals (e.g., “respect your evidence” and “respect your fallibility”), but violating one in favor of the other needn’t lead to any epistemic imperfection because each ideal provides a different reason that should be weighed against the other to yield an overall recommendation. Just as we needn’t accuse an agent of being morally imperfect for lying to avoid harm in spite of his violating an ideal against lying, we shouldn’t accuse an epistemic agent of being epistemically imperfect for violating a “respect your evidence” ideal in favor of a “respect your fallibility ideal,” in the right circumstances. And those circumstances will be ones in which the agent’s state-given epistemic reasons against her belief outweigh her object-given epistemic reasons for that belief.⁵³

⁵⁰ See Horowitz’s (2014) excellent overview of the work that recommends epistemic akrasia.

⁵¹ See Christensen (2010: 192-3, 204) and (2013: 95).

⁵² See Dancy (2004: 3-12).

⁵³ Admittedly, this relies on an account of weighing object-given and state-given reasons together, and I have no such account to offer here. But insofar as the distinction between these kinds of reasons is one we already needed, my appeal to this weighing relation brings with it no new commitment. However, this is definitely an element of the account that requires further development.

Finally, one might insist that a standard evidentialist/object-given framework can account for object-independence. Both Ron and Don have strong evidence that their belief resulted from a cognitive malfunction. Since this evidence isn't tied to any particular content, it is not object-dependent, and when a person has strong evidence of this sort, her evidence no longer supports that belief. What's the problem?

How plausible this line is partly depends on whether HOD has the other features described in §2 that allegedly distinguish it from undercutting defeat. This is because the SG-Hypothesis appears to have the explanatory advantage over the OG-Hypothesis with respect to those features. While the OG-Hypothesis is *consistent* with the retrospective aspect of HOD, the SG-Hypothesis actually makes *predictions* that cohere with it, and the OG-Hypothesis makes no such predictions. On the SG-Hypothesis, it's not surprising that HOE provides evidence that the defeated belief was epistemically irrational, as believers in the retrospective aspect claim. This hypothesis predicts that HOE provides subjects with evidence that there is some object-independent epistemic disvalue associated with the defeated belief. And evidence that this belief is epistemically irrational is arguably disvalue of that kind. On the OG-Hypothesis, the retrospective aspect is surprising, as no other defeating evidence works like this.

Things are worse for the OG-Hypothesis regarding the features identified by Christensen and Schechter because those features recommend an anti-evidentialist account of HOD. Thus, as an *evidentialist* account, the OG-Hypothesis is not even *consistent* with HOD's having these features. Moreover, it's worth noting that my arguments that HOD is object-independent and that this favors the SG-Hypothesis do not rely on claiming that HOD has the features identified by Christensen and Schechter. Yet, our arguments converge on an anti-evidentialist approach. This independent convergence provides additional support

to the SG-Hypothesis over the OG-Hypothesis. More than this, though, the SG-Hypothesis is not only a negative picture of HOD. Rather, it provides a positive proposal that explains why evidentialist accounts of HOD fail: capturing HOD requires state-given epistemic reasons – reasons an evidentialist will not countenance.

If HOD does not leave evidential support intact, or if it does not have a retrospective aspect, *and* if a plausible story explaining the object-independence of HOD on the OG-Hypothesis can be told, object-independence may not favor the SG-Hypothesis over the OG-Hypothesis. If, on the other hand, HOD has either of these other features, or if the OG-Hypothesis cannot explain object-independence, then the object-independence of HOD may indeed favor the SG-Hypothesis. This is partly why it is too early to tell what an account of HOD will look like. Whether certain alleged facts are data an account of HOD must capture is still unclear. Once we settle which data must be captured, as I have urged that we need to do and have started to do in this paper, we will be in a better position to provide a comprehensive account of HOD.

5. Conclusion

The defeat distinctness hypothesis is true: higher-order defeat is distinct from rebutting defeat and undercutting defeat. What an adequate theory of epistemic norms looks like depends on what distinguishing features higher-order defeat has. These features, therefore, must be studied and assessed.

I have identified a feature of higher-order defeat that distinguishes it from other kinds of defeat, namely, object-independence: whether a person should revise a belief in response to acquiring higher-order defeaters is independent of the object of that belief. Comparison of cases and reflection on the different types of defeat both suggest that higher-order defeat, but not these other types, has this feature.

Finally, I have gestured toward what a theory of rational belief that captures the object-independence of higher-order defeat might look like: a hybrid theory, essentially appealing to both object-given and state-given epistemic reasons. If everything I've argued here is correct, much more still needs to be known about higher-order defeat to adequately defend this theory. Even if object-independence does not *distinguish* higher-order defeat from other types of defeat, if *any* defeat is object-independent, the proposal that capturing object-independence requires state-given reasons may still be correct, and a hybrid theory may still be needed. In any case, a closer look at these details will advance our understanding of higher-order defeat and, thereby, advance our understanding of rational belief.

References

- Beddor, B. (forthcoming). Process Reliabilism's Trouble with Defeat. *The Philosophical Quarterly*.
- Christensen, D. (2010a). Higher-Order Evidence. *Philosophy and Phenomenological Research* 81: 185-215.
- Christensen, D. (2010b). Rational Reflection. *Philosophical Perspectives* 24: 121-140.
- Christensen, D. (2013). Epistemic Modesty Defended. In Christensen, D. & Lackey, J. (eds.) *The Epistemology of Disagreement: New Essays*. New York: Oxford University Press. 77-97.
- Dancy, J. (2004). *Ethics without Principles*. Oxford: Oxford University Press.
- Elga, A. (2005). On Overrating Oneself and Knowing It. *Philosophical Studies* 123: 115-124.
- Elga, A. (2013). The Puzzle of the Unmarked Clock. *Philosophical Studies* 164: 127-139.
- Feldman, R. (2005). Respecting the Evidence. *Philosophical Perspectives* 19: 95-119.
- Feldman, R. (2009). Evidentialism, Higher-Order Evidence, and Disagreement. *Episteme* 6: 294-312.
- Goldman, A. (1979/1992) What is Justified Belief? In Pappas, G., ed., *Justification and Knowledge*. Reprinted in *Liaisons* (1992). Cambridge: MIT Press. 105-126.
- Greco, D. (2013) A Puzzle about Epistemic Akrasia. *Philosophical Studies* Published Online: 1 January 2013.
- Horowitz, S. (2013) Epistemic Akrasia. *Nous* Published Online May 6 2013. 1-27.
- Kelly, T. (2010). Peer Disagreement and Higher Order Evidence. In Feldman, R. & Warfield, T. (eds.) *Disagreement*. New York: Oxford University Press. 111-174.
- Kelly, T. (2013). Disagreement and the Burdens of Judgment. In Christensen, D. & Lackey, J. (eds.) *The Epistemology of Disagreement: New Essays*. New York: Oxford University Press. 31-53.
- Lasonen-Aarnio, M. (2014). Higher-Order Evidence and The Limits of Defeat. *Philosophy and Phenomenological Research* 88: 314-345.
- Lewis, D. (1986). A Subjectivist's Guide to Objective Chance. *Philosophical Papers*, Vol. II. Oxford: Oxford University Press. 83-132.

- Lycan, W. (1988). *Judgment and Justification*. Cambridge: Cambridge University Press.
- Meacham, C. J. G. (2014). Impermissive Bayesianism. *Erkenntnis*, 79: 1185-1217.
- Parfit, D. (2011). *On What Matters Volume 1*. Oxford: Oxford University Press.
- Pollock, J. (1986). *Contemporary Theories of Knowledge*. Totowa, NJ: Roman Littlefield.
- Schechter, J. (2013). Rational Self-Doubt and the Failure of Closure. *Philosophical Studies* 163: 429-452.
- Schoenfield, M. (2014). Permission to Believe: Why Permissivism is True and What it Tells Us about Irrelevant Influences on Belief. *Nous* 48: 193-218.
- Schoenfield, M. (forthcoming a.) A Dilemma for Calibrationism. *Philosophy and Phenomenological Research*.
- Schoenfield, M. (forthcoming b.) Bridging Rationality and Accuracy. *The Journal of Philosophy*.
- Schroeder, M. (2011). Holism, Weight, and Undercutting. *Nous* 45: 328-344.
- Schroeder, M. (2012). The Ubiquity of State-Given Reasons. *Ethics* 122: 457-488.
- Sosa, S. & K. Sylvan (forthcoming). The Place of Reasons in Epistemology. In Star, D. (ed.) *The Oxford Handbook of Reasons and Normativity*.
- Titelbaum, M. (2015). Rationality's Fixed Point (Or: In Defense of Right Reason). *Oxford Studies in Epistemology* vol. 5. Oxford: Oxford University Press.