

Undermining Defeat and Propositional Justification

Giacomo Melis

University of Glasgow

Abstract

I extend the Higher-Order View of Undermining Defeat (*HOVUD*) defended in Melis (2014) to account for the defeat of propositional justification. In doing so, I also clarify the important notion of higher-order commitment, and I make some considerations concerning the defeat of externalist epistemic warrants.

Keywords: Defeaters, Undermining, Overriding, Propositional Justification, Higher-Order.

One of the things that epistemologists worry about is the relation of epistemic support, in virtue of which something gives a subject reason to believe that a proposition is true. While the notion of epistemic support might be accounted for in a variety of ways, for the sake of simplicity in what follows I will assume that what gives a subject reason to believe that a proposition p is true is the possession of a piece of evidence that supports p . We can thus say that evidence for p is a justifier for p , in that it provides the agent with justification to believe that p is true.

If justifiers—evidence—provide the epistemic agent with justification, defeaters take justification away from her. Just like justifiers, defeaters have epistemic force, but it is a force that speaks against believing a proposition, rather than in favour of it. For the purposes of this paper, let us understand defeaters as pieces of counterevidence—evidence that speaks against believing a proposition.

Since Pollock (1974: 42-43), defeaters are commonly distinguished in at least two different kinds. Say that p is a previously justified proposition for a subject S : *overriding* defeaters give S a reason to believe not- p ; *undermining* defeaters, on the other hand, give S merely a reason to give up p , without thereby giving a reason to believe not- p . The present article develops a previous contribution to understanding the way in which undermining defeaters work.

In Melis (2014), I defended a view according to which undermining defeaters require the subject to engage in some higher-order epistemic reasoning, while overriding defeaters do not. One limit of that account was that it applied only to a doxastic notion of justification and defeat. In this paper I extend the proposal

to propositional justification and defeat, and I elaborate the crucial notion of higher-order commitments. The plan is the following: first (§ 1) I will present the view already defended and (§ 1.1) I will provide some details about the notion of higher-order commitment, then (§ 2) I will briefly recall the distinction between propositional and doxastic justification and show why the account of undermining defeat presented does not apply to propositional justification. Finally (§ 3), I will extend the account to propositional justification.

1. The Higher-Order View of Undermining Defeat

Let us begin by considering some examples to illustrate the difference between undermining and overriding defeaters. In the propositional triads below, e represents the evidence that supports p , p the proposition the subject is justified in believing, and d the defeater.¹

- (1) $e = \langle \text{Adam says that Paul McCartney was the drummer of The Beatles} \rangle$
 $p = \langle \text{Paul McCartney played in the Beatles} \rangle$
 $d = \langle \text{Lauren tells me that Adam's knowledge in matters of pop music is poor} \rangle$
- (2) $e = \langle \text{I remember having left the book on the desk} \rangle$
 $p = \langle \text{The book is on the desk} \rangle$
 $d = \langle \text{I now see that the book is not on the desk} \rangle$
- (3) $e = \langle [S\text{'s apparent proof of } p] \rangle$
 $p = \langle [A\text{ seemingly logical theorem}] \rangle$
 $d = \langle \text{A logician tells } S \text{ that there is a mistake in the (apparent) proof} \rangle$
- (4) $e = \langle \text{All swans observed at } t_1 \text{ are white} \rangle$
 $p = \langle \text{All swans are white} \rangle$
 $d = \langle \text{At } t_2 \text{ a black swan is observed in Australia} \rangle$

The defeater d in cases (2) and (4) explicitly suggests that not- p , while in cases (1) and (3) it is compatible with the truth of p . In (2) and (4) the defeaters are overriders, in (1) and (3) they are underminers.

Taking the inspiration from different remarks made by Scott Sturgeon (2014) and Albert Casullo (2003: 45-46), I defended the following Higher-Order View of Undermining Defeat (*HOVUD*):

(HOVUD) Underminers suggest that something was wrong with the source of justification or with the justificatory process, and they operate their defeat by appealing to the higher-order commitment that the belief in question was based on that source or that process. If the suggestion is that the process, rather than the source, was defective, the defectiveness is to be understood as the occurrence of a mistake or some other disturbing event.²

In order to assess the explanatory power of the view, we need to go back to examples (1)-(4). Before doing that, however, I will make a few clarificatory remarks about some of the notions I have appealed to in (*HOVUD*).

¹ I do not mean to suggest that all evidence (or counter-evidence) is propositional, but only that for every piece of evidence, there is a proposition that can be used to represent it.

² A refinement of (*HOVUD*) proposed in fn. 16 of Melis 2014 has it that the suggestion made with respect to the source or the process can also be that they *might have been* defective.

By ‘source of justification’ I mean something of a rather large scale, like the agent’s five senses, her memory, her reasoning abilities, or someone’s testimony. Justificatory sources provide the evidence for p , and deliver justificatory processes. By ‘justificatory process’ I mean the activity that begins with the gathering of the evidence and that ends with the formation of the belief. The stage of evidence-gathering often involves interaction with things external to the mental life of the subject, but overall, justificatory processes are mental affairs internal to the agent.³

Higher-order commitments are potential or actual cognitive attitudes (*e.g.* beliefs, acceptances, presuppositions) of the agent towards propositions about the ways in which beliefs are justifiably formed, retained and abandoned—propositions that describe and sustain the justification of the agent’s doxastic state (the set of the agent’s beliefs) at a time t .⁴ In general, every epistemically respectable change in the set of the agent’s beliefs involves some higher-order commitments.⁵ The propositions towards which the *relevant agent* is committed are such that she would take them to be true (or just warranted) on reflection—at least for as long as she stands by the related piece of justification.⁶ We might say that, in a loose sense, the higher-order commitments are part of the agent’s justification, in that the agent could resort to them to defend the epistemic worthiness of her belief, if questioned. However, the agent need not be aware that the relevant commitments are in place in order to form, retain, and abandon beliefs in an epistemically worthy manner.

Let us consider an example. When I come to believe that it is ten o’clock by looking at my watch, I thereby form a justified belief which is sustained by commitments towards propositions like ‘the belief that it is ten o’clock was formed by looking at my watch’, ‘my watch is reliable’, or ‘I would have not trusted my watch if I had a good reason to believe it had stopped working’, and so on. However, I do not need to appreciate that such commitments are in place (*i.e.* that I am committed to take those higher-order propositions as true or warranted)⁷ in order to be justified in believing that it is ten o’clock. Still, if some exigent interlocutor were to push me to lay down my reasons to believe that it is ten o’clock, I would, maybe after some reflection, probably bring up some of the higher-order commitments.

We can now go back to (*HOVUD*), according to which the phenomenon of undermining defeat is articulated in two distinct parts. Firstly, the underminer needs a certain higher-order commitment about the source of justification or about the justificatory process to be in place; secondly, the underminer challeng-

³ That is not to say that the subject needs to be aware of the mental activity involved in the justificatory process.

⁴ In so far as such propositions overtly describe the epistemic activity of the subject—as opposed to merely implicitly expressing the subject’s engagement in it—they are higher-order propositions.

⁵ Exactly which ones depends on the details about the source of justification and justificatory processes involved.

⁶ More on the rather important notion of relevant agent below in § 1.1.

⁷ Details here depend on the specific cognitive attitude. For example, while believing would require taking the relevant higher-order propositions as true, accepting and presupposing might require only taking them as just warranted, or at any rate something which generally is not truth-guaranteeing.

es the epistemic worthiness of that very commitment. Let us consider the examples.

The commitment that has to be in place for the undermining defeat to be effective in example (1) is a commitment about the source of justification, the relevant proposition being: ‘the belief that Paul McCartney played in the Beatles was based on Adam’s testimony’. That the commitment has to be in place so that *d* can do its defeating job can be seen by noting that if I had taken my belief to have been justified by a different source—maybe someone else’s testimony—the acquisition of the information that Adam’s knowledge in matters of pop music is poor would not have had any defeating effect.⁸ That the epistemic worthiness of the commitment is challenged by *d* can easily be seen by considering the very content of *d*, whose general suggestion is that the source of justification is not reliable with respect to the subject matter.

By contrast, nothing of the sort goes on in the overriding defeat involved in examples (2) and (4). In both cases, the defeater *d* could have done its defeating work just as well if the relevant agent had taken the beliefs to have been justified by any other source rather than memory (case 2), or inductive reasoning (case 4). Thus, there is no need for a higher-order commitment concerning the source of justification or the justificatory process⁹ to be in place so that the defeat can be effective. Such examples suggest that overriding defeat, in general, does not rely on the presence of commitments about the way in which the relevant belief was formed or justified.

Case (3) is an example of undermining defeat in which the challenge raised by the underminer concerns the justificatory process rather than the justificatory source. In the case at hand, the justificatory source is provided by the agent’s proving abilities, as it were. Such source delivers the alleged proof that *p*, whose execution on the subject’s part constitutes the relevant justificatory process. The commitment that has to be in place and that is challenged in case (3) concerns the following proposition: ‘the belief that *p* was based on the execution of that specific proof’. If the subject had taken her proving abilities to have delivered a different computation, *d* would not have had any undermining effect. The epistemic worthiness of the commitment is challenged because *d* suggests that something went wrong in that justificatory process.¹⁰

The preceding paragraph offers the occasion for an interesting reflection about the difference in the way underminers and overrides impinge on higher-

⁸ This example should also help to clarify what it means for an underminer to *appeal* to some higher-order commitments: in general, an underminer *d* appeals to some higher-order commitment *c* when *d* could not do its defeating job, unless *c* was in place (that is, unless the relevant agent took *c* to be true or warranted on reflection).

⁹ Since justificatory processes are delivered by justificatory sources, if a defeater is effective regardless of the source of justification, it is to be expected that it is effective also regardless of the specific justificatory process delivered by the source.

¹⁰ One might think that the underminer challenges the epistemic worthiness of the source *via* the challenge raised to the justificatory process. That might happen in some cases, but it is easy to conceive of a scenario in which that is not what happens in example (3). Just suppose that the subject is generally good at proving theorems, and that she executed the calculation in optimal circumstances (she was sober, in a quiet room, etc.). Since the occasional mistake is compatible with the trustworthiness of the source, the challenge raised by *d* in such a construction of case (3) does not extend from the process to the source. See Melis (2014: §§ 2-3) for more details.

order reasoning. Indeed, there is a sense in which overrides suggest that something went wrong with the justificatory process too. Overrides suggest that the previously justified belief is false, and since a justificatory process that produces a false belief cannot be a successful one (believing what is true is one of the main epistemic goals), if the override is right, something must have gone wrong with the justificatory process. Thus, we have to acknowledge that overrides do have some implications in the realm of higher-order epistemic reasoning too. However, such an acknowledgement does not question the point made in (*HOVUD*).

The point made in (*HOVUD*) is that while the acceptance of an underminer (and the belief revision that follows) compels the agent to engage in some higher-order epistemic reasoning, the acceptance of an override does not. And that is compatible with the acceptance of overrides having some implications in the higher-order sphere. Let us have another quick look at the examples. The agent can accept the override in cases (2) and (4), and accordingly update her set of beliefs by replacing p with not- p , without considering any thought about the relevant justificatory sources or processes. Of course, if the agent is a very responsible epistemic agent, she might ask herself some questions about the epistemic worthiness of the sources or the processes. But she does not have to. On the other hand, in cases (1) and (3), the agent gives up the belief that p precisely as a result of the emergence of doubts about the trustworthiness of the justificatory sources or processes that have been raised by the relevant underminer.¹¹

Summing up, (*HOVUD*) does two things: on one hand it provides an account of how undermining defeat works; on the other, it explains the difference between underminers and overrides in terms of the impact they have on the sphere of higher-order reasoning. Contrary to what underminers do, overrides do not need to appeal to higher-order commitments about how the belief was formed, and do not need to challenge the epistemic worthiness of those commitments. If so, underminers force the subject to reflect about the ways in which her beliefs were formed, and thus cut at a deeper epistemic level than overrides do: unreflective agents can suffer overriding defeat, but they cannot suffer undermining defeat.¹²

¹¹ In Melis (2014: § 3), I gave more emphasis to another difference between underminers and overrides at this junction: while the suggestion made by overrides with respect to the faultiness of the justificatory process is compatible with the correct execution of the process—it is in the nature of justificatory processes delivered by fallible sources of justification that sometimes they fail to lead to truth—the suggestion made by the underminers is not. In other words, while overrides merely suggest that (as sometimes happens) the process failed to lead to truth, underminers suggest that a specific disturbing event has caused the process to fail. This point, reflected in the last clause mentioned in (*HOVUD*), is not crucial for the aim of extending (*HOVUD*) to propositional justification, and limits of space advice against expanding further on it.

¹² One might think that there might be cases in which defeaters work as both underminers and overrides. I do consider the possibility of such cases in Melis (2014: § 5).

1.1 More on Higher-Order Commitments¹³

Before moving on, I wish to clarify when exactly a higher-order commitment is in place. To that end, let us imagine that the agent has some misconceptions about the way in which she formed her belief. Say that the agent justifiably believes a proposition p , and that she explicitly believes (again, justifiably so) that she based her belief on source X while in fact she based it on source Y . In such a case, we can think of two kinds of higher-order commitments: those concerning propositions about the actual source Y —such as “the belief that p was based on source Y ”, or “source Y is reliable”, etc.—and those concerning propositions about the putative source X —such as “the belief that p was based on source X ”, or “source X is reliable”, etc. Are both kinds of commitments in place? Yes. For every justification, actual or merely putative, there are some higher-order commitments in place, and (*HOVUD*) is a proposal which is meant to account for the defeat of whichever justification (either actual or putative) is under scrutiny. But let us reflect on the case just described with some care.

One might raise the following worry. Since the agent described is completely unaware that she formed her belief in p on the basis of source Y , she would not be able to take a higher-order proposition like “source Y is reliable” to be true (or warranted) on reflection. How can the commitments about the actual source Y be in place then? It is at this juncture that the emphasis on the *relevant agent* becomes important. While in many ordinary circumstances the agent that would take the commitments to be in place is the actual agent, in cases where the actual agent has some misconceptions on the way she formed her beliefs (or simply ignores it without being in the position to figure it out—see next paragraph and fn. 14), the agent that would take a given commitment to be in place varies depending on the commitments (and the related justification) in question. If we consider the commitments about the *putative way* in which the belief was justified, then the relevant agent is the actual agent. In the case mentioned, it is the actual agent that, on reflection, would take the higher-order proposition “the belief that p was based on source X ” or “source X is reliable” to be true or warranted. However, if we consider the commitments about the *actual way* in which the belief was justified, then the relevant agent is the idealized agent who is like the actual agent in all respects, except that she has no misconceptions about the way in which the belief was formed. Of course, in the case described the actual agent might not be able to appreciate the defeat of what I have referred to as ‘actual justification’, but that does not mean that that justification is not defeated—in an externalist sense. And (*HOVUD*) has the tools to account for that defeat.

The relation between (*HOVUD*) and the defeat of externalist justification can be seen clearly by considering the following similar worry. Suppose that the subject S is (non-culpably) unaware of the fact that she has formed her belief in p on the basis of source X , and that the defeater, although it does nothing to show that X is the source of S 's belief, challenges the epistemic worthiness of X . Does not (*HOVUD*) in this case predict defeat when intuitively none occurs? It is true that (*HOVUD*) predicts that, in the case described, there is undermining defeat. I stand by that. I also acknowledge that there is an intuitive sense in which there is no defeat, but the intuitions at play here involve an internalist conception of

¹³ Thanks to two anonymous referees for alerting me to the issues discussed in this section.

epistemic justification, according to which a subject can only be justified if she is aware of the relevant justifying factors, as it were. In the case at hand, on an internalist conception of justification, *S* is not justified in believing that *p* on the basis of source *X* because *S* is not aware that she based her belief on source *X*. And, the objection goes, if there is no justification, there is no defeat. I have no quibbles with that. However, what the objection fails to mention is that, in an externalist sense—which does not require the subject to be aware of the relevant justifying factors—*S* is justified in believing that *p* on the basis of source *X* (assuming the conditions on the source posed by the relevant externalist account of justification are met). And such externalist justification is precisely the one that (*HOVUD*) predicts gets defeated.¹⁴

The lesson that we should draw from the cases illustrated in this section is that (*HOVUD*) is neutral with respect to at least some of the dimensions along which epistemic justification can vary. (*HOVUD*) is a theory of undermining defeat that explains how undermining defeat works in either internalist or externalist cases, and regardless of whether the justification in question is actual or putative. Let us now turn to the extension of (*HOVUD*) to propositional justification.

2. The Limit to Doxastic Justification

Before we see why (*HOVUD*) is limited to doxastic justification, and why that should be a reason of concern, let us recall the difference between propositional and doxastic justification.

It is common to account for the distinction between propositional and doxastic justification along the following lines: when a subject has a reason or evidence to believe that *p*, she has propositional justification *to* believe that *p*; when a subject has a reason or evidence to believe that *p*, and bases her belief that *p* on those reasons or that evidence, she is doxastically justified *in* believing that *p*.¹⁵ If a subject is doxastically justified in believing a proposition, she also has propositional justification to believe that proposition, but not vice versa: in other words, doxastic justification entails propositional justification, but not vice versa.

Here is a quick example: an agent that justifiably believes that Cagliari is in Sardinia, and that Sardinia is in Italy, has propositional justification *to* believe that Cagliari is in Italy, but it is when the agent (competently) draws the relevant inference that she is (doxastically) justified *in* believing that Cagliari is in Sardinia.

To see why (*HOVUD*) is limited to doxastic justification, we need to understand that while doxastic justification requires that the proposition justified is believed by the agent, propositional justification does not. Recall that according to (*HOVUD*), underminers work by appealing to some higher-order commitments about the way in which the relevant belief was originally formed. But propositional justification does not require belief formation, and thus does not require

¹⁴ To put the point in terms of the relevant agent, we might say that the higher-order commitments involved concern propositions that an *idealized agent*—alike the actual agent in all respects, except that she does not ignore how the belief was formed—would take to be true or warranted.

¹⁵ See Feldman (2002: 46) and Pollock and Cruz (1999: 35-36) for just two examples of this way of drawing the distinction.

any commitments about belief-formation. So, (*HOVUD*) does not apply to propositional justification.

To see why the advocate of (*HOVUD*) should be concerned by this limitation, consider the following variation of example (1). Suppose that Adam tells me many things and that, before I have a chance to consider them and update my set of beliefs, Lauren tells me that Adam is a compulsive liar. In such a case, I obtain an underminer for the justification I get from Adam's testimony, before I have a chance to form the beliefs whose higher-order commitments (*HOVUD*) predicts that the underminer calls into question. Thus, (*HOVUD*) cannot account for such a case. And the reason is that the underminer in question affects the propositional justification to believe propositions before the agent actually forms any belief on the basis of that justification.

The lesson seems to be that the phenomenon of defeat arises with respect to propositional justification. If so, any account that, like (*HOVUD*), is tied to doxastic justification does not generalize far enough.¹⁶

3. Extending (*HOVUD*) to Propositional Justification

In this section I will argue that (*HOVUD*) can be extended to propositional justification in virtue of the close relationship that exists between propositional and doxastic justification. In a nutshell, propositional justification involves higher-order commitments analogous to those involved in doxastic justification, and those commitments are what underminers call into question in the defeat of propositional justification.

But let us begin with the relationship between propositional and doxastic justification. We said that a subject that bases her belief in p on some evidence e is doxastically justified in believing p . However, the justificatory process that brings the subject to believe p on the basis of e is grounded in a relation between e and p that is in place regardless of whether the subject uses it to her avail. That relation of support between e and p is propositional justification. In this sense, doxastic justification arises from propositional justification.

Let me briefly expand on this. Consider the set of the propositions that constitute the subject's evidence. Those propositions epistemically support a number of other propositions, to which they are related in propositional justification. The use of that epistemic link to form a justified belief (doxastic justification) is available to the subject since the moment in which she acquires the evidence, but it might take a while before she follows that link to form an actual belief, if ever. Justificatory processes proceed along that pre-existing epistemic link and lead the subject to doxastic justification. What matters for our purposes is that the relations of (propositional) justification between the propositions in the set of the evidence and those in the set of what is supported by the evidence involve higher-order commitments. Of course, they are not commitments about actual belief-formation, but rather commitments about the ways in which e supports p .¹⁷ Those are the commitments that are challenged by undermining defeaters.

Let us consider the problematic case for (*HOVUD*) mentioned in the previous section. The (propositional) justification to believe the various propositions

¹⁶ Thanks to Declan Smithies for bringing this issue to my attention in correspondence.

¹⁷ They concern the source of justification, the acquisition of the evidence e provided by the source, and the relation of support between e and p .

that Adam is testifying involves the commitment towards the proposition that the source of justification is Adam's testimony. Such commitment does not concern belief-formation, but only the support relation that exists between the propositions that work as justifiers (propositions of the form 'Adam says that p ') and the propositions justified (the various corresponding p 's). For the acquisition of the information that Adam is a compulsive liar to undermine the (propositional) justification to believe what Adam says, that commitment has got to be in place: should the relevant agent take the propositional justifications in question to be due to Peter's testimony rather than Adam's, the information that Adam is not to be trusted would not have any undermining effect. Moreover, it is the epistemic worthiness of that very commitment that is called into question by the undermining defeater, which suggests that Adam is not a reliable source. This is all very much in the spirit of (*HOVUD*).

Here is a formulation of a version of (*HOVUD*) that applies to propositional justification.

(*HOVUD-prop*) Underminers suggest that something was wrong with the source of justification or with the grounds that support p , and they operate their defeat by appealing to the higher-order commitment that p is supported by that source or those grounds. If the suggestion is that something was wrong with the grounds, rather than with the source, the defectiveness is to be understood as the occurrence of a mistake or some other disturbing event that spoiled the epistemic worthiness of the grounds.

Of course, the main difference between (*HOVUD*) and (*HOVUD-prop*) is that the former talks about commitments concerning the formation of a justified belief, and the latter talks about commitments concerning the formation of a justification which need not result in belief formation. Just like in the case of doxastic justification, the commitments involved in propositional justification concern propositions that the relevant agent would take to be true on reflection (as long as she sticks to the corresponding justification). However, the agent that would take the higher-order propositions to be true or warranted on reflection is, once again, an idealized one: it is the idealized agent that, on the basis of the same body of evidence available to the actual agent, goes on to form all the beliefs that are supported by that body of evidence (including those that, for whatever reason, the actual agent fails to form).

A second difference worth noticing is that the notion of justificatory process included in (*HOVUD*) has been substituted with that of grounds in (*HOVUD-prop*). The reason is that justificatory processes are largely mental and have to do with belief-formation, and thus cannot play a role in (*HOVUD-prop*). However, propositional justification offers something that underlies the justificatory process involved in doxastic justification, and I have called that 'the grounds for p '. I understand grounds as a non-mental analogue of justificatory processes: the pre-existing epistemic path that goes from the subject's acquisition of e to the conferral of the positive epistemic status to p . Such a notion of grounds can thus refer both to the acquisition of the evidence (from which the epistemic path begins) and to the relation between the evidence and the proposition supported p (which constitute the remaining of the epistemic path). Just like justificatory processes, justificatory grounds have their origin in a source of justification.

Let us now briefly consider the previous examples of undermining defeat and suppose that the justification in question is propositional rather than doxastic. (*HOVUD-prop*) works in a way just parallel to (*HOVUD*). In (1) the underminer suggests that the source was not reliable with respect to the subject matter, and the higher-order commitment under attack is that the (propositional) justification to believe that Paul McCartney played in the Beatles has its source in Adam's testimony. In example (3) the underminer suggests that something was wrong with the grounds, and the commitment under attack is that the (propositional) justification to believe *p* is given by that alleged proof (as opposed to *the agent's execution* of that alleged proof, which constitutes the justificatory process).

4. Conclusion

To conclude, I acknowledge that, in an important sense, the phenomenon of defeat arises at the propositional level. Since the phenomenon of defeat concerns justification, and, at least in the way explained above, propositional justification comes before doxastic justification, it is no surprise that the phenomenon of defeat concerns primarily propositional justification. Yet, as I hope to have shown, (*HOVUD*) can easily be extended to account for the defeat of propositional justification. More generally, (*HOVUD*) promises to have the tools to account for undermining defeat regardless of several of the dimensions along which justification can vary: propositional vs. doxastic, actual vs. putative, externalist vs. internalist.¹⁸

References

- Casullo, A. 2003, *A Priori Justification*, Oxford: Oxford University Press.
 Feldman, R. 2002, *Epistemology*, Upper Saddle River, NJ: Prentice Hall.
 Melis, G. 2014, "Understanding Undermining Defeat", *Philosophical Studies*, 170, 433-42.
 Pollock, J. 1974, *Knowledge and Justification*, Princeton: Princeton University Press.
 Pollock, J. and Cruz, J. 1999, *Contemporary Theories of Knowledge*, 2nd ed., New York: Rowan & Littlefield.
 Sturgeon, S. 2014, "Pollock on Defeasible Reasons", *Philosophical Studies*, 169, 105-18.

¹⁸ I would like to thank Declan Smithies and Carrie Jenkins for very helpful discussions, and everyone who attended the Philosophy Work In Progress seminar at the University of Aberdeen in February 2014, where an early version of this paper was presented. I am especially grateful to two anonymous referees for raising several precious questions, which led to a substantial improvement of the article.