Beliefs and Blameworthiness

Liz Jackson, Kansas State University

Abstract: In this paper, I will present an analysis of epistemic blameworthiness. First, I will give Michael Bergmann's definition of epistemic blameworthiness, and propose a modification to one of his conditions. Then, I will argue that his definition is missing a further condition, specifically a 'control' condition, so called because it is about whether one has control over one's doxastic attitudes. I will propose a specific formulation of the control condition, and evaluate the revised definition with respect to potential counterexamples.

I. Introduction

It is not uncommon to say things like, "she should have believed that" or "he ought to have known that." Behind these common phrases lies the assumption that we are responsible for at least some of our beliefs. Given this assumption, one might wonder what it means to fail to live up to one's epistemic responsibilities. Roughly, I take *epistemic blameworthiness* to be a failure to fulfill some of our duties related to our beliefs.

This raises further questions. Specifically, what does it take for one to be blameworthy for a particular doxastic attitude?[1] In this paper, I will present an analysis of epistemic blameworthiness. First, I will give Michael Bergmann's definition of epistemic blameworthiness, and propose a slight modification to one of his conditions. Then, I will argue that his definition is missing a further condition, specifically a 'control' condition, so called because it is about whether one has control over one's doxastic attitudes. I will propose a specific formulation of the control condition, and evaluate the revised definition with respect to potential counterexamples.

II. Epistemic blameworthiness

Bergmann's Account

What is the proper understanding of epistemic blameworthiness? Bergmann defines epistemic blameworthiness as follows:

> **"EBW:** S's believing p at t is something for which she is epistemically blameworthy iff either (i) S believes at t that she ought not to believe p or (ii) S's failure to believe at t that she ought not to believe p is relevantly due to some other doing or failure of hers for which she is epistemically blameworthy."[2]

First, I will motivate conditions (i) and (ii) with an example. Tom has good evidence that Company X is exploiting innocent people overseas in order to produce a cheaper product. *Ceteris paribus*, Tom would believe that Company X committed human rights abuses, but

---

[1] S has a doxastic attitude toward proposition p iff S believes that p, withholds belief that p, or disbelieves that p.

[2] Michael Bergmann. *Justification without Awareness.* (Oxford, UK: Oxford University Press, 2006) 92.

Tom loves the cheap products produced by Company X. Tom convinces himself that *Company X is not exploiting its workers*. At the same time, he remains committed to an inferential standard such that he also believes that, given the evidence, he ought not to believe that X is not exploiting its workers. I take it that we would want to say of Tom that he is blameworthy for his failure to believe that X is exploiting its workers. This fits Bergmann's analysis, because Tom fulfills condition (i), so he is blameworthy.

Alternatively, assume again that Tom has substantial evidence that Company X is exploiting people. However, in this case, let's also assume Tom believes A, *I should believe everything my mother tells me*, and she assures him that Company X is not exploiting people. Because of this, Tom believes B, *Company X is not exploiting its workers*, although the only reason he believes this is the testimony of his mother. If we stipulate that Tom is epistemically blameworthy for his belief that A, and A is the basis for his belief B, then it appears that Tom is also epistemically blameworthy for his belief B. Tom fulfills condition (ii) of Bergmann's definition, so Tom's intuitive blameworthiness for B in this case also fits Bergmann's analysis.

Still, there are two initial difficulties with Bergmann's account, as is. First, I want to note that Bergmann's analysis only applies to the attitude of *believing p*, when it is intuitive that one can also be blameworthy for withholding belief in p[3] or disbelieving p.[4] Bergmann may have only been interested in questions about when someone is blameworthy for holding a belief, but it would be helpful to have an analysis that applies to all doxastic attitudes.[5]

Second, consider the following scenario: Sarah disrespects her teacher and so she fails to listen in class or read the syllabus. Sarah's teacher is a jerk, so the disrespect (which one could characterize as a complex belief) is not something for which she is blameworthy. However, Sarah is blameworthy for letting the disrespect to cause a failure of attention. Because of her lack of attention, she does not know that the exam is today so she fails the exam. Hence she is blameworthy for not having the belief E, *the exam is today*. She does not believe that she should believe E, so she does not fulfill condition (i); at the same time, she does not fulfill condition (ii) because she is not *epistemically* blameworthy for her failure to prevent her disrespect from causing her lack of attention. Therefore, Sarah's case is a counterexample to Bergmann's analysis; (i) and (ii) are not necessary for blameworthiness.

Intuitively, Sarah is *epistemically* blameworthy for failing to believe E because she is *morally* blameworthy for her failure to prevent a particular causal relation – she is responsible for letting her disrespect cause her lack of attention.[6] More generally, it is possible for a

---

[3] S withholds belief that p iff S has considered p and neither believes p nor ~p. For example, it would be rational to withhold belief that *there is an even number of stars*.

[4] S disbelieves that p if S believes ~p.

[5] This is not necessarily a defect or problem with Bergmann's definition; I am just interested in a different kind of analysis.

[6] Another case where one might be morally blameworthy for preventing causal relations might be one where we could intervene to prevent a death, i.e. if a child was drowning in a pond right in front of us. For those who are dubious about the possibility of moral responsibility for singular or actual causal relations, note that the do/allow distinction as traditionally rendered holds that we are more morally responsible for one kind of causal connection—the doings—than for others, the allowings.

person to be morally or pragmatically[7] blameworthy for an action (not a belief), which, in turn, causes a belief for which she is then blameworthy. Therefore, I want to suggest that condition (ii) should read "some other doing or failure of hers for which she is *blameworthy*," (rather than *epistemically blameworthy*). When we make this change, Sarah will fulfill our new condition (ii), and Sarah will be blameworthy.

This change also has a second advantage. One may be worried that Bergmann's definition suffers from circularity.[8] Bergmann claims that clause (i) of his definition functions as a base clause and clause (ii) functions as a recursive clause, so that all blameworthiness that results from clause (ii) being satisfied must ultimately trace back to clause (i) being satisfied.[9] In other words, all beliefs you are blameworthy for are beliefs that can be traced back to the belief that you should not believe something. But consider the following case: S believes that P. S would believe *I should not believe P* if it were not for another belief S has, Q. S would also believe *I should not believe Q* if it were not for S's belief that P. In this case, because S neither believes *I should not believe P* nor *I should not believe Q*, S's failure with respect to P and Q cannot be traced back to S's belief that she should not believe something. For these reasons, there are worries that Bergmann's definition may fail to be a successful inductive definition, but changing *epistemically blameworthy* to *blameworthy* in clause (ii) will eliminate these worries.

In summary, we want to modify our definition such that it (1) applies to all doxastic attitudes, not just *believing that p* and (2) avoids our counterexample and circularity problems by changing *epistemically blameworthy* to *blameworthy* in clause (ii). We can edit his definition to include these two changes:

> **EBW₁:** S's doxastic attitude toward p at t is something for which she is epistemically blameworthy iff either (i) S believes at t that she ought not to hold that attitude toward p or (ii) S's failure to believe at t that she ought not to hold that attitude toward p is relevantly due to some other doing or failure of hers for which she is blameworthy.[10]

There are various other grounds on which one might challenge Bergman's account, or developments of it such as the modified analysis immediately above. For example, one might appeal to cases in which one's beliefs about what one should believe are mistaken. If one grants that there are objective facts about doxastic duties, then it will be possible to violate those duties even if one does not believe that one has violated them. But there is no such possibility on Bergmann's account. However, I will lay these concerns aside for now, as I am concerned to explore one particular source of unease: Bergmann's analysis lacks a control condition.

---

[7] I take pragmatic blame to be, roughly, a failure to be prudent or efficient - a failure to meet or work toward one's goals.

[8] Bergmann claims that EBW is meant to be a definition of epistemic blameworthiness, not just necessary and sufficient conditions.

[9] Michael Bergmann. *Justification without Awareness.* (Oxford, UK: Oxford University Press, 2006) 90.

[10] Michael Bergmann. *Justification without Awareness.* (Oxford, UK: Oxford University Press, 2006) 92.

A Control Condition

Consider the following case: An undergrad, Fred, is taught Hume's problem of induction, and so forms the belief that he ought not to believe p, *the future will imitate the past* (fulfilling condition (i)).[11] However, as Fred leaves the classroom, he finds himself thirsty, and so goes to get a drink of water. He proceeds to take many similar actions that show he still believes that p. We do not hold Fred accountable for believing p because we do not think it is psychologically possible for him to not believe p. Because Fred does not have a choice about whether or not to believe that p, he is not epistemically blameworthy for this belief, even though condition (i) is fulfilled; EBW_1's analysans is not sufficient for epistemic blameworthiness.

This counterexample shows that what Bergmann's analysis is lacking is *a control condition*. If we grant the commonly held meta-ethical principle that 'ought implies can,' it would seem that one who lacks control over a doxastic attitude cannot be blamed for that attitude. Hence it is reasonable to assume that some kind of control over a doxastic state is necessary for epistemic blameworthiness with regard to that state. Given this, we can edit Bergmann's definition again:

> **EBW_2:** S's doxastic attitude toward p at t is something for which she is epistemically blameworthy iff (1) either (i) S believes at t that she ought not to hold that attitude toward p or (ii) S's failure to believe at t that she ought not to hold that attitude toward p is relevantly due to some other doing or failure of hers for which she is blameworthy <u>and</u> (2) S had control over her doxastic attitude toward p at t.

This definition looks better, and it appears to deal with our counterexample. Fred does not have control over his belief that *the future will mimic the past*, and for this reason, he does not fulfill condition (2). Fred's case is not a counterexample to EBW_2.

What Kind of Control?

To ensure EBW_2 is sufficiently informative, it may be helping to make clause (2) more specific. We can borrow from Alston's analysis of control in his essay "The Deontological Conception of Doxastic Justification."[12] Alston distinguishes three main kinds of control over our beliefs we might have that could satisfy this condition: direct control, long-range control, and indirect influence. In this section, I intend to do two things: one, describe each of the three kinds of control, and two, argue that each kind is, at very least, possible.

Direct Control

What distinguishes direct control over the other kinds of control is that it is immediate and short-term, rather than a long-term control of a belief over time. It is the

---

[11] Here, I am assuming Fred's reasoning is something like this: "I am not justified in believing that p; therefore, I ought not to believe that p." I am neither claiming that Hume advocates this inference, nor that this inference is sound. I merely meaning to stipulate a case.

[12] William Alston, *Epistemic Justification: Essays in the Theory of Knowledge* (New York: Cornell University Press, 1989) 119-142.

ability to bring about a doxastic attitude "right away, in one uninterrupted intentional act."[13] But can we ever have direct control over our beliefs? Alston thinks we rarely do, if ever.[14]

However, I disagree with Alston on this point. For example, consider a story told by Clifford in "The Ethics of Belief." A particular ship owner was about to set sail in his ship, Providence. However, Providence was old, not built well, and had been repaired many times; for these reasons, he had doubts that she was seaworthy, and thought she might need to be totally overhauled before she was safe to sail. But the ship owner managed to overcome these doubts before the voyage, reminding himself she had safely completed many other voyages, including ones with serious storms. He chose to trust Providence, and made the decision to believe she would protect her passengers and take them safely to their destination. He chose to dismiss all of his doubts about Providence, and in doing so, he "acquired a sincere and comfortable conviction that his vessel was thoroughly safe and seaworthy; he watched her departure with a light heart."[15]

In this example, the ship owner has inclinations both to trust Providence and to distrust her. He clearly wavers between believing *Providence is reliable* and *Providence is not reliable*, but he chooses to believe the former in a short, uninterrupted act, so this is apparently a reasonable instance of direct control. Wolterstorff notes that this story does not seem bizarre; in fact, it seems like we could come up with a host of similar examples. Direct control is more plausible and more common than Alston thinks.[16]

Long-Range Control

A second kind of control we might have is long-range control. This is "the capacity to bring about a state of affairs, C, by doing something (usually a number of different things) repeatedly over a considerable period of time, interrupted by activity directed to other goals."[17] We have long-range control over things like our weight and our blood pressure; we may also have a similar type of power to influence our beliefs. For example, it seems plausible that I can set out on a project to get myself to believe God exists. I can study arguments for God's existence, spend time with theists, find smart theists to address my questions and objections, etc.

Alston is suspicious that we can reliably influence our beliefs via long-range control. While he acknowledges that this sometimes works, he nevertheless doubts the success rate for this is substantial.[18] Because of the low success rate, Alston does not want to count this as legitimate control. However, Wolterstoff points out that we often use long-range control not to acquire or get rid of beliefs, but rather maintain or strengthen one we already have: "to maintain our belief in Marxism, to maintain our atheism, to hold fast to our Presbyterianism."[19] This is common and seems to be much more successful. Additionally, it

---

[13] Ibid, 129.

[14] Ibid, 125.

[15] Clifford, William. "The Ethics of Belief." *Contemporary Review,* 29, (1877): 289.

[16] Nicolas Wolterstorff, *Practices of Belief.* (Cambridge, UK: Cambridge University Press, 2010) 76.

[17] William Alston, *Epistemic Justification: Essays in the Theory of Knowledge* (New York: Cornell University Press, 1989) 134.

[18] Ibid, 135.

[19] Nicolas Wolterstorff, *Practices of Belief.* (Cambridge, UK: Cambridge University Press, 2010) 76.

does seem like sometimes we can use long-range control to acquire or get rid of beliefs, such as the example of believing in God. While this may not always be successful, it is more common and more fruitful than Alston acknowledges.

Indirect Influence

A final category of control that we may have is what Alston calls "indirect influence." This type of control is different than the others in that it does not involve altering a doxastic attitude *toward a specific proposition*. Rather, it refers to the control we have over the things that influence our beliefs and belief-forming habits. This includes, "how long I look for relevant evidence or reasons, reflect on a particular argument, seek input from other people, search my memory…training myself to be more critical of gossip, instilling in myself a stronger disposition to reflect carefully before making a judgment on highly controversial matters…"[20] Supposing there are intellectual obligations, one would presume that they include obligations to do more rather than less of each of the things in the quoted list; doing these things seems to help us seek truth and avoid falsehoods. Of all three types of control, this one seems the most clearly psychologically possible.

Given these three categories of control, we can edit our definition of blameworthiness again, making it even more specific:

> **EBW₃:** S's doxastic attitude toward p at t is something for which she is epistemically blameworthy iff
> (1) either (i) S believes at t that she ought not to hold that attitude toward p or
>   (ii) S's failure to believe at t that she ought not to hold that
>   attitude toward p is relevantly due to some other doing or failure
>   of hers for which she is blameworthy <u>and</u>
> (2) S had (i) direct control or
>   (ii) long-range control or
>   (iii) indirect influence over her doxastic attitude toward p at t.

III. Proposed counterexamples to EBW₃

We have significantly edited Bergmann's definition, but even with these modifications, does our definition suffer from counterexamples? We will consider several potential counterexamples to the EBW₃.

(1) One species of counterexamples to EBW₃ appeals to pragmatic oughts. For example, let's say there is a belief-removal machine that you can hook up to your brain and can remove any belief you desire. However, this machine is extremely expensive. Suppose Bob believes proposition r and also has the deep conviction he should not believe r. Because of this machine, Bob has long-range control over r – let's stipulate this is the only way Bob could make himself not believe r. But the belief-removal machine is so expensive that Bob would have to work for 25 years in order to save enough money to buy the machine. Since the machine is so expensive, it is supposed to be intuitive that Bob is not blameworthy for

---

[20] William Alston, *Epistemic Justification: Essays in the Theory of Knowledge* (New York: Cornell University Press, 1989) 138.

his belief that r. Buying the machine is pragmatically impractical and inefficient, and some say this cleanses Bob of his epistemic blame.[21]

However, I think the proper understanding of the case is to say that if Bob does not buy the machine, Bob is still epistemically blameworthy but he is pragmatically blameless. It is unhelpful to talk about what Bob 'ought' to do in a situation like the above without qualifying which type of 'ought' we are talking about. As Richard Feldman says, "[It is very unclear] that there is such a thing as *just plain ought*, as opposed to the various kinds of oughts philosophers have succeeded in distinguishing."[22] Given this, is the reasonable to say that two different species of oughts are making demands of Bob; relative to his epistemic duties he should take one course of action, and relative to his pragmatic duties he should take another course of action.

(2) A second group of counterexamples for EBW$_3$ involves moral oughts. Borrowing from counterexample (1), let's again say Bob believes proposition r and also has the deep conviction he should not believe r. There is an evil demon who has the power to control Bob's beliefs, and he makes Bob an offer: if Bob will brutally murder 10,000 children, then the demon will cause Bob to no longer believe r. Intuitively, Bob is not blameworthy for continuing to believe r, because the only alternative involves doing something that is seriously morally wrong.

We can respond to this counterexample similarly to the way we responded to the first counterexample: If Bob does not murder the children, he is epistemically blameworthy but morally blameless. This raises some other interesting questions about the correct course of action when two kinds of 'ought's conflict each other. In cases (1) and (2), one ought is much more demanding than the epistemic ought so, intuitively, that ought 'overrides' the epistemic ought. However, there may be cases where one ought does not as clearly override another, and the correct course of action is unclear.[23] This is an interesting area that could be researched further, and has already been addressed by a few philosophers.[24]

(3) Let's take a case similar to Tom's second scenario on p. 3 in which we stipulated that Tom was blameworthy for his belief *I should believe everything my mother tells me.* However, in this case, let's suppose he is not blameworthy for this belief; it results from an honorable, deep respect for his elders. Let's also suppose his mother firmly tells him to not believe p, so Tom forms the belief q, *I should not believe p.* At the same time, suppose Tom has mounds of overwhelming evidence for p, so he is not blameworthy for believing p. Apparently, Tom is not blameworthy for his belief that p, but at the same time, believes he should not believe

---

[21] Thanks to Amelia Hicks and Katrina Prichard for this case.

[22] Richard Feldman, "The Ethics of Belief." Philosophy and Phenomenological Research, 60, no. 3. (2000): 692.

[23] Some philosophers argue that moral oughts always override other oughts. See Feldman on Hall and Johnson in Feldman's "The Ethics of Belief" (2000), p. 692. Others, such as Dogherty (2012) argue that all oughts are ultimately moral. Finally, some, such as Feldman himself, argue that "there is no meaningful question about whether epistemic oughts 'trump' or are trumped by other oughts" (2000, 694).

[24] See Eugene Mills, "The Unity of Justification," *Philosophy and Phenomenological Research* 58 (1998): 27-50*,* and Richard Feldman, "The Ethics of Belief," *Philosophy and Phenomenological Research* 60, No. 3. (May, 2000): 667-695.

that p. If we suppose Tom has control over his belief that p, then this is a potential counterexample to EBW_3.[25]

The defender of EBW_3 might respond in several ways. First, he might claim that, although Tom has overwhelming evidence for p, if Tom has the deep conviction that q (*I should not believe p*), then Tom is blameworthy for believing p; Tom should follow his convictions. As long as Tom chooses to maintain his belief that q, EBW_3's defender can simply maintain that Tom is blameworthy for not following his convictions.

A second factor we must consider when thinking about this case is that ordinary folk, when presented with overwhelming evidence for p, would suspend their belief that q (*I should not believe p*). When we are thinking about this case, we are thinking about ordinary folk, so our intuitions about the case are not clear. Therefore, while this third counterexample might put some pressure on EBW_3, it is not a devastating counterexample; there are several potential responses.

IV. Conclusion

In conclusion, I propose EBW_3 as a potential definition for epistemic blameworthiness. I am unaware of any serious counterexamples to EBW_3, and I think this definition furthers our understanding of what it means for one to be epistemically blameworthy. Additionally, exploring this definition has brought up some interesting areas for further research, such as how epistemic, moral, and pragmatic oughts fit together, or if there can ever be a thing as "just plain ought."[26]

---

[25] Thanks to Salvatore Florio for this case.

Works Cited

Alston, William. *Epistemic Justification: Essays in the Theory of Knowledge*. New York: Cornell University Press, 1989.

Bergmann, Michael. *Justification without Awareness*. Oxford, UK: Oxford University Press, 2006.

Clifford, William. "The Ethics of Belief." *Contemporary Review,* 29, (1877): 289-309.

Dougherty, Trent. "Reducing Responsibility: An Evidentialist Account of Epistemic Blame" *European Journal of Philosophy,* 20, no. 4 (2012): 534-547.

Feldman, Richard. "The Ethics of Belief." *Philosophy and Phenomenological Research,* 60, no. 3. (2000): 667-695.

Mills, Eugene. "The Unity of Justification." *Philosophy and Phenomenological Research,* 58, (1998): 27-50.

Wolterstorff, Nicolas. *Practices of Belief*. Cambridge, UK: Cambridge University Press, 2010.