A responsibilist account of knowledge

Xingming Hu

Nanjing University, China

Forthcoming in *Philosophical Quarterly*

Correspondence to: xingminghu99@gmail.com

This paper argues for a responsibilist account of knowledge: S knows that p iff S believes the truth that p (rather than one of the alternatives to p) because S forms/retains the belief in a way that is ultima facie epistemically responsible. This account implies that knowing that p requires neither having evidence that favors p over ~p, nor possessing reliabilist virtues, nor exhibiting responsibilist virtues or motives.

Keywords: knowledge; evidentialism; reliabilism; responsibilism; virtue epistemology; epistemic blame

Many philosophers (e.g. Chisholm 1977; Bonjour 1985; Zagzebski 1996; Conee & Feldman 2004; Corlett 2008; Axtell 2008; Williams 2008; Greco 2010; Peels 2017; Goldberg 2018; Brogaard 2023) hold that S knows that p only if S forms the belief that p in an epistemically responsible way. This paper draws on this insight and proposes a new analysis of knowledge: S knows that p iff S believes the truth that p (rather than one of the alternatives to p) because S forms/retains the belief in a way that is *ultima facie* epistemically responsible.[1] Call it *procedural responsibilism*. I argue that procedural responsibilism is a viable competitor to some influential forms of evidentialism and virtue epistemology.

[1] Williamson (2000: 185) argues that knowledge is a primitive, unanalyzable concept. However, an analysis of knowledge does not have to be reductive. It can also be connective (cf. Strawson 1992).

I will proceed as follows. Section 1 will clarify what it means to be ultima facie epistemically responsible and show how procedural responsibilism would address the Gettier problem. Section 2 will address the worry that procedural responsibilism collapses into an influential form of evidentialism. I will argue that procedural responsibilism denies the evidentialist principle that S knows that p at t only if S believes that p on the basis of evidence that favors p over not-p at t. Accordingly, it can better handle some problems (e.g. the skeptical problem) than the form of evidentialism defended by Conee and Feldman (2004). Section 3 will tackle the worry that procedural responsibilism reduces to a form of virtue epistemology incorporating both responsibilist and reliabilist elements. I will argue that procedural responsibilism implies that knowledge requires neither responsibilist virtues (or motives) nor reliabilist virtues. Accordingly, it offers a way to resolve problems facing various forms of virtue epistemology. Finally, I will respond to the concern that there might be counterexamples to procedural responsibilism.

1. Clarifying procedural responsibilism

Procedural responsibilism states that S knows that p iff S believes the truth that p (rather than one of the alternatives to p) *because* the way S forms/retains the belief is *ultima facie* epistemically responsible. In this section, I will first explain the notion of epistemic responsibility. Then I will clarify, via a Gettier case, what it takes to be *ultima facie* epistemically responsible.

S forms/retains a belief in an epistemically responsible way iff S is doxastically responsible yet epistemically blameless for forming/retaining the belief. Whether S is doxastically responsible for forming/retaining the belief that p is a matter of whether S can be appropriately appraised (e.g. blamed, praised, or neutrally appraised) for forming/retaining the belief that p. Consider Keith Lehrer's (1990) case of Mr. Truetemp, a person who has a device implanted in their brain that accurately detects the ambient temperature and produces a true belief about it. However, Truetemp is unaware of this device and forms beliefs about the temperature seemingly out of the blue. Intuitively,

Truetemp's true belief does not amount to knowledge. According to procedural responsibilism, this is because Truetemp's belief-forming process lacks doxastic responsibility, thereby failing to be epistemically responsible. Note that doxastic responsibility for holding a belief should be distinguished from having direct, voluntary control over its formation. For example, someone who sees a tree automatically believes that a tree is present. Although this belief formation is involuntary—seeing the tree essentially forces the belief—the person is still considered doxastically responsible. This responsibility does not stem from an ability to choose the belief directly in that moment. Instead, as some philosophers (e.g. Peels 2017: 141-142) suggest, it relates to the agent's broader capacity to influence their beliefs through voluntary actions. For instance, the person could have chosen to look elsewhere, thereby avoiding the perceptual experience that leads to the belief.

One might be doxastically responsible for a belief yet form/retain it in a blameworthy way.[2] Imagine a medical scientist who, after years of research on a new drug for a rare disease, believes that it is highly effective and safe based on preliminary data and animal trials. However, when human trials begin, several patients suffer severe adverse reactions, some even fatal. Despite this alarming new evidence, the scientist dismisses it and clings to her original belief, relying solely on the earlier results. Intuitively, she is not epistemically blameless in holding onto this belief in the face of such strong counter-evidence.

Strictly speaking, being epistemically blameless is not sufficient for being epistemically responsible. For someone might hold a belief without being doxastically responsible for it, and in such cases, they could be considered blameless for holding that belief. However, for simplicity, I will use 'epistemically blameless' to refer exclusively to situations where the individual is doxastically responsible but not blameworthy. In this sense, epistemically responsible belief = epistemically blameless belief.

[2] Here I use 'blame' in a broad sense. Being blameworthy is equivalent to being criticizable or deserving criticism. Cf. Boult (2024).

Sometimes, a person forms a belief in an epistemically blameless way, based on the information available to them. However, if they possessed all relevant information, the original way they form the belief might be epistemically blameworthy. For example, consider Alice, a high school student who enjoys reading her generally reliable local newspaper during breakfast. One morning, she reads an article about a well-known charity raising funds for a new community centre, emphasizing its positive reputation and past achievements. Trusting the report, she forms the belief that donating to the charity is a good idea. Unbeknownst to Alice, the charity is corrupt, and the article contains misinformation. Although her belief is false, the process by which she formed it is epistemically blameless given her limited information. However, if Alice had been aware of the charity's corruption, relying solely on that misleading article to form her belief would be epistemically blameworthy, as it would involve disregarding clearly contradictory evidence.

The way S forms the belief that p is *ultima facie* epistemically responsible iff it is epistemically blameless, and there are no (undefeated) defeaters: had S known all the relevant facts, the original way S forms the belief would still be epistemically blameless.[3] According to procedural responsibilism, if the way S forms the belief that p is epistemically blameless but not *ultima facie* epistemically responsible, then S does not know that p. Thus, we have a ready solution to the Gettier problem. Consider a canonical Gettier scenario: The city tower clock has always been reliable, and locals depend on it for the time. When Russell, a local philosophy professor, passes by the city tower this morning, he glances up and sees the clock reads 8:00. Unhurried and with no cause for suspicion, he does not bother to double-check. He naturally concludes that it is 8:00 AM. As it happens, the time is 8:00 AM, so his belief is true. However, unbeknownst to Russell, the clock actually stopped exactly 12 hours ago. Given its consistent reliability and his lack of any reason to suspect a malfunction, the way Russell forms his belief is epistemically blameless. Nevertheless, his true belief does not qualify as knowledge.

[3] Paxson and Lehrer's (1969) library book theft case is about defeated or misleading defeaters. But sometimes one is not blameless if one fails to consider a misleading defeater, as I will discuss later.

According to procedural responsibilism, this is because the way Russell forms the belief is not *ultima facie* epistemically responsible: Intuitively, had Russell known that the clock is not working, he should not form the belief in the original way, that is, he should not believe that it is 8:00 AM simply by looking at the clock's display and implicitly assuming that it is running accurately this morning. Otherwise, his beliefs would be baldly incoherent.[4]

So far, I have drawn upon intuitions regarding whether the agent in a particular case is epistemically blameworthy/blameless. One might wonder what makes the agent blameworthy/blameless. Here is a plausible hypothesis: S is epistemically blameworthy for forming/retaining a belief in a certain way iff doing so violates an epistemic duty S has, and this violation is not excusable. If S has no such epistemic duty, or S is excusable for failing to fulfill it, then S is epistemically blameless. We may distinguish two types of epistemic duty: (1) basic duty that everyone has merely in virtue of their status as epistemic subjects (i.e. people who can be held responsible for their beliefs) and (2) role-based duties that certain people have in virtue of the roles they play in their epistemic communities. As Goldberg (2018) suggests, both types of epistemic duty seem to arise from what others are entitled to expect. E.g. all normal people have the basic epistemic duty to avoid bald incoherence because we are entitled to expect each other to do so. People like scientists and doctors have a role-based epistemic duty to (say) study the latest significant developments within their discipline, because we are entitled to expect them to be up-to-date. (Such entitlements seem to be grounded in our epistemic practice. I will elaborate on this in Section 2.)

Thus, the medical scientist is epistemically blameworthy because she has an inexcusable failure to fulfil the role-based epistemic duty to carefully study extensive clinical trial data and adjust her beliefs accordingly. She has such a duty because patients are entitled to expect her to do so. By contrast, the high school student is epistemically

[4] This is similar to the no-defeaters approach to the Gettier problem. For responses to the criticisms of this approach, see de Almeida & Fett (2016).

blameless for believing that donating to the charity is a good idea, because she has no epistemic duty to investigate the newspaper article's accuracy: No one could reasonably expect her to do so. For similar reasons, the philosophy professor does not have any epistemic duty to check whether the city tower clock is working this morning. Yet both the high school student and the philosophy professor have the basic epistemic duty not to form beliefs in an obviously incoherent way, because we are entitled to expect them to avoid bald incoherence. It also seems that people are not excusable for failing to avoid bald incoherence as long as they can be held doxastically responsible. Thus, if they had acquired the relevant new information (e.g. the corruption of the charity and the malfunction of the clock) but continued to form their beliefs in the original ways, they would be epistemically blameworthy.

Finally, I'd like to note that my argument for procedural responsibilism does not presuppose the above account of epistemic blame (or duty), which must be tested by our intuitions about particular cases of epistemic blamelessness/blameworthiness. Rather, my argument directly rests on such intuitions.

2. Against evidentialism

One might worry that procedural responsibilism collapses into a form of evidentialism. Specifically, the way S forms/retains a belief in p is epistemically responsible iff (and because) S believes that p on the basis of sufficient evidence: S's total evidence sufficiently favors p over ~p.[5] Thus, the responsibilist account of knowledge amounts to the claim that S knows that p iff S believes the truth that p on the basis of sufficient evidence, and there are no (undefeated) defeaters. This view has been influential (Conee & Feldman 2004; cf. Lehrer & Paxson 1969). One might think procedural responsibilism is nothing more than a restatement of this view.

[5] Some might think that the way S forms/retains a belief in p is epistemically responsible iff S is doxastically justified in believing that p. Cf. Peels (2017). My account does not appeal to the concept of justification.

In what follows, I will first explain why procedural responsibilism does not collapse into evidentialism. Then I will argue that procedural responsibilism can avoid some problems that evidentialism faces.

Procedural responsibilism would deny the evidentialist principle that the way S forms/retains a belief in p is epistemically responsible only if S's total evidence favors p over ~p. To see that this principle is false, consider the following case:

> **Twins:** You are browsing through the aisles of the college bookstore. Suddenly, you spot a familiar face near the fiction section. You believe that the person you see is Judy, a classmate you met yesterday at the orientation mixer. You are right. However, you are unaware of the fact that Judy has an identical twin sister, Trudy, who is also in the bookstore. Due to relying solely on observation, you cannot rule out the possibility that the person you see is Trudy: If the person you see were Trudy, you would still believe that she is Judy.

Intuitively, the way you form the belief that the person you see is Judy is epistemically blameless (and thus responsible) even though your total evidence does not favor <the person you see is Judy> over <the person you see is Trudy>. It is unreasonable to blame or criticize you for failing to consider and rule out the possibility that the person you see is Trudy. After all, you only met Judy yesterday at the orientation mixer. You have no epistemic duty to consider and rule out this possibility.

One might object that you should not form the belief that the person you see is Judy merely based on observation. A more reliable approach involves incorporating her testimony. It would be easy for you to go to the person and ask her whether she is Judy. Your belief that the person is Judy should be based on both what you see and what she tells you.

I agree that you may easily adopt this more reliable belief-forming method, but relying solely on your observation is not epistemically blameworthy. No one is entitled

to expect you to do more in such a context. More specifically, the epistemic expectations others are entitled to have of you seem to be grounded in our epistemic practices. A legitimate epistemic practice is a social practice with internal rules that can adequately promote truth. To adequately promote truth, the internal rules of any epistemic practice would require us to avoid bald incoherence and other simple fallacies. But as Fleisher (2025) notes, different epistemic practices might have different internal rules. This is because practices are adapted to their specific circumstances. They often aim to answer different kinds of questions and operate with different resources (like conceptual schemes, background assumptions, methods). Factors like the specific topic, the social setting, and practical constraints like time pressure influence which rules are most suitable or adequate for promoting truth in that context. In addition, there are often multiple different, potentially inconsistent sets of rules that could be adopted to adequately promote truth. To function effectively and cooperatively, members of a practice need to settle on one specific set of rules so they can support each other's efforts rather than interfere. Thus, selecting and adhering to a specific rule set from the various permissible options is partly conventional: It involves partially arbitrary choices made previously by individuals to solve coordination problems. The internal rules are enacted socially through the acceptance of practitioners, which can be either implicit or explicit. Practitioners exhibit acceptance by expecting each other to follow the rules and holding each other responsible. The *Twins* is a case of everyday epistemic practice in casual settings. In this practice, the rule 'trust observation' suits the casual, quick nature of a bookstore encounter. It is a convention—your community implicitly accepts it for everyday IDs. Adding testimony might be more reliable, but it's not required—practical constraints (time, effort) and the low stakes (it's just a classmate) make observation adequate. The practice does not impose a duty to ask unless there's a cue (e.g. doubt or stakes rise). Without Trudy on your radar, the expectation to 'do more' is not part of the rule set your peers enforce.

So far, I have shown that the evidentialist principle, which states that the way S forms/retains a belief in p is epistemically responsible only if S's total evidence favors p

over ~p, is false. Next, I will argue that rejecting the evidentialist principle allows procedural responsibilism to better address certain problems that trouble evidentialism.

Let us begin with the skeptical problem. On Conee and Feldman's (2004) interpretation, skepticism is best understood as an evidentialist challenge, questioning whether our ordinary beliefs can meet the evidentialist standards of knowledge. According to evidentialism, S knows that p only if S's belief that p is based on evidence favoring p over ~p. Yet, the skeptic argues that my ordinary beliefs cannot satisfy this condition *even if* I live in the real world where my cognitive faculties are reliable. While attempting to refute skepticism, Conee and Feldman (2004: 305) admit that it is hard to meet the skeptical challenge. Suppose that I form the true belief that there are people in front of me via clear observation, and I have no counterevidence against this belief. Even if I live in the real world where my cognitive faculties are reliable, it is difficult to provide a detailed plausible explanation of why my evidence favors the real-world hypothesis (i.e. I live in the real world where my cognitive faculties are reliable, and the people I see are real) over various skeptical hypotheses (e.g. I am a brain in vat, and whatever I seem to see is unreal). Conee and Feldman (2004: 306) note that evidentialists might not be able to overcome this difficulty: 'If no [evidentialist] response … were to be correct, then skepticism would be true after all.'[6]

By contrast, procedural responsibilism can offer a neat response to skepticism. To see this more clearly, consider a variant of *Twins*:

> **Possible Twins:** As in the *Twins* case, you believe that the person is Judy solely based on observation, and you do not consider the possibility that she might have an identical twin sister. The only difference from the *Twins* case is this: Judy actually has no twin sister, nor does anyone else in the city look very much like her. To be sure, there is a possible world, W, where Judy has an identical twin sister called Trudy, and you mistake Trudy for Judy in the bookstore. You are

[6] For a related discussion, see Beebe (2009).

unable to rule out this possibility: Were you in the world W, you would still believe that the person is Judy (and thus that you are not in W).

In both *Twins* and *Possible Twins*, the way you form the belief that the person you see is Judy is epistemically responsible even though your total evidence does not favor <the person you see is Judy> over <the person you see is Trudy>. However, there is a significant difference between the two cases:

In *Twins*, the way you form the true belief is not *ultima facie* epistemically responsible. If you later happen to learn *the fact* that Judy has a twin sister, Trudy, who is also in the bookstore, then, intuitively, forming the belief that the person you see is Judy based solely on observation would no longer be epistemically blameless. Learning the fact generates an epistemic duty to rule out the possibility that the person is Trudy before forming your belief, for your epistemic community is entitled to expect you to rule out this possibility after acquiring the new information. If you persist in relying solely on observation, you would fail to rule out this possibility.

By contrast, in *Possible Twins*, the way you form the true belief is *ultima facie* epistemically responsible. Suppose you later learn that while Judy actually has no twin sister, there is a *possible* world (W) in which Judy has a twin sister called Trudy, and in which the person you see in the bookstore is Trudy. Even with this new information, you may still form the belief that the person you see is Judy based solely on observation. This is because, intuitively, the original way of forming the belief would remain epistemically blameless. To be sure, merely relying on observation makes you unable to rule out the possibility that you are in W: If you were in W, you would still believe that you are not in W. But you do not have any epistemic duty to rule out this possibility, for it would be unreasonable to expect you to do so.

This suggests a response to the skeptical challenge. Suppose I live in the real world: I am a regular person with well-functioning cognitive faculties. I form the true belief that there are people in front of me based on my visual experiences, and this belief is not

incoherent with my other beliefs. Then this belief is knowledge because the way I form the belief is *ultima facie* epistemically responsible: I form the belief in an epistemically blameless way, and there are no undefeated defeaters. Admittedly, there is a remote possible world in which I am a brain in a vat. Furthermore, I cannot rule out the possibility that I am currently in this world, because if I were, my visual experiences would still lead me to believe there are people in front of me, and thus that I am not in this world. However, even being aware of this possibility imposes no epistemic duty on me to rule it out. It would be unreasonable to expect me to do so, even if I am an epistemologist (Indeed, the internal rules of epistemological practice do not require epistemologists to rule out such skeptical possibilities). Thus, procedural responsibilism has a ready solution to the skeptical problem. (Note that skepticism claims that I cannot know anything about the external world *even if* I live in the real world where my cognitive faculties are reliable. To refute skepticism, we do not have to prove that I know certain things about the external world. Rather, it is sufficient to show that *if* I live in the real world where my cognitive faculties are reliable, then I can know something about the external world.)

Some might object that the very fact that procedural responsibilism provides a neat solution to the skepticism problem shows that procedural responsibilism is false. For it cannot explain why the skeptical problem has been a persistent topic in philosophy. Indeed, Conee and Feldman (2004) suggest that it is an *advantage* of evidentialism to fail to provide a neat solution to the skeptical problem, because evidentialism well explains why the skeptical problem has been a persistent topic: If the evidentialist standards of knowledge are correct, then it is really difficult to show how our ordinary beliefs can be knowledge.

But Conee and Feldman's argument seems flawed. Here is an alternative explanation of why the skeptical problem has been persistent: (a) Many philosophers assume the evidentialist principle that S knows that p only if S's total evidence favors p over ~p, and (b) this principle had not been adequately questioned before 1970. Indeed, influential

theories explicitly denying it—such as the relevant alternatives theory, the sensitivity theory, and the safety theory—emerged only afterwards. (a) and (b) do not entail that evidentialism is true. This explanation seems no less plausible than Conee and Feldman's view that it is the truth of evidentialism that explains why the skeptical problem has been persistent.

In addition, procedural responsibilism can avoid two other problems for evidentialism. Evidentialism claims that knowledge requires justification, and it endorses a synchronic theory of justification: Whether one's belief is justified at a certain time merely depends on the evidence one has at that time. Accordingly, (i) If one previously formed a true belief on sufficient evidence but forgets the evidence now, then one's belief is unjustified and falls short of knowledge. (ii) Even if one happens to form a true belief on sufficient evidence at t because one did not gather evidence responsibly, one's belief might still be justified and amount to knowledge at t.

However, both (i) and (ii) seem to have implausible consequences. Goldman (1999: 280) provides a counterexample to (i):

> **Forgotten Evidence:** Last year, Sally read a story about the health benefits of broccoli in the 'Science' section of the *New York Times*. She then justifiably formed a belief in broccoli's beneficial effects. She still retains this belief but no longer recalls her original evidential source (and has never encountered either corroborating or undermining sources). Nonetheless, her broccoli belief is still justified, and, if true, qualifies as a case of knowledge.

While evidentialists may add a plug-in or two to address this case (cf. Conee and Feldman 2004), procedural responsibilism can provide a straightforward explanation of why Sally still knows. It is because the way she retains the belief is *ultima facie* epistemically responsible: She is epistemically blameless in retaining the belief, and there are no undefeated defeaters. (Following Fleisher (2025), we may say that Sally engages in an epistemic practice where people rely on trusted media to know scientific facts, with

rules tailored to accessibility and practicality. Sally follows the rules—trust the article, don't worry about forgetting it—which work well enough for everyday truth.)

The implausibility of (ii) has also been widely discussed (cf. Kornblith 1983; Axtell 2008; Baehr 2009; Cloos 2015; Goldberg 2018). Consider the following case:

> **Ignoring Evidence:** Smith is a scientist who believes a true scientific theory on the basis of evidence E at t1. Suppose that E is sufficient evidence for Smith to believe the theory at t1. Yet at t2, some leading scientists co-publish a paper that challenges the theory. The paper causes most scientists to believe that the theory is false because they do not recognize that the evidence the paper presents is misleading. But Smith firmly believes the theory at t2. For he has become intellectually arrogant and refuses to follow any recent research on the topic. He is unaware of the paper as well as other scientists' new opinion on the theory. Had he read the paper, he would, like most scientists, think the theory is false.

Intuitively, Smith does not know that the theory is true at t2. Yet evidentialism seems to imply that Smith knows, for at t2, E is still sufficient evidence for him to believe the theory. Further, there is no undefeated defeaters: The paper only presents misleading information.

In contrast, responsibilism can easily explain why Smith does not know: It is because the way Smith has retained his belief in the theory is epistemically irresponsible in the first place. He has an inexcusable failure to fulfil his role-based epistemic duty to be up-to-date in his field.[7]

[7] Some evidentialists claim that Smith's retention of his belief is prudentially irresponsible, but not epistemically so. For recent objections to this claim, see McWilliams (2023) and Flores & Woodard (2023).

3. Against virtue epistemology

One might also worry that procedural responsibilism collapses into a form of virtue epistemology that integrates both virtue reliabilist and responsibilist components. John Greco defends such a form of epistemology. He argues that knowledge is a kind of success for which the knower deserves credit, for an agent knows that p iff she believes the truth that p 'because she is intellectually able and because she has exercised her abilities' (Greco 2010: 97). Greco uses 'intellectual ability' and 'intellectual virtue' interchangeably. His account of intellectual virtue is modeled on Aristotle's account of moral virtue, which involves 'both a motivational component and a reliable success component: a virtuous person is motivated toward the good, but is also successful in achieving it' (Greco 2010: 43). Accordingly, Greco claims that S's belief that p is intellectually virtuous iff two conditions are satisfied:

(a) S's belief that p is epistemically responsible in the sense that 'S's believing that p is properly motivated, …[i.e.] S's believing that p results from intellectual dispositions that S manifests when S is motivated to believe the truth' (*ibid*.).
(b) S is objectively reliable in believing that p: 'S's believing that p results from intellectual dispositions that reliably produce true belief' (*ibid*.).

Now one might contend that the way S forms/retains a belief is *ultima facie* epistemically responsible iff both (a) and (b) are true. To be sure, the way S forms/retains a belief might be epistemically responsible yet unreliable, as the new evil demon problem shows. But one might think if the way S forms/retains a belief is *ultima facie* epistemically responsible, then it must be reliable. For had S known that the original way she forms/retains the belief is unreliable, forming/retaining the belief in the original way would be epistemically irresponsible: She ought to suspend judgment. Thus, procedural responsibilism is equivalent to Greco's virtue epistemology.

In what follows, I will explain why procedural responsibilism is not equivalent to Greco's virtue epistemology. Given Greco's view, both (a) and (b) are necessary for

knowledge. I will show that procedural responsibilism entails that neither (a) nor (b) are necessary for knowledge. I will also argue that procedural responsibilism offers a better solution to the lottery problem than other versions of virtue epistemology that also reject (a) and (b) as necessary conditions for knowledge.

Let's begin with Condition (a), which seems to imply that if S is not motivated to believe the truth about whether p, then S does not know that p. This view is shared by two variants of responsibilist virtue epistemology: (i) Robust Virtue Responsibilism: Knowledge requires possessing and exercising responsibilist virtues, which involve proper motives;[8] (ii) Moderate Virtue Responsibilism: While one can have knowledge without fully developed responsibilist virtues, knowledge entails belief arising from acts of responsibilist virtues, which involve proper motives (Zagzebski 1996).

However, the view that knowledge requires proper motives cannot explain passive knowledge, as Baehr (2011) notes. Imagine a scenario where I'm working late at my desk one night, and suddenly the power goes out, plunging the room into darkness. As a result, I would immediately and automatically form the belief that the room's lighting has changed. This belief would arise involuntarily, regardless of what my intellectual motives happen to be at that moment.

By contrast, procedural responsibilism can explain passive knowledge because it denies that knowledge requires proper motives. Specifically, while it claims that knowledge requires epistemic responsibility, it defines epistemic responsibility in terms of epistemic blamelessness rather than epistemic praiseworthiness.[9] One can form a belief in an epistemically blameless way without being motivated to believe the truth, as the above case of passive knowledge shows. (For similar reasons, procedural responsibilism denies that knowledge requires possessing responsibilist virtues. Accordingly, it can

[8] Code (1987) and Montmarquet (1993) seem to suggest this view though they do not provide an explicit definition of knowledge.
[9] Weatherson (2008) argues that S is epistemically responsible in forming/retaining a belief only if S is epistemically *praiseworthy* (rather than merely blameless) in forming/retaining the belief. For a reply, see Peels (2017).

avoid the situationist challenge, which claims that most people do not possess responsibilist virtues [cf. King 2014]).

Still, one might ask: If Greco understood epistemic responsibility in terms of epistemic blamelessness rather than proper motives, would his view be equivalent to procedural responsibilism?[10] The answer is negative, for procedural responsibilism also rejects the idea that knowledge requires reliable abilities. Consider the following case:

> **Unreliable Memory:** Ron suffers a brain tumor, which leads to unreliable memory about his childhood: Even if Ron clearly 'remembers' an event in his childhood, there is 60% chance that the event never happened. Aware that his memory is unreliable, Ron does not suspend judgment. Instead, Ron adopts the following policy: Believe that p about this childhood on the basis of memory iff (i) p is not morally or practically significant, (ii) it seems to him that he clearly remembers that p, and (iii) he has no other evidence against p (e.g. no childhood friends testify against p). Adopting this policy, which is still unreliable, leads him to believe that he aced his 4th grade final exam. His belief is true: He does remember that he aced his 4th grade final exam (cf. Lai 2024).

Intuitively, Ron knows that he aced his 4th grade final exam. Procedural responsibilism provides a ready explanation: He believes the truth because the way he forms the belief is *ultima facie* epistemically responsible. Intuitively, there is nothing blameworthy in his way of forming the belief, and there are no undefeated defeaters (note that Ron is fully aware that his memory is unreliable). However, the way Ron forms the belief is unreliable. It follows that Ron's success in believing the truth is not attributable to his intellectual virtue, which requires more than 50% reliability. Thus, Unreliable

[10] Elsewhere in his 2010 book, Greco (2010: 77) appears to downplay the motive condition and defines ability solely in terms of reliability.

Memory is a counterexample to any version of reliabilist virtue epistemology claiming that knowledge in all contexts requires more than 50% reliability.[11]

Procedural responsibilism would enjoy a greater advantage over reliabilist virtue epistemology if scientific knowledge is often achieved via an unreliable method. Scientists employ inference to the best explanation, the method of choosing the hypothesis or theory that best explains the available data. Some philosophers have argued that inference to the best explanation is unreliable. For example, John Turri (2015: 536) writes,

> Why think the best explanation might not be very likely? Let 'D' name a robust set of well confirmed data and 'H' the hypothesis that best explains it. What makes H the best explanation? Take the set of possible worlds where D obtains. Call these the 'D-worlds'. Take the set of worlds where H explains D. Call these the 'H»D worlds'. Now suppose a plurality of D-worlds are H»D worlds. This makes H the most likely and so best explanation of D. But H»D is still not likely given D (in the same sense of likelihood). A plurality of D-worlds are H»D-worlds, but most D-worlds are not H»D-worlds.

If Turri is right, then scientists can sometimes know, via unreliable reasoning, that H explains D. This would pose a serious challenge for reliabilist virtue epistemology, as Turri (2015) notes. But procedural responsibilism can easily explain why scientists can sometimes know that H explains D via inference to the best explanation: It is because forming the true belief via this method is *ultima facie* epistemically responsible. Even if the method is less than 50% reliable and scientists realize that it is less than 50% reliable, they may still form/retain the belief via the method: There is nothing blameworthy in continuing to use it. This aligns with Fleisher's epistemic practice account. Scientists choose a specific set of rules from multiple permissible options that adequately promote truth. Even rules with less than 50% reliability can be adequate. In some cases, such rules

[11] Sosa (2015: 175–176) explicitly endorses this claim though he agrees that knowledge in some contexts might require a higher degree of reliability than knowledge in other contexts.

may be the most reliable options available for addressing particular questions. In others, scientists may recognize more reliable alternatives but opt for less reliable rules due to practical constraints, such as time pressure or implementation challenges. Fleisher (2025) highlights that the threshold for adequacy is context-dependent, shaped by factors like coordination needs, the research field, the questions at hand, and situational limitations.[12]

One might point out that Turri's abilism, which states that knowledge is true belief resulting from the exercise of cognitive abilities, can also explain unreliable knowledge and passive knowledge, for cognitive abilities do not have to be more than 50% reliable (Turri and Greco use 'abilities' differently), and exercising such abilities does not involve having proper motives. However, abilism cannot well accommodate our intuition about the lottery case:

> **Lottery:** You are very good at reasoning. After drawing a lottery, you recognize that the odds of winning are extremely low. Thus, before the official results are announced, you form the belief that you will lose. Suppose you are correct: Your number is 19651211, while the winning number will be 20080405.

Here, your success in believing the truth that you will lose is attributable to your intellectual virtues (i.e. reliable cognitive abilities). Thus, Turri's abilism, like Sosa's (2015) reliabilist virtue epistemology, implies that you know that you will lose. Yet, intuitively, you do not really know that you will lose before the official results are announced. As Beddor & Pavese (2020) note, this intuition is widely attested among both philosophers and non-philosophers.

In contrast, procedural responsibilism can accommodate this intuition: You do not know that you will lose because the way you form the belief is not *ultima facie* epistemically responsible: Had you known that *the* winning number will be 20080405, forming the belief merely based on the statistical fact (that the odds of winning are

[12] As Fleisher (2025) notes, his account implies that the consequentialist account of epistemic duty is false: The internal rules of a practice may forbid epistemic trade-offs. Cf. Berker (2013).

extremely low) would be epistemically blameworthy. This is because the statistical fact equally supports that 20080405 will not be the winning number. Had you known that *the* winning number will be 20080405, you should form the belief that you will lose based on the fact that your number is different from the winning number, otherwise you would be inconsistent: by applying the original belief-forming method, you would believe that 20080405 is not the winning number.

One might think if Turri's abilism incorporates a safety condition, then it can accommodate our intuition about the lottery case.[13] Admittedly, Sosa (2015) thinks that your belief in the lottery case is safe, because he holds that a belief is safe iff it holds true in most of these close worlds where the agent has that belief. In the lottery case, the odds are astronomically against you: In most of these close worlds where you believe that you will lose, you are correct. By contrast, Beddor & Pavese (2020: 61) propose that S's true belief is safe iff 'in all relevantly close worlds where S forms a sufficiently similar belief, S's belief is true.' The relevantly close possible worlds are those in which the conditions are at least as normal for the task at hand as those that obtain in the world of evaluation. In the lottery case, the task at hand for you is evaluating the likelihood of your ticket winning based on the statistical odds. In a close possible world where your ticket actually wins, the conditions can be just as normal for the task at hand as the world in which your ticket loses. Thus, your belief that you will lose is unsafe. Turri's abilism can incorporate a safety condition along the lines suggested by Beddor and Pavese.

However, Beddor & Pavese's account of safety seems to lead to inductive skepticism. Consider a case provided by Sosa (2015):

> **Chute:** You drop a bag of rubbish into the garbage chute next to your high-rise apartment, and a few moments later you form the true belief that the rubbish is

[13] Some philosophers think virtue epistemology should incorporate a safety condition. For example, Pritchard (2010) proposes that knowledge is a safe true belief as a result of exercising our intellectual virtues, while Beddor & Pavese's (2020) defines knowledge as a maximally skillful performance, and skillfulness (and by extension, intellectual virtue), in turn, is defined in terms of safety.

> now in the basement. You form this belief via inductive reasoning: You know that the rubbish chute is in fact very reliable in this regard – indeed, it has never failed to deliver rubbish to the basement, over a long history – and it is well maintained and serviced. Moreover, there is nothing wrong with the rubbish chute on this occasion, nor any reason for you to worry about the reliability of the rubbish chute in this specific instance.

Intuitively, you know that the rubbish is now in the basement. Indeed, this seems to be a paradigm case of inductive knowledge. However, your belief is unsafe according to Beddor & Pavese's account. In a close possible world where the rubbish is stuck in the middle, the conditions can be just as normal for the task at hand as the world in which the rubbish is now in the basement.[14]

Procedural responsibilism can accommodate our intuition about Chute. You know that the rubbish is now in the basement because you form the true belief in an *ultima facie* epistemically responsible way: The way you form the belief is clearly blameless, and there are no undefeated defeaters. The possibility that the rubbish is stuck in the middle is not a defeater: You already realize this possibility when you form the belief via the *inductive* reasoning. In addition, even if you later acquire stronger evidence (e.g. you go to the basement and see in person that the rubbish is there), this stronger evidence would not undercut the strength of your old statistical evidence, and it thereby would not make the original way you form the belief epistemically blameworthy. (If you later believe that the rubbish is the basement on the basis of both the statistical evidence and your observation, that would be epistemic *supererogation*. You are not epistemically blameworthy if your belief is still based on mere statistical evidence.) Here Chute is different from the lottery

[14] Beddor & Pavese (2020: 72) admit that their normality-based account of safety has difficulties in addressing the chute case, Comesaña's Halloween Party case, Baumann's bank robber case, etc. For recent objections to the safety approach, see Kelp (2016), Zhao (2020), Zhao (2022), and Goldstein & Hawthorne (2024).

case, where knowing that the winning number will be 20080405 would destroy the effect of your old statistic evidence.

4. Conclusion

I have argued that procedural responsibilism is a viable competitor to some influential forms of evidentialism and virtue epistemology. In closing, I'd like to briefly address the worry that there might be counterexamples to procedural responsibilism. My response is that my argumentative strategy is similar to the one Beddor & Pavese (2020) endorse in their paper 'Modal virtue epistemology', where they propose a normality-based account of safety. They note that their account cannot easily accommodate our intuitions about certain cases. But they emphasize that they merely claim that their normality-based account can 'ward off at least some of the standard counterexamples to modal accounts' (Beddor & Pavese 2020: 72). Their claim is 'not that a normality-based framework, on its own, handles all of the apparent counterexamples to the necessity of modal robustness for skillfulness and knowledge' (*ibid.*). Similarly, my claim is not that procedural responsibilism can, on its own, handle all of the potential counterexamples. Rather, I merely aim to show that procedural responsibilism enjoys advantages over some influential forms of evidentialism and virtue epistemology. This modest goal also explains why I define epistemically responsible belief in terms of doxastic responsibility and epistemic blamelessness *without* giving a detailed account of these concepts: To achieve this goal, appealing to an intuitive understanding of doxastic responsibility and epistemic blamelessness is sufficient. (Compare: Sosa (2022: 6) explicitly notes that while the concept of manifestation is key to his virtue epistemology, he does not give 'a fuller explication of the notion of manifestation.' Rather, he takes the concept of manifestation as primitive (cf. Sosa 2015). He suggests that without a detailed account of manifestation, he can still show that his virtue epistemology is better than some alternative accounts of knowledge.)

Acknowledgements

I am grateful to Anqi Fan, Bin Zhao, Haicheng Zhao, and two anonymous reviewers for their helpful feedback.

References

Axtell, G. (2008) 'Expanding epistemology: A Responsibilist Approach', *Philosophical Papers*, *37*/1: 51-87.

Baehr, Jason (2009) 'Evidentialism, Vice, and Virtue', *Philosophy and Phenomenological Research,* 78/3: 545-67.

Baehr, J. (2011) *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford: OUP.

Beddor, Bob & Pavese, Carlotta (2020) 'Modal Virtue Epistemology', *Philosophy and Phenomenological Research,* 101/1: 61-79.

Beebe, James R. (2009) 'The Abductivist Reply to Skepticism', *Philosophy and Phenomenological Research,* 79/3: 605-36.

Berker, S. (2013) 'The Rejection of Epistemic Consequentialism', *Philosophical Issues*, 23/1: 363-87.

BonJour, L. (1985) The Structure of Empirical Knowledge. Cambridge, MA: Harvard University Press.

Boult, C. (2024) 'Degrees of Epistemic Criticizability', *Philosophical Quarterly,* 74/2: 431-52.

Brogaard, Berit (2023) 'Virtue Epistemology and Epistemic Responsibility', in Luis R. G. Oliveira (ed.) *Externalism about Knowledge*, 213–46. Oxford: OUP.

Chisholm, R. (1977) *Theory of Knowledge*, 2nd edition. Englewood Cliffs, NJ: Prentice-Hall.

Cloos, Christopher Michael (2015) 'Responsibilist Evidentialism', *Philosophical Studies,* 172/11: 2999-3016.

Code, Lorraine (1987) *Epistemic Responsibility*. Hanover, NH: University Press of New England and Brown University Press.

Conee, E. and Feldman, R. (2004) *Evidentialism: Essays in Epistemology*. Oxford: Clarendon Press.

Corlett, J. A. (2008) 'Epistemic Responsibility', *International Journal of Philosophical Studies*, *16*/2: 179-200.

de Almeida, Claudio & Fett, J. R. (2016) 'Defeasibility and Gettierization: A Reminder', *Australasian Journal of Philosophy,* 94/1: 152-69.

Fleisher, Will (2025) 'Epistemic practices: A Unified Account of Epistemic and Zetetic Normativity', *Noûs,* 59/1: 289-314.

Flores, Carolina & Woodard, Elise (2023) 'Epistemic Norms on Evidence-gathering', *Philosophical Studies,* 180/9: 2547-71.

Goldberg, S. (2018) *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford: OUP.

Goldstein, Simon & Hawthorne, John (2024) 'Safety, Closure, and Extended Methods', *Journal of Philosophy,* 121/1: 26-54.

Greco, J. (2010) *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. New York: CUP.

Kelp, C. (2016) 'Epistemic Frankfurt Cases Revisited', *American Philosophical Quarterly,* 53/1: 27-37.

King, Nathan L. (2014) 'Responsibilist Virtue Epistemology: A Reply to the Situationist Challenge', *Philosophical Quarterly*, 64/255: 243–53.

Kornblith, Hilary (1983) 'Justified Belief and Epistemically Responsible Action', *Philosophical Review,* 92/1: 33-48.

Lai, Changsheng (2024) 'Remembering Requires No Reliability', *Philosophical Studies,* 181/1: 43-63.

Lehrer, Keith & Paxson, Thomas (1969) 'Knowledge: Undefeated Justified True Belief', *Journal of Philosophy,* 66/8: 225-237.

Lehrer, K. (1990) *Theory of Knowledge*. London: Routledge.

McWilliams, Emily C. (2023) 'Evidentialism and Epistemic Duties to Inquire', *Philosophical Quarterly,* 73/4: 965-982.

Montmarquet, James A. (1993) *Epistemic Virtue and Doxastic Responsibility*. Lanham: Rowman and Littlefield.

Peels, Rik (2017) *Responsible Belief: A Theory in Ethics and Epistemology*. New York: OUP.

Pritchard, D (2010) 'Knowledge and Understanding', in D. Pritchard, A. Millar, & A. Haddock (ed.) *The Nature and Value of Knowledge: Three Investigations*, 3-86. New York: OUP.

Sosa, E. (2015) *Judgment and Agency*. New York: OUP.

Sosa, E. (2022) 'John Greco's *The Transmission of Knowledge*', *Synthese,* 200/4: 1-11.

Strawson, P. F. (1992) *Analysis and Metaphysics: An Introduction to Philosophy*. New York: OUP.

Weatherson, B. (2008) 'Deontology and Descartes's Demon', *Journal of Philosophy,* 105/9: 540-69.

Williams, M. (2008) 'Responsibility and Reliability', *Philosophical Papers*, 37/1: 1-26.

Williamson, T. (2000) *Knowledge and Its Limits*. New York: OUP.

Zagzebski, L. T. (1996) *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: CUP.

Zhao, B. (2022) 'A Dilemma for Globalized Safety', *Acta Analytica,* 37/2: 249-61.

Zhao, H. (2020) 'Knowledge without Safety', *Synthese,* 197/8: 3261-78.