

Reliable Knowledge

A Reply to Turri

JONATHAN DIXON

doi:10.48106/dial.v74.i3.03

Jonathan Dixon. 2020. "Reliable Knowledge: A Reply to Turri." *Dialectica* 74(3): 495–508.
doi:10.48106/dial.v74.i3.03.



Reliable Knowledge

A Reply to Turri

JONATHAN DIXON

Recently John Turri (?) has argued, contra the orthodoxy amongst epistemologists, that reliability is not a necessary condition for knowledge. From this result, Turri (? , ? , ?) defends a new account of knowledge—called *abilism*—that allows for unreliable knowledge. I argue that Turri’s arguments fail to establish that unreliable knowledge is possible and argue that Turri’s account of knowledge is false because reliability must be a necessary condition for knowledge.

Many epistemologists agree that knowledge must be reliably produced. For example, Goldman holds that justification is necessary for knowledge and that justification “is a function of the reliability of the process or processes that cause it” (1979, 345); Sosa holds that knowledge is produced by a disposition “that would in appropriately normal circumstances ensure (or make very likely) the success of any relevant performance issued by it” (2007, 29); and Williamson claims that “no reason has emerged to doubt the intuitive claim that reliability is necessary for knowledge” (2000, 100).¹ Recently John Turri (2015b) argued against this orthodoxy by providing two theoretical arguments for the possibility of unreliably produced knowledge. If either of Turri’s arguments is sound then all accounts of knowledge that require reliability are false and most epistemologists have been on the wrong track in understanding the nature of knowledge. Realizing this, Turri (2015a, 2017, 2016a, 2019) defends a new account of knowledge, called *abilism*, which allows for knowledge to be unreliably produced.

¹ See Turri (2015b, footnote 1) and Goldman and Beddor (2021) for a more complete list of epistemologists who think that knowledge must be reliably produced. However, there are a few contemporary philosophers who indirectly deny this claim. Sartwell (1991, 1992) argues that true belief alone is sufficient for knowledge; and because Hetherington (1998, 1999, 2016) argues that victims of Gettier-cases do possess the relevant knowledge, *a fortiori*, he holds that lucky (and so unreliable) processes can produce knowledge.

After providing some background and clarifying terms in § 1, in § 2 and § 3 I explain why each of Turri's (2015b) theoretical arguments for unreliable knowledge fail. And I conclude in § 4 with reasons why abilism is false and why reliability must be a necessary condition for knowledge.

1 Background and Clarifying Terms

Turri's (2015b) theoretical arguments for unreliable knowledge rely on what is called an *achievement account of knowledge*. This is roughly the family of views which hold that an agent S has knowledge of P just in case S's true belief in P manifests S's cognitive achievement.² While there are many ways of spelling-out the details of this account of knowledge and there are many challenges to this family of views,³ I will set these issues aside and grant for the sake of argument that knowledge is a kind of cognitive achievement. My arguments below show that *even if* we grant this, both of Turri's (2015b) arguments for the possibility of unreliable knowledge fail.

The next thing I should explain is what Turri means by "reliability" and "achievement." Turri's definition of "reliability" is in line with how it is standardly understood: a process, disposition, or ability is (epistemically) reliable when and only when (significantly) more than half of its produced beliefs are true; and a process, disposition, or ability is (epistemically) *unreliable* when and only when less than half of its produced beliefs are true (2015b, 530).⁴ While Turri (2015b) does not provide a definition of "achievement," the important thing for Turri is that *achievements need not be reliably produced* because "achievement can issue from even highly *unreliable* ability" (2015b, 531). An agent has an unreliable ability to Φ iff in using this ability to Φ the agent fails to Φ most of the time. For example, a novice musician who plays a chord for the first time, a child who takes his first step or speaks his first sentence, and a rookie golfer who makes par for the first time are all examples of achievements for Turri even though they fail to achieve their desired ends most of the time (2015b, 531–532). In sum, for Turri, achievements involve

2 This general account of knowledge is part of the ongoing research program in epistemology called *performance-based epistemology* which is exemplified by virtue epistemology. See Sosa (2007), Zagzebski (2009), and Greco (2010) for more details on this research program, virtue epistemology, and achievement accounts of knowledge. See Bradford (2015) for a general account of achievements.

3 For example, see Pritchard (2008, 2009) and Lackey (2007, 2009).

4 See Alston (1995) for a detailed characterization of epistemic reliability.

simply attaining one's intended outcome through one's (un)reliable process, disposition, or ability. I will also assume this understanding of "reliability" and "achievement" in what follows.

Lastly, it is worth pointing out that Turri's account of achievement is unique among those who hold an achievement account of knowledge because it does not require that achievements manifest one's *competence* which involves the reliability of processes, dispositions, or abilities (see Sosa 2007, 2015; Zagzebski 2009; and Greco 2010). Turri (2016b) explicitly points out this omission and Turri (2015a, 2017, 2016b; and 2019) endorses this as a beneficial feature of his achievement account of knowledge because it avoids problems Turri sees for these authors' accounts of knowledge.

2 Against Turri's First Argument

Turri's first argument for the possibility of unreliable knowledge is

1. Achievements don't require reliable abilities. (Premise)
2. If achievements don't require reliable abilities, then unreliable knowledge is possible. (Premise)
3. So unreliable knowledge is possible. (From 1 and 2) (2015b, 531)⁵

Turri supports the [first premise](#) by referencing the examples he provides of achievements issuing from unreliable abilities mentioned above. Turri supports the [second premise](#) by saying that if knowledge is a kind of intellectual achievement and achievements generally do not necessarily issue from reliable processes, abilities, or dispositions, then "absent a special reason to think otherwise, we should expect [knowledge] to share the profile of achievements generally" (2015b, 532). In short, Turri's argument attempts to shift the burden of proof on those who believe reliability is a necessary condition of knowledge to show why knowledge, as an intellectual achievement, cannot issue from unreliable abilities.

Turri's [first argument](#) fails to convincingly shift the burden of proof because it faces a dilemma: Either the [first premise](#) is false or the argument as a whole begs the question. The [first premise](#) is false if it is interpreted to mean "all achievements don't require reliable abilities." There are many achievements

⁵ Turri (2015b, 2015a, 2017, 2016a, 2016b; and 2019) never specifies how "possibility" should be understood. I likewise will not assume any particular account of "possibility."

that require reliable abilities. More specifically, achieving some goal often requires reliably performing some action. For example, winning a competitive darts or archery tournament often requires one to reliably hit their intended mark.⁶ Indeed, achieving the goal of performing some action with 90%+ accuracy (e.g. hitting a bullseye in archery, hitting a baseball, playing a piece of music, or walking) requires performing this action with 90%+ accuracy. So, the proper interpretation of the **first premise** must be something like “*some* achievements don’t require reliable abilities.” However, if this interpretation is placed back into the argument above then it begs the question. The **second premise** would now read “if *some* achievements don’t require reliabilities, then unreliable knowledge is possible.” But since Turri has said nothing against the possibility that knowledge is the kind of intellectual achievement that requires reliability (like the ones listed above), Turri has not provided adequate reasons to think that *knowledge* is the kind of achievement that can be unreliably produced—which is the purpose of the argument. So, in order for **this argument** to conclude “unreliable knowledge is possible,” it must beg the question and consequently fails to shift the burden of proof.

Turri anticipates and responds to this dilemma⁷ by claiming that it can be avoided if we interpret the **first premise** as a proposition “about dominant tendencies, or what is typical, or what is natural and normal for a kind” (2015b, 534). For example, the propositions that “humans don’t have eleven fingers” or “cats don’t have two faces” express tendencies about how humans and cats’ anatomy are typically constituted (Turri 2015b, 534). Although there are exceptions to these claims, these exceptions do not render these claims false when these claims express such tendencies. So, if **premise one** is understood as a tendency proposition, Turri claims his argument “would still be plausible because, as already mentioned, we would expect knowledge to fit the profile of achievements generally, unless we’re given a special reason to think otherwise” (2015b, 534).

This response still fails for the reasons mentioned above. Even if we grant that **premise one** is a tendency proposition, Turri has not established that achievements have a general tendency to be unreliable. As argued above, there are a large number of achievements that require reliability. Turri’s few examples of unreliable achievements are insufficient to establish that **premise**

6 However, in order to achieve some goal, one need not reliably achieve that goal (e.g. to win the archery competition one need not reliably win the archery competition). Thanks to an anonymous reviewer for helping to clarify this.

7 Turri (2015b, fn. 7) attributes this dilemma to Bruce Russell.

one is a tendency proposition. Furthermore, Turri has provided no positive reason to think that *knowledge* is kind of achievement that can be unreliably produced—which (again) is the purpose of the argument. So, **Turri's first argument** fails to shift the burden of proof because it either has a false premise or begs the question.

A better strategy for Turri to establish that unreliable knowledge is possible is to take a more direct route by providing an example where one intuitively knows some proposition P even though one's true belief that P was formed by an unreliable cognitive process, i.e. one that produces more false than true beliefs. This is what Turri's second argument for unreliable knowledge attempts to do. In § 4 I will take on the burden of proof and argue that reliability is a necessary condition for knowledge.

3 Against Turri's Second Argument

Turri's second and more direct argument for the possibility of unreliable knowledge involves explanatory inference (aka, inference to the best explanation or IBE). As Turri notes, IBE is used in scientific reasoning and in everyday life to provide probable explanations for a set of data or certain phenomena. What best explains the fact that humans and chimpanzees have so many anatomical similarities? We have a common ancestor. What best explains the appearance of a new jug of milk in the fridge? My spouse bought it at the store. Turri claims that this kind of reasoning supports the possibility of unreliable knowledge:

The epistemic efficacy of explanatory inference supports the view that unreliable knowledge is possible. Inference to the best explanation yields knowledge if the explanation that we arrive at is true. But even when it is true, the best explanation might not be very likely. So our disposition to infer to the best explanation might not be reliable. So unreliable knowledge is possible. (2015b, 536)

That is, even though IBE is often unreliable, the explanations it provides (when true) can yield knowledge. More specifically, some hypothesis "H" can

best explain a set of data “D” in our world even if there is a greater number of (nearby) possible worlds where D obtains and H is false (Turri 2015b, 536–537).⁸

To illustrate this argument, Turri provides a case study involving the television show *House M.D.* Gregory House (the protagonist) is a world-renowned medical doctor who has an incredible ability to diagnose patients where other doctors have failed. Simply put, he is the best of the best. However, despite being the best, House misdiagnoses patients a lot. Indeed, nearly every episode follows the same structure where House misdiagnoses the patient several times before coming to the right diagnosis just in the nick of time to save the patient’s life. Turri contends that House’s method for diagnosing patients is IBE—House infers a hypothesis/diagnosis that best explains the data/symptoms. And with each failed diagnosis House gains new insights to symptoms that inform his subsequent diagnoses. Given this description of House’s track record, Turri argues that House’s reliability is considerably less than .5. But despite House’s unreliability, when he ends up correctly diagnosing his patient “House knows what disease that patient has” (Turri 2015b, 538). In short, this case study shows that IBE “can yield knowledge, even though it doesn’t yield the correct verdict most of the time” (Turri 2015b, 539). Turri summarizes his second argument as

1. If House knows, then unreliable knowledge is possible.
(Premise)
2. House knows. (Premise)
3. So unreliable knowledge is possible. (From 1 and 2)

The argument is valid. Line 1 is supported by the fact that House’s method usually produces false beliefs. Line 2 is supported by intuition, and by the fact that millions of viewers, including trained epistemologists, detect no incoherence in the story line, week after week, over many seasons. (2015b, 539)

I believe that both premises of **Turri’s second argument** are false because Turri misrepresents House’s medical abilities and knowledge. While Turri is right that House’s diagnostic track record is well below .5, Turri takes the lesson here to be that, despite his track-record, “House knows” the correct diagnosis when he gets it right via IBE because House has a special ability to figure

8 There is controversy about whether IBE can provide explanations and/or produce knowledge (e.g. Fraassen 1989). So, Turri’s second argument has the important caveat that one must first accept that *IBE can produce explanations / knowledge* before this argument can be persuasive.

out the right diagnosis more often than any other doctor. This misrepresents House's abilities because, contra Turri, House is remarkable at getting the right diagnosis not because he knows the correct diagnosis more often than any other doctor, but because he has a remarkable ability to propose novel diagnostic hypotheses worthy of consideration and testing. But this ability to come up with possible explanations of patient's symptoms does not itself allow House to know that his diagnoses are correct *until* the treatment actually works (or when the reliable test results confirm his diagnosis).⁹

To illustrate these points, consider the following case that parallels Turri's House example:

Jessica has very poor eyesight and is legally blind without her glasses. However, despite her eyesight, Jessica has a special ability to correctly identify pictures without her glasses. While others who are similarly handicapped can only identify pictures 5% of the time on average, Jessica is able to correctly identify such images 25% of the time on average. Now imagine that Jessica is presented with an image of a basketball that she, and others with her eyesight, phenomenologically describes as a blurry spot of reddish orange. Without her glasses Jessica infers incorrectly three times in a row that the picture is of an orange fruit, the Sun, and then a Lego piece. After each incorrect answer or hypothesis Jessica is told new information about the image that reveals why her answers were incorrect, e.g. it is not a fruit for her orange fruit hypothesis, it is an object you can touch for her Sun hypothesis, and it is an object that is bigger than a Lego piece. After all of this Jessica then answers correctly, but is not yet told that she is correct.

The crucial question to now ask is: At this point, does Jessica *know* what the picture is of? Intuitively, the answer is no. While Jessica, like House, has a special ability to get it right more often than her peers, this is not because she knows the correct answer more often, but because she is better at coming up with worthy hypotheses.¹⁰ And, like House, Jessica does not know

⁹ I am indebted to Hilary Kornblith for a discussion on these points.

¹⁰ This misrepresentation of House's abilities is related to another debate concerning the nature of IBE. There is a tradition, going back to Pierce, of distinguishing abduction from IBE. Traditionally, abduction is concerned with hypothesis construction while IBE is concerned with selecting the hypothesis that is most likely to be true from a set of hypotheses. And traditionally, it is a feature

her hypothesis is correct until it's confirmed. Thus, **premise two** of **Turri's argument** is false because before the proposed treatment works (or when a reliable test result confirms a diagnosis) House does not know whether his hypothesized diagnosis is correct. **Premise one** is also false because if we plug this understanding of what House knows back into the antecedent of this premise, it renders the consequent false. That is, if "House knows" is understood to be true only after his hypothesized diagnosis has been tested and confirmed, then House's knowledge is not an instance of unreliable knowledge.¹¹

4 Why Reliability is a Necessary Condition for Knowledge

So far, I have argued that Turri (2015b) has not provided adequate reasons to reject the orthodox view that knowledge requires reliability. In this final section I will directly argue against Turri's (2015a, 2017, 2016a, 2019) abilist account of knowledge¹² and argue that reliability must be a necessary condition for knowledge.

of abduction to be unreliable in order to produce a variety of hypothesis to be tested. Furthermore, there has been a recent trend to conflate these two (see McAuliffe 2015 for a defense of these points). On this understanding, House does not use IBE but uses abduction and Turri conflates these two when he writes "House and his team explicitly reason abductively" (2015b, 537) and "House's method for trying to solve the case *just is* to employ inference to the best explanation" (2015b, 540, his emphasis).

- 11 Turri's argument that IBE can produce unreliable knowledge also concerns two other recent debates about the nature of IBE. First, it is related to the issue of whether van Fraassen's Bad Lot objection shows that IBE is an unreliable inference form. The Bad Lot Objection argues that IBE is an inadequate inference form because it has no way of discerning whether a set of hypotheses are all false and so would lead to a false conclusion in these cases. Schupbach (2014) recently argues that van Fraassen's Bad Lot Objection does not establish that IBE is an unreliable inference form any more than we can show that modus ponens is unreliable by plugging in false premises (see Dellsén 2017 for a response). If this were true, then it would be inappropriate to characterize either IBE or modus ponens as unreliable. Second, many have argued that IBE is a heuristic that approximates objective Bayesian reasoning. On this understanding, IBE's primary function is to locate the "most probable available explanatory hypothesis to serve as a working hypothesis in an agent's further investigations" (Dellsén 2018). This understanding dovetails with my explanation of House's use of IBE above. In short, these issues concerning the nature and reliability of IBE impact and potentially undermine the soundness of Turri's argument that IBE can unreliably produce knowledge.
- 12 Turri (2015b) also calls his account of knowledge *ecumenical reliabilism*. While there are small differences between these accounts, they can safely be considered together since they both hold that unreliable knowledge can occur when agents successfully achieve their desired ends *through* their abilities. This claim is the focus of the rest of the paper.

Turri defines *abilism* in the following ways:

Abilism defines knowledge as true belief manifesting the agent's cognitive ability or powers (2016a, 225);

Knowledge is approximately true thin belief manifesting *cognitive ability* (2015a, 321; and 2017, 164);

Knowledge is an accurate representation produced by cognitive ability (2019).¹³

Turri's terminology of cognitive abilities "producing" or "manifesting" true beliefs serves to explain why certain unreliable processes can produce knowledge. Turri (2016b) takes the following example from Sosa (2007) to elucidate these concepts: An archer hitting a bullseye manifests her athletic ability only when her hitting the bullseye is *based on or the result of or because of* her abilities. If a gust of unexpected wind interferes with the arrow's path and causes the arrow to hit the bullseye, then the bullseye was not a result of the archer's abilities. But unlike Sosa, Turri does not require that our cognitive abilities be reliable (see § 1). This also fits with his account of achievements explained in § 1 above: Achievements involve simply attaining one's intended outcome *through* one's (un)reliable ability. In my own words, Turri holds that S knows or intellectually achieves P iff P is true, and S believing P is the result of or manifests S's (un)reliable cognitive abilities.¹⁴

One tempting argument to make against any account of knowledge that allows for the possibility of unreliable knowledge is that such accounts would implausibly allow for *lucky* knowledge. Turri's account of knowledge seems especially vulnerable to this objection since it seems that the novice archer who achieves a bullseye on her first try has beginner's luck even though she achieved the bullseye, in some sense, *through* her abilities. In response, Turri agrees that lucky knowledge is implausible but he denies that abilism allows for lucky knowledge:

The fact that someone cannot reliably produce an outcome does not entail that it's "just luck" when she does produce it. Unreliable

13 The second and third definitions indicate that Turri now holds that neither belief nor truth is a necessary condition for knowledge. I will ignore these aspects of Turri's account in this paper since they do not affect my arguments below.

14 Recall from § 1 that this account of achievement is distinctive because it does not require achievements manifest one's competence which involves manifesting one's abilities *reliably* (see Sosa 2007, 2015; Zagzebski 2009; and Greco 2010).

performers usually still have *some* ability or power to produce the relevant outcome. *Unreliability* does not equal *inability*. (2015b, 533)

While Turri does not explicate the different kinds of luck at issue here,¹⁵ the ideas are clear enough to be intuitively compelling. The novice archer who hits the bullseye through their unreliable abilities (e.g. through effort and concentration) does not succeed just by luck; while the archer who hits the bullseye because of a gust of wind does succeed by luck. Likewise, for Turri, *intellectual* achievements that issue from one's unreliable cognitive abilities are not lucky in the way that achieving a true belief through, say, guessing is lucky. Despite his poor track-record, when House correctly diagnoses a patient through his great diagnostic ability, he does so in a way that an avid fan of *House M.D.* does not when they guess the correct diagnoses. Because many unreliable processes *manifest one's ability* while lucky processes do not, Turri argues that his account of knowledge does not allow for lucky knowledge.

In essence, Turri is making the following argument:

1. Not all unreliable cognitive processes are lucky.
2. Some of the processes in (1) are non-lucky but unreliable cognitive processes that manifest one's cognitive ability.
3. Some of the processes in (2) can produce knowledge.
4. Thus, unreliable knowledge is possible.

I agree with Turri that unreliability does not equal inability and that, per [premise one](#), we should not think that all unreliable processes are just lucky processes. To deny these claims is to implausibly deny that there are nascent cognitive abilities. I also agree with Turri that, per [premise two](#), his account of knowledge does not allow for lucky knowledge. However, the key issue is whether [premise three](#) is true because if it is, then abilism is true and unreliable knowledge is possible.¹⁶

¹⁵ See Pritchard (2005) for an analysis of the different kinds of epistemic luck.

¹⁶ For various theoretical reasons, many authors would deny this premise. For example, Pritchard (2012) and Kelp (2013) argue that an ability condition on knowledge should be combined with a safety condition in order to deal with counterexamples typically leveled against ability conditions (e.g. fake barn cases). And since safety is a kind of reliability condition, Turri is presumably committed to rejecting these accounts of knowledge. Others have suggested that when the ability condition on knowledge is properly unpacked it entails an anti-luck condition precisely because it entails a reliability condition. For example, Sosa (2015), Carter (2016), and Beddor and Pavese (2020) have suggested, on different grounds, that the best version of a cognitive ability condition

To see why [premise three](#) is false it is important to first realize that the [Jessica example](#) in § 3 is one instance of someone who fits Turri's definition of abilist/unreliable knowledge but intuitively fails to have knowledge. Jessica's true belief that the blurry picture in front of her is of a basketball is the result or manifestation of her unreliable cognitive ability to recognize such images (i.e. 25% average accuracy) but she fails to have knowledge until she is told her belief is true. [Premise three](#) is false because counterexamples like this can be generalized to show that unreliable/abilist knowledge is impossible. In short, I argue that this unreliable/abilist knowledge is impossible because any agent that is in a sufficiently favorable epistemic position to have unreliable/abilist knowledge will fail to have knowledge. And as was shown in § 3, Jessica is in such a sufficiently favorable epistemic position for unreliable/abilist knowledge but she intuitively fails to have knowledge.

One might object that Jessica is not in a sufficiently favorable epistemic position to have unreliable/abilist knowledge. Firstly, an objector could argue that knowledge can be unreliably achieved only above some threshold of unreliability (e.g. above 40%). So, while Jessica is very reliable in comparison to her peers, she still only has 25% reliability and falls below this threshold for unreliable knowledge. Additionally, one could object that our intuitions about the Jessica case may be compromised by the fact that Jessica's unreliability is caused by her sub-par eyesight or malfunctioning ability to see. Indeed, what makes the House case compelling is that House's unreliability is not caused by a sub-par or malfunctioning ability (since he is the best of the best) but because of the difficulty of his job—i.e. diagnosing unusual patients. So, for these reasons one could argue that the Jessica case is not a convincing counterexample to abilism and the possibility of unreliable knowledge.

In response, I claim that additional examples can be constructed to avoid these pitfalls that nevertheless show that unreliable/abilist knowledge is impossible:

Ashley is a professional singer. While Ashley does not have perfect pitch, after many years of studying, practicing, and performing she has gained some ability to accurately identify notes played on a piano. Specifically, Ashley is able to accurately identify what single note is played by listening alone with almost 50% average accuracy.

for knowledge entails a safety condition. Presumably Turri is also committed to rejecting all of these arguments. My argument below undermines Turri's [third premise](#) directly via novel counterexamples without relying on any of these theoretical considerations.

In contrast, the average lay person is almost never able to correctly identify the right note since they have no ability to recognize which of the 12 possible notes is played. Those with perfect pitch are able to recognize which note is played with near 100% accuracy. Imagine that you are watching Ashley practice her ability over the period of half an hour. In this time, you see her correctly identify what note is played on average almost 50% of the time. Furthermore, you notice that when Ashley is wrong, she is never more than a musical half-step from the right answer (e.g. if the answer is A#, Ashley answers A; or if the answer is F, Ashley answers E).


Unlike Jessica, Ashley is much more reliable at almost 50% and, like House, does not have a sub-par or malfunctioning ability. You could say that she *nearly* has perfect pitch since her answers indicate that even when she is wrong, she is still tracking the correct pitch. But even with this great ability to identify pitches by auditory means alone, imagine that Ashley is played a Db note on a piano and correctly answers Db, but is not yet told that her answer is correct. At this point, does Ashley *know* that the note is a Db? Intuitively, Ashley does *not* know the answer is Db, and I contend the only explanation for this intuition is that despite her nascent perfect pitch ability she is still unreliable at identifying pitches. Thus, abilism is false because examples like this show that one can have a true belief that manifests one's unreliable cognitive abilities without having knowledge.

So, to reiterate, examples like this also show that unreliable knowledge is impossible since such agents are in sufficiently favorable epistemic conditions to have this kind of knowledge, but intuitively still fail to have knowledge. Furthermore, I contend that many more examples can be constructed to support the intuition that unreliable agents like Jessica and Ashley fail to have knowledge. In summary, I am making the following argument:

1. If those in sufficiently favorable epistemic positions to have unreliable/abilist knowledge fail to have knowledge, then unreliable abilist/knowledge is impossible.
2. Ashley, Jessica, etc., are in sufficiently favorable epistemic positions to have unreliable/abilist knowledge but fail to have knowledge.
3. Thus, unreliable/abilist knowledge is impossible.

In conclusion, Turri has not established that unreliable knowledge is possible and there are decisive reasons for thinking knowledge requires reliability.*

Jonathan Dixon

 0000-0003-4040-7289

Wake Forest University

dixonj@wfu.edu

References

- ALSTON, William P. 1995. "How to Think About Reliability." *Philosophical Topics* 23(1): 1–29, doi:10.5840/philtopics199523122.
- BEDDOR, Bob, and Carlotta PAVESE. 2020. "Modal Virtue Epistemology." *Philosophy and Phenomenological Research* 101(1): 61–79, doi:10.1111/phpr.12562.
- BRADFORD, Gwendolyn. 2015. *Achievement*. Oxford: Oxford University Press, doi:10.1093/acprof:oso/9780198714026.001.0001.
- CARTER, J. Adam. 2016. "Robust Virtue Epistemology as Anti-Luck Epistemology: A New Solution." *Pacific Philosophical Quarterly* 97(1): 140–155, doi:10.1111/papq.12040.
- DELLSÉN, Finnur. 2017. "Reactionary Responses to the Bad Lot Objection." *Studies in History and Philosophy of Science* 61: 32–40, doi:10.1016/j.shpsa.2017.01.005.
- . 2018. "The Heuristic Conception of Inference to the Best Explanation." *Philosophical Studies* 175(7): 1745–1766, doi:10.1007/s11098-017-0933-2.
- FRAASSEN, Bas C. van. 1989. *Laws and Symmetry*. Oxford: Oxford University Press, doi:10.1093/0198248601.001.0001.
- GOLDMAN, Alvin I. 1979. "What Is Justified Belief?" In *Justification and Knowledge. New Studies in Epistemology*, edited by George Sotiros PAPPAS, pp. 1–24. Philosophical Studies Series n. 17. Dordrecht: D. Reidel Publishing Co. Reprinted in Goldman (2012, 29–49).
- . 2008. "Reliabilism." In *The Stanford Encyclopedia of Philosophy*. Stanford, California: The Metaphysics Research Lab, Center for the Study of Language and Information, <https://plato.stanford.edu/archives/sum2008/entries/reliabilism/>. Version of April 21, 2008; reprinted in Goldman (2012, 68–94).
- . 2012. *Reliabilism and Contemporary Epistemology. Essays*. New York: Oxford University Press, doi:10.1093/acprof:oso/9780199812875.001.0001.

* I would like to thank the audiences at the 3rd Annual Philosopher Convergence and the 2019 Eastern APA for their valuable feedback on previous versions of this paper. Thanks to Heather Spradley for her insightful commentary at the Eastern APA. I would also like to thank Hilary Kornblith, Christopher Meacham, Sophie Horowitz, Cameron Gibbs, and especially Elizabeth C. Huffer for providing written comments and discussing these issues at length with me.

- GOLDMAN, Alvin I., and Bob BEDDOR. 2021. "Reliabilist Epistemology." In *The Stanford Encyclopedia of Philosophy*. Stanford, California: The Metaphysics Research Lab, Center for the Study of Language and Information, <https://plato.stanford.edu/entries/reliabilism/>. Revision, May 21, 2021, of Goldman (2008).
- GRECO, John, ed. 2008. *The Oxford Handbook of Skepticism*. Oxford Handbooks. New York: Oxford University Press, doi:10.1093/oxfordhb/9780195183214.001.0001.
- . 2010. *Achieving Knowledge. A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- HETHERINGTON, Stephen Cade. 1998. "Actually Knowing." *The Philosophical Quarterly* 48(193): 453–469, doi:10.1111/1467-9213.00114.
- . 1999. "Knowing Failibly." *The Journal of Philosophy* 96(11): 565–587, doi:10.2307/2564624.
- . 2016. *Knowledge and the Gettier Problem*. Cambridge: Cambridge University Press, doi:10.1017/CBO97811316569870.
- KELP, Christoph. 2013. "Knowledge: The Safe-Apt View." *Australasian Journal of Philosophy* 91(2): 265–278, doi:10.1080/00048402.2012.673726.
- LACKEY, Jennifer. 2007. "Why We Don't Deserve Credit for Everything We Know." *Synthese* 158(3): 345–361, doi:10.1002/9781119420828.ch13.
- . 2009. "Knowledge and Credit." *Philosophical Studies* 142(1): 27–42, doi:10.1007/s11098-008-9304-3.
- MC AULIFFE, William H. B. 2015. "How Did Abduction Get Confused with Inference to the Best Explanation?" *Transactions of the Charles Sanders Peirce Society* 51(3): 300–319, doi:10.2979/trancharpeirsoc.51.3.300.
- PRITCHARD, Duncan. 2005. *Epistemic Luck*. Oxford: Oxford University Press, doi:10.1093/019928038X.001.0001.
- . 2008. "Greco (2008) on Knowledge: Virtues, Contexts, Achievements." *The Philosophical Quarterly* 58(232): 437–447, doi:10.1111/j.1467-9213.2008.550.x.
- . 2009. "Apt Performance and Epistemic Value [on Sosa (2007)]." *Philosophical Studies* 143(3): 407–416, doi:10.1007/s11098-009-9340-7.
- . 2012. "Anti-Luck Virtue Epistemology." *The Journal of Philosophy* 109(3): 247–279, doi:10.5840/jphil201210939.
- SARTWELL, Crispin. 1991. "Knowledge is Merely True Belief." *American Philosophical Quarterly* 28(2): 157–165.
- . 1992. "Why Knowledge is Merely True Belief." *The Journal of Philosophy* 89(4): 167–180, doi:10.2307/2026639.
- SCHUPBACH, Jonah N. 2014. "Is the Bad Lot Objection Just Misguided?" *Erkenntnis* 79(1): 55–64, doi:10.1007/s10670-013-9433-8.
- SOSA, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge. Volume I*. Oxford: Oxford University Press, doi:10.1093/acprof:oso/9780199297023.001.0001.
- . 2015. *Judgment and Agency*. Oxford: Oxford University Press, doi:10.1093/acprof:oso/9780198719694.001.0001.

- TURRI, John. 2015a. "From Virtue Epistemology to Abilism: Theoretical and Empirical Developments." In *Character. New Directions from Philosophy, Psychology, and Theology*, edited by Christian B. MILLER, R. Michael FURR, Angela KNOBEL, and William FLEESON, pp. 315–331. Oxford: Oxford University Press, doi:10.1093/acprof:oso/9780190204600.001.0001.
- . 2015b. "Unreliable Knowledge." *Philosophy and Phenomenological Research* 90(3): 529–545, doi:10.31219/osf.io/6acht.
- . 2016a. "A New Paradigm for Epistemology: From Reliabilism to Abilism." *Ergo* 3(8): 189–231, doi:10.3998/ergo.12405314.0003.008.
- . 2016b. "Knowledge as Achievement, More or Less." In *Performance Epistemology. Foundations and Applications*, edited by Miguel Ángel FERNÁNDEZ VARGAS, pp. 124–135. New York: Oxford University Press, doi:10.1093/acprof:oso/9780198746942.001.0001.
- . 2017. "Epistemic Situationism and Cognitive Ability." In *Epistemic Situationism*, edited by Abrol FAIRWEATHER and Mark ALFANO, pp. 158–167. Oxford: Oxford University Press, doi:10.1093/oso/9780199688234.001.0001.
- . 2019. "Virtue Epistemology and Abilism on Knowledge." In *The Routledge Handbook of Virtue Epistemology*, edited by Heather BATTALY, pp. 309–316. Routledge Handbooks in Philosophy. Milton, Abingdon: Routledge.
- WILLIAMSON, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press, doi:10.1093/019925656X.001.0001.
- ZAGZEBSKI, Linda Trinkaus. 2009. *On Epistemology*. Belmont, California: Wadsworth.