

Skepticism, Fallibilism, and Rational Evaluation

Michael Hannon
University of Nottingham

Introduction

According to many historical philosophical figures, knowledge must be based on infallible foundations. These foundations have been characterized in different ways; e.g., as “cognitive impressions” by the ancient Stoics, as “clear and distinct perceptions” by Descartes, and as “the given” element in experience by C. I. Lewis and other twentieth-century philosophers (Reed 2012: 585). In each case, it has been assumed that these foundations are infallible in that they preclude error on the part of the knower. To have knowledge, in other words, we must have justification that *guarantees* that our belief is true.¹ This is infallibilism. It is the view that knowledge demands the highest degree of justification.

In contemporary epistemology, it is widely accepted that infallibilist theories of knowledge are “doomed to a skeptical conclusion” (Cohen 1988: 91).² We humans are fallible creatures that can rarely guarantee the truth of our beliefs; indeed, almost no belief can be rationally supported or justified in a way that removes all possible doubt. By demanding infallibility, we would prevent ourselves from knowing almost anything. This is an unwelcome skeptical result.

To avoid skepticism, many contemporary epistemologists endorse a fallibilist view of knowledge. A fallibilist believes that we can know things on the basis of justification that is less than fully conclusive. As Jim Pryor says, “a fallibilist is someone who believes that we can have knowledge on the basis of defeasible justification, justification that does not *guarantee* that our beliefs are correct” (2000: 518). This view is attractive because it allegedly avoids the skeptical consequences of infallibilist conceptions of knowledge. It is for this reason that fallibilism is typically regarded as the only serious option in epistemology. As Harvey Siegel says, “we are all fallibilists now” (1997: 164).³

Although fallibilism is almost universally accepted in epistemology, the nature of fallibilist knowledge is still poorly understood. There are at least two reasons for this. First, it is unclear how to formulate fallibilism precisely.⁴ Second, it is surprisingly difficult to describe the level of fallible justification required for

¹ I will use ‘justification’ broadly to include the related notions of ‘evidence’, ‘probability’, ‘warrant’, and ‘reliability’.

² Some have denied that infallibilism must lead to skepticism. For example, Fred Dretske (1981), Timothy Williamson (2000), Wayne Davis (2007), and Ram Neta (2011) each defend a version of infallibilism that allegedly has non-skeptical results. If these authors are right, then the step from infallibilism to skepticism is not an inevitable one. In contrast, Brown (2018) argues that infallibilists can avoid skepticism only at the cost of problematic commitments concerning evidence and evidential support. However, my focus in this chapter is not on whether infallibilism leads to skepticism, so I will set this issue aside.

³ Michael Williams agrees: “We are all fallibilists nowadays” (2001: 5). And Stewart Cohen says, “the acceptance of fallibilism in epistemology is virtually universal” (1988: 91).

⁴ See Reed (2002) and Brown (2018) for a discussion.

knowledge in a clear and non-arbitrary way.⁵ Despite this lack of precision, however, it is clear that contemporary fallibilists typically endorse the following two claims: first, knowledge is compatible with our cognitive fallibilities as inquirers; second, we typically meet the level of justification required for knowledge. In other words, fallibilists are not usually skeptics. As Stephen Hetherington writes in his encyclopaedia entry on fallibilism,

it is fallibilist epistemologists (which is to say, the majority of epistemologists) who tend not to be skeptics . . . Generally, those epistemologists see themselves as thinking about knowledge and justification in a comparatively realistic way — by recognizing the fallibilist realities of human cognitive capacities, even while accommodating those fallibilities within a theory that allows perpetually fallible people to have knowledge and justified beliefs. (Hetherington 2019)

But fallibilism does not necessarily escape skepticism. A theory might be fallibilist while still espousing standards too demanding to be regularly met. In other words, it is coherent to be both a *fallibilist* and a *skeptic*. A ‘fallibilist skeptic’ is someone who endorses the following two theses: first, the level of justification required for knowledge is “less than fully conclusive”, so we need not guarantee the truth of our beliefs to have knowledge (i.e. fallibilism); second, many of our ordinary knowledge claims are nevertheless false or unjustified (i.e. skepticism).

This might seem like a puzzling combination of ideas. In contemporary epistemology, fallibilism and skepticism are often depicted as opposing views: we embrace fallibilism to escape skepticism, and to deny fallibilism is to risk skepticism. However, fallibilism alone does not guarantee that most of our ordinary knowledge claims are true. In this paper, I will defend a version of skepticism that is compatible with fallibilism and supported by recent work in psychology. In particular, I will argue that we often cannot properly trust our ability to rationally evaluate reasons, arguments, and evidence (a fundamental knowledge-seeking faculty). We humans are just too cognitively impaired to achieve even fallible knowledge, at least for many beliefs.

High-standards Skepticism

Fallibilists often complain that skeptics wrongly impose impossibly high standards for knowledge. The skeptic claims, for instance, that knowledge requires one to be *absolutely certain* and that absolute certainty is impossible or rare (see Unger 1975).⁶ This view demands that we justify our beliefs to the highest possible degree. I will therefore call it ‘high-standards skepticism’.

⁵ BonJour (2010) makes this objection. See Hetherington (2006) and Hannon (2017) for replies.

⁶ ‘Certainty’ is ambiguous (cf. Reed 2008). A belief is *psychologically* certain when one is supremely convinced of its truth. A belief is *epistemically* certain when it has the highest possible epistemic status. Some infallibilists say that psychological certainty is necessary for knowledge, while others say that epistemic certainty is necessary.

Descartes seems to endorse a version of high-standards skepticism in the *Meditations*. He writes:

Reason now leads me to think I should hold back my assent from opinions which are not *completely certain and indubitable* just as carefully as I do from those which are patently false. So, for the purpose of rejecting all my opinions, it will be enough if I find in each of them at least *some* reason for doubt. (Descartes 1998: 59)

Scholars often interpret Descartes as endorsing the following idea: the bar for knowledge is set so high as to demand infallibility or absolute certainty. We also find this idea in Descartes' *Rules for the Direction of the Mind*, where he writes: "All knowledge is certain and evident cognition". Although Descartes was himself not a skeptic, he seems to imply that knowledge requires us to meet a very demanding standard.⁷ Likewise, the skeptic seems to presuppose an infallibilist principle like this one: if I know that *p*, then I can eliminate all grounds for doubting it. Put another way, Descartes and the skeptic seem to suggest that knowing *p* requires one to have evidence or reasons sufficient to rule out all the alternative possibilities to *p*. This is a version of 'high-standards' skepticism because it requires knowers to justify their belief to the highest possible degree.⁸

A common response to such skeptical arguments is to treat them as depending on too stringent a conception of knowledge (Reed 2012: 585). That is, we may agree with the skeptic that hardly any belief can be justified to the highest degree. After all, our cognitive faculties are too imperfect to establish the truth of a proposition with 100% certainty. But who cares? So what if nothing meets this incredibly high standard? After all, many things are probably true and it is reasonable for us to believe those things. Indeed, even granting that nothing is certain, why conclude that we have no knowledge? While the skeptic believes that we cannot know what we cannot confirm with 100% certainty, fallibilists have a more moderate view. Absolutely certainty is not required for knowledge. As Baron Reed puts it, "our faculties are still very good . . . they allow us to achieve a more modest sort of cognitive success. Fallibilism, then, takes that modest success to be knowledge" (2012: 585). Duncan Pritchard expresses a similar thought:

it doesn't seem at all credible that the bar for knowledge should be set so high as to demand infallibility or absolute certainty (or, for that matter, indubitability). On the contrary, our everyday conception of knowledge seems to leave us perfectly happy with the idea that knowledge can be fallible and not absolutely certain (and thus to a degree dubitable) while being *bona fide* knowledge nonetheless. (2019: 36)

I think this reaction is exactly right. If the skeptic sets the bar for knowledge too high, then we should reject that standard (see Hannon 2019a). The epistemic contextualists have adopted this line of

⁷ I do not actually agree with this interpretation of Descartes, but I set this issue aside. See Pasnau (2017) for a discussion of this point.

⁸ Stroud (1984) seems to interpret Descartes as endorsing a 'high standards' view.

reasoning.⁹ Their solution to skepticism involves two basic elements: first, in ordinary contexts we often meet the reasonable (fallible) epistemic standard for knowledge; second, in skeptical contexts the standards to know are much higher.¹⁰ As DeRose puts it,

In some contexts, “S knows that p” requires for its truth that S have a true belief that p and also be in *a very strong epistemic position* with respect to p, while in other contexts, an assertion of the very same sentence may require for its truth, in addition to S’s having a true belief that p, only that S meet some lower epistemic standard. (2009: 3)

A core feature of contextualism is that we need not meet the skeptic’s very high epistemic standard to have knowledge in daily life. This is fallibilism. The contextualist is a fallibilist who rejects (or at least confines) the infallible standard assumed by the skeptic.

While this response to skepticism is *prima facie* plausible, it only gains purchase if we satisfy some reasonable epistemic standard that ordinarily suffices for knowledge. After all, the contextualist line is precisely that in everyday contexts our knowledge claims are true because the standards are not too demanding. However, the most challenging skeptical arguments do not simply claim that we fail to meet some extraordinary standard for knowledge. Rather, they claim that we do not meet even ordinary (fallible) standards. I will consider this view in the next section.

Hard-hitting Skepticism

The contextualist portrays the dispute between the skeptic and the non-skeptic as a difference between using stricter standards and more lax ones. On this interpretation, the skeptic may be accused of merely imposing abnormal requirements on our familiar concept of knowledge. As Hans-Johann Glock says, “the skeptic . . . is like someone who claims that there are no physicians in London, since by ‘physician’ he understands someone who can cure any disease within twenty minutes” (2010: 100).

But this misportrays the skeptic’s view. As Bryan Frances writes, “the skeptic isn’t complaining that our knowledge doesn’t satisfy some super-duper high-octane condition that only a philosopher could love” (2008: 243). Rather, the skeptic is arguing that it is much more difficult than we realized for a belief to qualify as knowledge *even by ordinary standards*.¹¹ In presenting her argument, the skeptic raises doubts about whether we actually satisfy the very same epistemic standards that we have always thought we satisfied, not some unattainably high standard. Thus, any solution that characterizes the skeptic as “raising the standards” or smuggling in abnormal requirements would mischaracterize the view.

⁹ See Cohen (1988), DeRose (1995), and Lewis (1996) for early statements of this view.

¹⁰ I have not indicated semantic ascent by putting quotation marks around ‘know’ and ‘knowledge’. I state my discussion in the object language for ease of exposition.

¹¹ Feldman (1999; 2001), Klein (2000), Kornblith (2000), Bach (2005), Hazlett (2014), and Pritchard (2019) raise versions of this objection.

What about the skepticism of Descartes' *Meditations*? As previously mentioned, several philosophers have interpreted the Cartesian skeptic as imposing high standards or presupposing infallibilism.¹²

But this connection between infallibilism and skepticism may be a red herring. As John Greco (2008: 116), Allan Hazlett (2014: 90), and others have pointed out, the Cartesian skeptical argument does not essentially depend on infallibilism. The Cartesian skeptical argument runs as follows:

1. I know that I have hands only if I can know that I'm not deceived by a demon (about whether I have hands).
2. I can't know that I'm not deceived by a demon (about whether I have hands).
3. Therefore, I don't know that I have hands.

None of the motivations for the first or second premise presuppose infallibilism. The Cartesian skeptic is not claiming that you cannot be *certain* that you're not deceived by a demon, and therefore know very little. That *would* presuppose infallibilism. The idea behind the Cartesian skeptical argument is, as Hazlett says, "that you've got no way of knowing whether you're deceived, and therefore know very little" (2014: 90).

Nevertheless, many people are unwilling to grant the skeptical premise that we don't (or can't) know that we're not deceived. Instead of allowing the skeptic to use claims like "I can't know that I'm not being deceived" as premises in her reasoning, we should instead expect the skeptic to *convince* us that we can't know this.¹³ As Greco writes,

Is that premise [that I cannot know that I am not a brain in a vat] initially (or pretheoretically) plausible? It seems to me that it is not. In fact, it seems to me that it is initially obvious that I *do* know that I am not a brain in a vat. (2008: 111)

Whatever we think of this response to the skeptic, there are two important lessons to draw out from the current discussion. First, to portray the skeptic as demanding that we justify our beliefs to the highest degree is to mischaracterize their view. The really worrying form of skepticism is not 'high standards' skepticism but rather what I call 'hard-hitting skepticism'. The hard-hitting skeptic says that our beliefs fail to qualify as knowledge *by ordinary standards*.

The second lesson is this: a common source of resistance to skepticism is the implausibility of far-fetched scenarios involving evil demons, brains in vats, and so forth. As Frances (2008) observes, this is one thing that commonly bothers undergraduates in philosophy. They object:

why on earth do some philosophers take the BIV [brain in a vat] hypothesis to pose any threat at all to our beliefs, given that those very same philosophers think that there is no real chance that the

¹² Unger (1975) and BonJour (2010) also seem to defend versions of high-standards skepticism.

¹³ Pryor (2000) attempts to formulate a version of skepticism that does not require this controversial premise.

BIV hypothesis is true? Sure, the BIV hypothesis is formally inconsistent with my belief that I have hands, so if the former is true, then my belief is false. But so what? Why should that bare inconsistency matter so much? The students would understand the fuss over the BIV hypothesis if there were some decent reason to think that the BIV hypothesis was really true. (Frances 2008: 225)

A more threatening and significant type of skepticism, then, would meet two conditions. First, it would not characterize the skeptic as demanding infallibility or the highest possible degree of justification. Second, it would not be based on doubts that are “purely philosophical” or “merely academic threats” (Frances 2008: 228). Rather, it would involve what C. S. Peirce calls a “real” doubt. (We don’t really doubt the existence of the external world.) In the next section, I will outline a sceptical argument that meets these two criteria and thereby generates a “real” skeptical threat.¹⁴

Skepticism and Rational Evaluation

This section will outline a type of contingent real-world skepticism that has not received much attention. Unlike ‘high-standards’ skepticism, this view does not demand infallibility or extraordinarily high justification. Unlike Cartesian skepticism, it does not rely on remote possibilities, such as the possibility that one is a brain in a vat or radically deceived by an evil demon.¹⁵ Rather, I will present an empirically informed, scientifically respectable skeptical hypothesis that targets many of our most cherished beliefs. More specifically, I will outline a type of skepticism that targets *many beliefs that are the product of reasoning*.

I am using ‘reasoning’ quite broadly to refer to cases in which we have evidence and draw conclusions on the basis of rationally evaluating the evidence. Put differently, I am thinking of cases in which we *exercise our rational capacities to acquire knowledge*. This contrasts with more basic and immediate cases of perceptual knowledge, which require little to no reflection or reasoning.

Consider the following example.¹⁶ Irena suspects that the death penalty is not an effective tool to reduce the murder rate, but she doesn’t have much evidence either way. To become more informed, she decides to look for evidence about the deterrent effects of capital punishment. After searching several reputable websites and reading multiple academic articles, Irena concludes that the evidence strongly indicates that the death penalty does not reduce the murder rate. Irena is intelligent, articulate, and is now able to

¹⁴ Pyrrhonian skepticism in no way depends on the problematic assumptions outlined above. The pyrrhonist did not claim the standards for knowledge were incredibly high, nor did they appeal to far-fetched skeptical scenarios.

¹⁵ Here I follow Frances’s (2005) strategy to outline what he calls a “live” skeptical hypothesis. Frances says that a hypothesis is live when it satisfies five conditions: (i) it has been through a significant evaluation in the community by experts over many years; (ii) it is judged actually true or as likely to be true as any relevant possibility by a significant number of experts; (iii) the judgment of those experts has been reached in an epistemically responsible way; (iv) those experts consider there to be several independent sources of evidence for the hypothesis; (v) many of those experts consider the hypothesis to be a live possibility (see Frances 2005: 18-9).

¹⁶ This example is adapted from Kornblith (1999: 181-2).

present reasons for her belief. Moreover, she now believes that the reasons for which she holds her belief are the reasons she is now able to present.

This might seem like a reliable way to form beliefs, but we have abundant evidence that this process of reasoning is likely beset by a host of cognitive biases and reasoning errors. For example, the general human tendency to *rationalize* is incredibly common.¹⁷ This occurs even when people offer what in fact turn out to be sound arguments. Though they may have successfully hit on a good argument, they are still rationalizing because they would have held the same view with or without good reasons, i.e. on the basis of non-rational considerations.

Suppose that Irena has views about the morality of the death penalty that she holds on grounds independently of her views about its deterrent effects. These moral commitments will likely drive her views about the deterrent effects of the death penalty. For instance, Irena will inflate the quality of the studies that present evidence in favour of her moral view and she will be far more critical of (and more likely to downplay the significance of) studies that provide evidence that cut against her moral view. In other words, her views about the death penalty's effectiveness will not be the result of her understanding the relevant data. Quite the opposite: her understanding of the relevant data will be the product of her moral beliefs. Moreover, Irena will be completely unaware of this fact. She will sincerely believe that her reasons for belief are different from her actual reasons.

Irena is not a unique case. A vast amount of work in cognitive psychology indicates that we all frequently interpret and filter evidence in ways that fit with our antecedent worldview.¹⁸ For example, we selectively expose ourselves to evidence that confirms our pre-existing beliefs and avoid information that conflicts with them. This is known as *selective exposure* (Nickerson 1998). We also tend to uncritically accept (and better remember) evidence that is favourable to our view, whereas we are far more critical (and forgetful) of counterevidence. We “routinely rationalize the facts, figures, and arguments that [we] cannot effortlessly discount, depreciate, denigrate, or deny” (Lodge and Taber 2013: 59). This general human tendency to accept confirming evidence without much scrutiny and subject disconfirming evidence to highly critical evaluation is known as *biased assimilation* or *confirmation bias* (Lord et al. 1979). When this occurs, two people can look at the exact same body of evidence and yet walk away with radically different conclusions about what the evidence shows, thereby drawing undue support for their initial positions.¹⁹

In general, our belief-forming processes are often corrupted by an array of normal human cognitive and affective tendencies. Following Aaron Ancell (2019: 411), I will call these “sources of unreason”. Sources of unreason include: our prejudices and biases (both implicit and explicit); the tendency of self-interest and

¹⁷ For a defense of this view, see Kornblith (1999), Tavis and Aronson (2007), Haidt (2012), Sperber and Mercier (2017).

¹⁸ See Gilovich (1993) for an overview.

¹⁹ These biases are sometimes collectively referred to as “motivated reasoning”, which is “the tendency to seek out, interpret, evaluate, and weigh evidence and arguments in ways that are systematically biased toward conclusions that we ‘want’ to reach for reasons independent of their truth or warrant” (Ancell 2019: 418). See also Kunda (1990); Ditto, Pizarro, and Tannenbaum (2009); Lodge and Taber (2013).

group-interest to distort our judgments; stubbornness and dogmatism; bad reasoning; the desire to reduce cognitive dissonance; and psychological comfort. Our reasoning is especially prone to error or bias when it comes to beliefs that *matter to us*; e.g., our moral, political, and personal beliefs that are partly constitutive of our identity (Haidt 2012). These ‘identity-constitutive beliefs’ include any belief that reflects one’s conception of “who they are, of what sort of people they are, and how they relate to others” (Hogg and Abrams 1988: 2). Moreover, these corrupting motivations are not transparent to us. We believe that the reasons we present are the reasons for which we hold our beliefs, but we are often wrong about this. We think we are motivated by the desire for truth, but our other motives are concealed from our view (Wilson 2002).

This raises a general doubt about human reasoning. As Ancell writes, “Our reasoning capacities are beset by an array of built-in cognitive and affective biases that make it very difficult—often practically impossible—for us to think clearly and objectively about issues that affect our interests and arouse our passions” (2019: 418). This would not be so worrying if we were able to introspectively discern our own biases and cognitive shortcomings. However, we are typically introspectively blind to them (Pronin, Lin, and Ross 2002).²⁰ While it might *seem* to us that we are being impartial and unbiased, we are often “twisting the argument and evidence to make them fit the conclusions we ‘want’ to reach” (Ancell 2019: 419). We suffer from what psychologists call *the illusion of objectivity* (Kunda 1990).

In addition to the general human tendencies to rationalize, assimilate information in biased ways, and selectively expose ourselves to favourable information, there are many other biases and cognitive errors to which we humans are prone. For instance, hundreds of studies have confirmed the existence of a *hindsight bias*, which occurs when people who know the outcome of an event judge it to be more probable than people who are ignorant of the outcome (Roese and Vohs 2002). This is known to affect judgments about topics as diverse as terrorist attacks, medical diagnoses, and accounting decisions. In addition, we all have *implicit biases* where we unconsciously and automatically associate concepts with one another. Such biases are especially pernicious when we associate certain traits (e.g., dangerous) with members of particular social groups.²¹ Further, there is considerable evidence that many of our most cherished beliefs are shaped by *irrelevant influences* that do not bear on the truth of what we believe (Vavova 2018). For instance, the fact that you were raised in one community rather than another seems epistemically irrelevant to what you ought to believe about God, morality, or politics. But factors like upbringing are

²⁰ As Ballantyne (2019: 131) writes: “A central idea in psychology is that most biases are not reliably detected by introspection (Nisbett and Wilson 1977; Wilson and Brekke 1994; Kahneman 2003). We typically can’t figure out whether we are biased by merely gazing into our minds. Biases normally ‘leave no trace’ in consciousness. As Timothy Wilson and Nancy Brekke quip, ‘Human judgments—even very bad ones—do not smell’ (1994: 121). From the inside, biased attitudes seem just like unbiased ones.”

²¹ Jennifer Saul (2013) has argued, compellingly, that implicit biases present a skeptical challenge to the ordinary ways that we assess reasons, arguments, and evidence. While my argument is similar to Saul’s, it has wider scope. Saul presents a challenge to beliefs that may be influenced by implicit biases, whereas I raise doubts about our rational capacities more generally. That said, we both agree that (a) the rational evaluation of evidence and arguments is often corrupted by problematic psychological tendencies, (b) this influence operates below the level of consciousness, and (c) this provides a reason for skepticism about such beliefs.

known to guide our convictions on these and other, less charged, topics.²² (While this is not itself a psychological bias, it does further illustrate the vast extent to which our beliefs are not the product of a reliable belief-forming process.)

These are just a handful of our many cognitive shortcomings. It would be impossible to provide a comprehensive overview of all the epistemically irrational biases in human reasoning that should be of interest to epistemologists. However, these few examples illustrate a general worry about human reasoning. We fail on a variety of cognitive dimensions and, as a result, we have a distorted sense of the plausibility of our own beliefs. In other words, our epistemic situation is likely much worse than we think. Moreover, these cognitive shortcomings are not rare but rather are the norm. There is plenty of evidence that rationalization is extremely widespread, especially when it comes to beliefs that matter to us; e.g., our political beliefs, moral beliefs, religious beliefs, or any other issue that affects our interests or stirs our emotions (Haidt 2012).

Astute observers of human nature anticipated these psychological findings. Francis Bacon wrote, “The human understanding when it has once adopted an opinion (either as being the received opinion or as being agreeable in itself) draws all things else to support and agree with it” (1620: XLVI). Bertrand Russell said, “It is a law of our being that, whenever it is in any way possible, we adopt beliefs as will preserve our self-respect” (2004 [1928]: 51). British essayist William Hazlitt noted, “The narrowness of the heart warps the understanding, and makes us weigh objects in the scales of our self-love, instead of those of truth and justice” (1824: 34). And John Locke remarked on the deplorable state of the human mind in his posthumously published work, *Of the Conduct of the Understanding*:

There are several weaknesses and defects in the understanding, either from the natural temper of the mind or ill habits taken up, which hinder it in its progress to knowledge. Of these there are as many possibly to be found, if the mind were thoroughly studied, as there are diseases of the body, each whereof clogs and disables the understanding to some degree, and therefore deserves to be looked at and cured. (1996 [1706], 187; §12)

Although Locke was keenly aware of our intellectual imperfections, he was optimistic about our ability to overcome them. He prescribes that we impartially self-examine our own beliefs to root out “the prejudices imbibed from education, party, reverence, fashion, interest” (1996 [1706] 184 §10). In other words, he thinks that rigorous self-study will help to expose our biases. But this seems overly optimistic. Self-judgment is often clouded by rationalization and prejudice, so attempts to watch for signs of bad reasoning will go undetected and leave us thinking that we’ve actually tuned them out. As Nathan Ballantyne writes, “The feeling that we’ve done our best to be unbiased will encourage us to think we are unbiased, but that feeling should not be trusted” (2019: 131-2).

²² Sometimes factors like upbringing will affect our convictions in perfectly rational ways, since they expose us to different experiences and thus different bodies of evidence. However, there are also many cases in which factors like upbringing are epistemically irrelevant.

A Skeptical Argument

These biases and psychological shortcomings pose a significant epistemic threat to our beliefs. When our beliefs are shaped by sources of unreason, this casts doubt on the epistemic quality of those beliefs. This implies a form of skepticism. Consider the following skeptical argument, which I will call *skepticism from unreason*:

1. Many of our beliefs are shaped by sources of unreason.
2. When a belief is shaped by sources of unreason, we should significantly reduce our confidence in it.
3. Therefore, we should significantly reduce our confidence in many of our beliefs.

We can also derive a form of skepticism about knowledge:

4. We have a reason to significantly reduce our confidence in many of our beliefs.
5. If we have a reason to significantly reduce our confidence in a belief, then that belief does not amount to knowledge.
6. Therefore, many of our beliefs do not constitute knowledge.

This does not yet give us a very worrying skeptical result. While it does show that many of our beliefs do not constitute knowledge, it leaves open the possibility that we can *tell* when the process of rational evaluation generates justified (or known) beliefs and when it goes awry. To derive a worrying skeptical conclusion, we must not only show that some of our beliefs are shaped by sources of unreason. (We knew *that* already.) Rather, it must also be true that we are unable to *distinguish* the beliefs that are shaped by sources of unreason from beliefs that have not been corrupted by our cognitive biases. (If we were able to determine which beliefs are the result of twisting the evidence from those that are the product of even-handed rational evaluation, the skeptical result would be much less threatening.) As mentioned, however, the biasing processes that lead us to rationalize, to selectively filter our evidence, and so forth, typically take place behind the scenes.²³ Thus, we often cannot cancel out the threat posed by our psychological shortcomings, for we lack some epistemic feature (e.g., appropriate evidence, reliability, or what not) that would otherwise cancel out this threat. As Ansell writes,

we lack reliable means of detecting, avoiding, and correcting for these biases, [so] we cannot plausibly expect that otherwise reasonable people will never be led astray by them. Indeed, we must expect the opposite; the . . . reasoning of sincere and conscientious people will often be warped by their biases, interests, partisan loyalties, and so on. Because such warped views are liable

²³ Externalists (e.g., reliabilists) might take this as a point in their favour, since they will claim that we needn't be *aware* of the reliability of our belief-forming processes. However, externalists are not clearly at an advantage here because the relevant point is that our beliefs are often the result of an *unreliable process*—not that we lack introspective awareness of when we are rationalizing.

to be unreasonable, it follows that sincere and conscientious people will often hold unreasonable views. (2019: 411)

All this should lead us to doubt that our reasoning faculties are reliable routes to knowledge. Much empirical research over the past fifty years reveals the disturbingly expansive range to which fallibility enters our cognitive lives.²⁴ This challenges the rational standing of our beliefs.

One might try to safeguard much of our knowledge from these empirical findings by arguing as follows: these findings do not undermine the *justification for our beliefs* but rather they show that we are often more confident than we actually should be, and thus we have *irrational credences*. However (the objection continues), it is not obvious how irrational credences translate to the justification of beliefs, which are ultimately relevant for knowledge.²⁵

To illustrate, let's assume a simple minded-threshold view where to have a belief is to have a credence above a certain threshold. Given that, a belief should count as justified when it is rational to have a credence above the threshold. Now it could easily be that it is rational to have a credence above the threshold even when it is not rational to have the exact credence one has. Thus, even if we grant that our credences are rarely rational, it may still be that our beliefs are justified most of the time. The only cases where a bias affects the justificatory status of a belief are cases where the respective bias leads us to cross the threshold. For instance, let's say the threshold is 0.8 and my rational credence is 0.75, but due to a bias, my actual credence is 0.85. Here I end up with an unjustified belief, given the threshold view. The conditions described, however, seem relatively specific, which may lead one to conclude that they are rarely instantiated—or so the objection goes.

This objection highlights an important difference between externalist vs. internalist conceptions of justification. According to the above objection, an agent may have an irrational credence that results from an inability to properly evaluate the strength of their evidence; yet the agent may nonetheless achieve a level of justification that suffices for their belief to be knowledge. However, many epistemologists will argue that these various biases make it impossible for one to determine how strong one's justification actually is, and thus one cannot tell whether their justification is good enough for knowledge. For a hardline externalist, what matters might be whether our actual justification is good enough for knowledge. But those with internalist leanings, like myself, will claim that our inability to tell how good our justification actually is provides us with a relevant defeater for knowledge. After all, biases like selective exposure and biased assimilation lead us to draw undue support for our views by filtering and processing the relevant evidence in epistemically problematic ways. As a result, the justification for our beliefs is often the result of twisting the evidence and arguments to make them fit with conclusions we want to reach. Thus, our

²⁴ This may reflect the malaise of our allegedly 'post-truth' era. It is not so much the idea of truth or the existence of various truths that has come under attack, but rather the notion that there can be any such thing as objective inquiry into it (Blackburn 2019).

²⁵ Thanks to Alexander Dinges for raising this objection.

evidence is likely worse than we think. Moreover, we cannot tell whether it is good enough to qualify as knowledge. This provides us with a reason to doubt whether our beliefs do qualify as knowledge, in addition to doubting that our credences are rational.

While I have argued that our beliefs often *are* shaped by sources of unreason, my skeptical argument does not actually require our beliefs to *have been* shaped by any such source. All that is required is that the relevant belief is one that reasonably could have been influenced by a source of unreason *and we cannot tell whether this has occurred in the relevant case*. When this happens, we have a *defeater* even for those beliefs that are, in fact, based on good reasons or evidence. To illustrate, recall my example of Irena and the death penalty. In this case, I admitted the possibility that she in fact hit upon good arguments for her resulting belief. Still, the worry is that humans have a general tendency to rationalize and Irena lacks the ability to tell whether she holds her belief about the death penalty's deterrent effects on the basis of good reasons or instead due to her views about the morality of the death penalty.²⁶

These pessimistic facts about reasoning and rationalization go some distance toward making sense of “fallibilist skepticism”. This view is *skeptical* because it provides a real, live hypothesis that targets reasoning as a source of knowledge. It is also *fallibilist* because this new skeptical hypothesis does not merely deny *certainty* or ‘high standards’ knowledge. Rather, it generates a reason to deny the entirely modest amount of epistemic warrant we ordinarily expect to know something. Even though many of our beliefs may be true, the justification we have for these beliefs is much less than any of us have supposed in our anti-skeptical moments. In short, the epistemic quality of our position is much worse than we thought.

The Scope and Force of Skepticism

This view has a narrower scope than some traditional forms of skepticism. The scope of skepticism is determined by the set of proposition it targets, where these propositions are said to be unknowable, unjustified, or those about which we should suspend judgment. Some skeptics target all claims about the external world. Others target our knowledge of the mental lives of others. Still others target only propositions about the future, the past, or religious matters. The type of skepticism I am outlining, by contrast, targets the propositional contents of many beliefs yielded by rational evaluation.²⁷ These beliefs

²⁶ One might take this as a reason in favor of externalist. For what it's worth, I don't think we should adopt philosophical positions simply because we don't like skeptical conclusions. (That said, in Hannon 2019b I argue that we can reject skepticism because it is impractical in a particular way.) In any case, adopting a version of externalism that allows us to *call* these beliefs ‘justified’ or ‘knowledge’ would still leave untouched my real concern, namely, that we care about being able to tell how good our epistemic position actually is. I find the externalist retreat of little practical value in responding to this concern, even though it may allow us to continue calling certain beliefs ‘justified’ or ‘knowledge’.

²⁷ We can also put this point in terms of distrusting a putative *source* of knowledge. While the traditional skeptic targets beliefs yielded by sense perception, my skeptic targets those produced by reasoning or rational evaluation.

are targeted on the grounds that they are likely corrupted by error and bias, especially when these beliefs are partly constitutive of our identity.²⁸

Not all beliefs that are the product of reasoning will be subject to the concerns I have raised. Suppose, for example, that I come to believe that we have enough milk for the next week on the grounds that we have four litres and we tend to go through about a litre every two days. A belief like this is not (obviously) likely to be shaped by “sources of unreason”.²⁹ Many of our beliefs formed through reasoning will be of this form. Still, the literature in psychology makes clear that the extent to which our beliefs *are* shaped by sources of unreason is indeed quite broad. In addition to pretty much all of our moral, political, philosophical, and religious beliefs, it will also concern other “cherished beliefs” (e.g., related to family members, sports teams, professional employment, etc.), beliefs that are influenced by our implicit biases and prejudices (see Saul 2013 for the extent to which this is troubling), beliefs influenced by ‘ingroup’ and ‘outgroup’ bias (the research on ‘minimal group paradigms’ reveals that even arbitrary distinctions between groups, such as preferences for certain paintings, beliefs about whether hotdogs are sandwiches, or the color of their shirts can trigger a tendency to favor one's own group), or other issues that affect our interests, stir our emotions, or challenge our worldview.

Although this version of skepticism has narrower scope than traditional (i.e. radical) skepticism, I believe it packs more punch. Instead of relying on the far-fetched possibility that an evil demon is controlling our minds (or that we are brains in vats, etc.), we have the much more live (and empirically supported) hypothesis that our cognitive biases are leading us astray. It is not just the *possibility* of cognitive error or bias that is raised; the research in psychology and cognitive science suggests that it is *very likely* that we are doing these things quite often.³⁰ As a result, this type of skepticism is more likely to engender doubt and to inspire behaviour. This contrasts with the fairly ineffective change in belief and behaviour brought about by traditional forms of skepticism. As Hume observed, the activities of ordinary life were sufficient to dispel traditional skeptical doubts:

Most fortunately it happens, that since reason is incapable of dispelling these clouds, nature herself suffices for that purpose, and cures me of this philosophical melancholy and delirium . . . I dine, I play a game of back-gammon, I converse, and am merry with my friends; and when after three or four hour's amusement, I wou'd return to these speculations, they appear so cold, and strain'd, and ridiculous, that I cannot find in my heart to enter into them any farther. (2003: bk.1, pt.4, sec.7)

²⁸ While I have focused on beliefs that are the product of reasoning, the concerns I have raised may also target some perceptual beliefs, beliefs based on memory, and the acceptance of testimony. In the case of perceptual beliefs, it is well known that our ‘group identity’ can shape our perceptions of reality. This explains why sports fans who cheer for different teams will perceive games differently (see Hastorf and Cantril 1954). In the case of memory, we suffer from ‘rosy retrospection’ (Mitchell et al. 1997) and hindsight bias (Roese and Vohs 2002). And in the case of beliefs based on testimony, the literature on epistemic injustice provides compelling evidence that we may routinely give too little (or too much) credibility to certain sources of testimony (see Fricker 2007).

²⁹ Thanks to Martin Smith for this point and example.

³⁰ Saul (2013) makes a similar point.

Similarly, C. S. Peirce wrote, “Let us not pretend to doubt in philosophy what we do not doubt in our hearts” (1893: CP 5.265). The most forceful and interesting skeptical challenges provide us with doubts that are genuinely compelling — i.e. they present us with good reasons to think we are very likely making errors and they challenge our ability to inquire responsibly.³¹ Skepticism about rational evaluation does precisely this.

Final Thoughts

Is this type of skepticism really all that troubling? You might think the doubt cast on our beliefs is fairly localized. Saul considers this objection in her article on skepticism and implicit bias. She writes,

It seems, at first, to be like the sort of doubt we experience when we discover how poor we are at probabilistic reasoning. We have extremely good reason to think we’re making errors when we make judgments of likelihood. But this sort of doubt doesn’t trouble us all that much because we know exactly when we should worry and what we should do about it: if we find ourselves estimating likelihood, we should mistrust our instincts and either follow mechanical procedures we’ve learned or consult an expert (if not in person, then on the internet). This kind of worry is one that everyone can accept without feeling drawn into anything like skepticism. (Saul 2012: 250)

It may seem as though skepticism about rational evaluation is like this. One might argue, for example, that we should not worry about this type of skepticism because it only concerns a small set of our total beliefs. Moreover, one might argue that we are *aware* of these various biases and cognitive shortcomings (after all, they are well documented in psychology), so we can guard against them.

There are at least two problems with this reply. First, we should *very often* be worried about these biases influencing our judgments. As the literature in psychology makes clear, these biases tend to be triggered whenever an issue affects our interests or stirs our emotions. Thus, we should be worried about the rationality of our beliefs in a variety of situations. Second, even though we are aware of the existence of these biases, it is unlikely that we can overcome them. Although a variety of debiasing strategies have been proposed (see Larrick 2004; Lilienfeld et al. 2009; Jolls and Sunstein 2006), Kristoffer Ahlstrom-Vij (2013) identifies two obstacles with the attempt to debias ourselves. First, we are not motivated to engage in debiasing because we do not view ourselves as biased (which is very common). Second, people who are

³¹ Frances’s (2005) work on ‘live sceptical scenarios’ provides an instructive comparison. However, the reasons he offers for his brand of skepticism strike me as less compelling. As Saul (2013: 254) observes, “The hypotheses in question are things like eliminativism about belief and error theory about colour. And the reasons for thinking that they are still live is that some sensible people who know a great deal endorse (or might endorse) these theories on the grounds of compelling scientific or philosophical reasons.” But this falls short of the general doubts about human reasoning that I have outlined. My claim is that *we all have very good reason to believe* that we are frequently making errors that have their root in motivated reasoning. This is stronger than Frances’s claim that a hypothesis is “live”, by which he means (roughly) that sensible and knowledgeable people might endorse it on the basis of good reasons. I borrow this point from Saul (2013), who makes a similar observation in the context of her own skeptical argument that is based on implicit biases.

persuaded to engage in debiasing efforts run the risk of both overcorrection and undercorrection (Wilson 2002).

Some may prefer to interpret the evidence from psychology in a less pessimistic way. According to an alternative view, many of our so-called biases are not really irrational at all. Rather, they reflect “bounded rationality” (e.g., Gigerenzer et al. 2001) or “instrumental rationality” (e.g., Kolodny and Brunero 2018). We act in ways that are *boundedly* rational when we employ cognitive shortcuts and rules that give rise to biases. As fallible beings with limited time and cognitive resources, it is rational for us to use heuristics and shortcuts in order to make the best decisions we can, given our limitations, even though this type of rationality gives rise to biases and errors in judgment.³² Additionally, many of our seemingly irrational beliefs may be the result of *instrumental* rationality, since these beliefs play an important role in achieving our goals. For example, it may be instrumentally rational for an individual to dogmatically hold on to certain political beliefs in the face of counterevidence because changing one’s mind is a psychologically difficult process that could potentially alienate the believer from their community and sense of self.

This line of reasoning is perfectly sensible, but it does little to blunt the force of the skeptic’s argument. The skeptic may simply argue that many of our boundedly rational beliefs do not amount to knowledge. Bounded rationality is achieved when we rely on fast and frugal heuristics to help us make sense of the world. In general, this is a good thing because the world is complex and we don’t always have the time to make a well-thought-out rational choice about a decision. Nevertheless, the skeptic will insist that she is targeting beliefs that are not sufficiently epistemically justified, even though such beliefs may be the result of a boundedly rational process. A similar point can be made about instrumentally rational beliefs. While it may be instrumentally rational to dogmatically hold on to certain beliefs, it is not *epistemically* rational to do so. The psychological costs associated with giving up such beliefs may be brutal, but this does not make the belief rational in an epistemic sense. Thus, this line of response does nothing to escape the skeptic’s clutches.

A final worry about my argument is that it is self-defeating. After all, haven’t I attempted to convince you to endorse my skeptical conclusion through rational argument? If so, then shouldn’t we doubt whether we know this conclusion?

As a skeptic, I am willing to admit that we do not know the conclusion of my argument. It is likely that philosophical argumentation is subject to various epistemic vices and cognitive biases, just as political reasoning, moral thinking, and religious belief are subject to these biases. However, to say that we do not know my conclusion is not to say that we *do* know that it is false. Rather, we are simply left uncertain (or at

³² For this objection, see Nick Hughes’s forthcoming paper, “Evidence and Bias”.

least lacking knowledge) as to whether the beliefs that are the product of rational evaluation are known or justified. This is still a skeptical conclusion.³³

Moreover, this claim is compatible with the idea that we have *some* epistemic justification to believe the conclusion of my argument. I need not argue that beliefs entirely lack justification when they are the product of rational evaluation. That is an especially strong form of skepticism that would run into the self-defeat problem, for the following reason: If rational evaluation is unjustified or unreliable, then the rational evaluation that ‘rational evaluation is unjustified or unreliable’ would itself be unjustified. But if it were unjustified, then we have good reason not to believe it or trust it. Still, we may throw the epistemic status of these beliefs into doubt without undermining their justification entirely. Thus, we may have some epistemic justification for believing that my conclusion is true, but without knowing that it is true. Perhaps we should simply suspend judgment on the matter.

Another way to avoid the self-defeat worry, which I do not pursue, is to suggest that my skepticism about rational evaluation does not apply to itself.³⁴ This would allow me to escape the self-defeat worry by claiming that my rational evaluation about the unreliability of rational evaluation is not itself unjustified. However, going this route would leave me with the challenge of convincing you that you should take my rational evaluation to be justified. In other words, my skeptical argument would provide a defeater for most beliefs that are the product of rational evaluation, but I would provide a defeater for why this argument applies to the rational evaluation of my argument. This may not be an insurmountable task, for I do not claim that rational evaluation is always or universally an unreliable process. Nevertheless, I have argued that we should *very often* be worried about these biases influencing our judgments; that we do not know exactly when we should worry; and that attempting to overcome these biases is often unhelpful.³⁵

In summary, I have argued that one of the main lessons from the literature on human psychology is that we should not trust ourselves as inquirers. Many beliefs that we take to be the product of the careful exercise of reason are likely biased and wrong. Put differently, an ordinary but fundamental belief-forming method is much less reliable than we thought. This gives rise to a form of skepticism about our rational capacities. This skeptical argument challenges what Chris Hookway calls “our cognitive instruments”, since

³³ Mark Walker suggests another version of the self-undermining objection. It runs as follows: “If we are systemically subject to sources of unreason, doesn’t this worry infect cognitive psychology itself? So, if my argument is correct, then we should reduce our confidence after reviewing the cognitive psychology literature. But now it seems I have less to worry about from cognitive psychology, since my confidence in their results is reduced.” In response to this worry, I am willing to admit that we should reduce our confidence in what we take the evidence about cognitive biases to show. However, this does not mean we should doubt whether such biases are operative. Rather, it suggests that the confidence with which we hold our beliefs about cognitive bias may not accurately reflect the evidence. In other words, our situation might be cognitively better than the evidence about bias suggests, *but it might also be worse*. We can’t tell which way our biases are pushing us. Are we weighing the evidence about bias too heavily or not heavily enough? Thus, we may have a reason to doubt that our confidence about what the evidence shows is adequately reflecting the evidence. But this doesn’t permit us to reduce our confidence in the existence or extent of cognitive bias.

³⁴ Adam Elga (2010: 179-82) argues that self-defeat objections can be avoided because, in general, methods can be exempted from self-application. I will not outline the details of his argument here, but the rough idea is that self-exemption is not ad hoc or arbitrary because, unless we exempt our methods from self-application, they can’t be coherent—that is, they will give inconsistent recommendations in possible cases. As Elga puts it, “in order to be consistent, a fundamental policy, rule, or method must be dogmatic with respect to its own correctness” (2010: 185).

³⁵ See Machuca (2017) for additional strategies to avoid the self-defeat objection.

it is generated by doubts about our ability to engage in rational, unbiased evaluation. Moreover, it is unclear how to develop techniques that will drive out what John Stuart Mill called “the fogs which hide from us our own ignorance” (1984 [1867]: 239).³⁶ Instead of attempting to lift the fog, I have simply attempted to illuminate it.

Acknowledgements. Thanks to Alexander Dinges, Nick Hughes, Christos Kyriacou, Martin Smith, and Mark Walker for valuable comments on a draft of this paper.

Bibliography

- Abrams, Dominic, and Michael Hogg (1988). Comments on the motivational status of self-esteem in social identity and intergroup discrimination. *European journal of social psychology* 18 (4): 317-334.
- Ahlstrom-Vij, Kristoffer (2013). Why we cannot rely on ourselves for epistemic improvement. *Philosophical Issues* 23 (1): 276-296.
- Ancell, Aaron (2019). The Fact of Unreasonable Pluralism. *Journal of the American Philosophical Association* 5 (4): 410-428.
- Bach, Kent (2005). The emperor's new 'knows'. In Gerhard Preyer & Georg Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press. pp. 51-89.
- Bacon, Francis (1620). *The New Organon*.
- Ballantyne, Nathan (2019). *Knowing Our Limits*. Oxford University Press.
- Blackburn, Simon (2019). How can we teach objectivity in a post-truth era? *New Statesman*. Accessed at: <https://www.newstatesman.com/culture/2019/02/how-can-we-teach-objectivity-post-truth-era>
- BonJour, Laurence (2010). The myth of knowledge. *Philosophical Perspectives* 24 (1):57-83.
- Brown, Jessica. (2018). *Fallibilism: Evidence and Knowledge*. Oxford University Press.
- Brunero, John & Kolodny, Niko (2013). Instrumental Rationality. *Stanford Encyclopedia of Philosophy*.
- Cohen, Stewart (1988). How to be a fallibilist. *Philosophical Perspectives* 2: 91-123.
- Davis, Wayne (2007). Knowledge claims and context: loose use. *Philosophical Studies* 132 (3):395-438.
- DeRose, Keith (2009). *The Case for Contextualism*. Oxford University Press.
- DeRose, Keith (1995). Solving the skeptical problem. *Philosophical Review* 104 (1):1-52.
- Descartes, Rene (1998). *Discourse on Method and Meditations on First Philosophy, 4th Ed*. Hackett.
- Ditto, Peter H., David A. Pizarro, and David Tannenbaum (2009). Motivated moral reasoning. *Psychology of learning and motivation* 50: 307-338.
- Dretske, Fred (1981). The pragmatic dimension of knowledge. *Philosophical Studies* 40 (3):363-378.

³⁶ I thank Nathan Ballantyne (2019: 64) for drawing my attention to this quote from Mill.

- Feldman, Richard (2001). Skeptical problems, contextualist solutions. *Philosophical Studies* 103 (1):61-85.
- Feldman, Richard (1999). Contextualism and skepticism. *Philosophical Perspectives* 13 (s13):91-114.
- Frances, Bryan (2008). Live Skeptical Hypotheses. In John Greco (ed.), *Oxford Handbook of Skepticism*. Oxford University Press. pp. 225-245.
- Frances, Bryan (2005). *Scepticism Comes Alive*. Oxford University Press.
- Gigerenzer, G, et al. 2001. *Simple Heuristics That Make Us Smart*. Oxford University Press
- Gilovich, Thomas (1993). *How We Know What Isn't So*. Free Press
- Greco, John (2008). Skepticism about the External World. In *The Oxford Handbook of Skepticism*. Oxford University Press. pp. 108-128.
- Glock, Hans-Johann (2010). The Development of Analytic Philosophy: Wittgenstein and After. In D. Moran (ed.), *The Routledge Companion to Twentieth Century Philosophy*. Routledge.
- Haidt, Jonathan (2012). *The Righteous Mind*. Penguin.
- Hannon, Michael (2019a). *What's the Point of Knowledge?* Oxford University Press.
- Hannon, Michael (2019b). Skepticism: Impractical, Therefore Implausible. *Philosophical Issues* 29 (1):143-158.
- Hannon, Michael (2017). A solution to knowledge's threshold problem. *Philosophical Studies* 174 (3):607-629.
- Hastorf, Albert H., & Cantril, Hadley. (1954). They saw a game; a case study. *The Journal of Abnormal and Social Psychology* 49(1): 129-134.
- Hazlett, Allan (2014). *A Critical Introduction to Skepticism*. Bloomsbury Academic.
- Hetherington, S. (2019). Fallibilism. *Internet Encyclopedia of Philosophy*.
- Hetherington, Stephen (2006). Knowledge's Boundary Problem. *Synthese* 150 (1):41-56.
- Hughes, Nick. Forthcoming. Evidence and Bias. *Routledge Handbook of the Philosophy of Evidence*. Routledge.
- Hume, David (2003). *A treatise of human nature*. Courier Corporation.
- Jolls, Christine, and Cass R. Sunstein (2006). Debiasing through law. *The Journal of Legal Studies* 35.1: 199-242.
- Kahneman, Daniel (2003). A perspective on judgment and choice: mapping bounded rationality. *American psychologist* 58.9: 697.
- Klein, Peter D. (2000). Contextualism and the Real Nature of Academic Skepticism. *Philosophical Issues* 10 (1):108-116.
- Kornblith, Hilary (2000). The Contextualist Evasion of Epistemology. *Noûs* 34 (s1):24-32.
- Kornblith, Hilary (1999). Distrusting reason. *Midwest Studies in Philosophy* 23 (1):181-196.
- Kunda, Ziva (1990). The case for motivated reasoning. *Psychological bulletin* 108.3: 480.
- Larrick, Richard P. (2004) Debiasing. *Blackwell handbook of judgment and decision making*. Blackwell. pp. 316-338.

- Lewis, David K. (1996). Elusive knowledge. *Australasian Journal of Philosophy* 74 (4):549-567.
- Lilienfeld, Scott O., Rachel Ammirati, and Kristin Landfield (2009). Giving debiasing away: Can psychological research on correcting cognitive errors promote human welfare? *Perspectives on psychological science* 4.4:390-398.
- Locke, John (1996) [1706]. *Some Thoughts Concerning Education and Of the Conduct of the Understanding*. Edited by Ruth Grant and Nathan Tarcov. Indianapolis, IN: Hackett.
- Lodge, Milton, and Charles S. Taber (2013). *The Rationalizing Voter*. Cambridge University Press.
- Lord, Charles G., Lee Ross, and Mark R. Lepper (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of personality and social psychology* 37(11): 2098.
- Machuca, Diego (2017). A neo-Pyrrhonian response to the disagreeing about disagreement argument. *Synthese* 194(5): 1663-1680.
- Mitchell, Terence R., et al. (1997). Temporal adjustments in the evaluation of events: The “rosy view”. *Journal of experimental social psychology* 33(4): 421-448.
- Neta, Ram (2011). A Refutation of Cartesian Fallibilism. *Noûs* 45 (4):658-695.
- Nisbett, Richard E., and Timothy D. Wilson (1977). Telling more than we can know: verbal reports on mental processes. *Psychological review* 84 (3): 231.
- Pasnau, Robert (2017). *After Certainty: A History of Our Epistemic Ideals and Illusions*. Oxford University Press.
- Peirce, Charles S. (1868). Some Consequences of Four Incapacities. *Journal of Speculative Philosophy* 2: 140-157.
- Pritchard, Duncan. (2019). *Skepticism: A Very Short Introduction*. Oxford University Press.
- Pronin, Emily, Daniel Y. Lin, and Lee Ross (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin* 28.3:369-381.
- Pryor, James (2000). The skeptic and the dogmatist. *Noûs* 34 (4): 517-549.
- Reed, Baron (2012). Fallibilism. *Philosophy Compass* 7 (9): 585-596.
- Reed, Baron (2008). Certainty. *Stanford Encyclopedia of Philosophy*.
- Reed, Baron (2002). How to think about fallibilism. *Philosophical Studies* 107 (2): 143-157.
- Roese, Neal J., and Kathleen D. Vohs (2012). Hindsight bias. *Perspectives on psychological science* 7.5: 411-426.
- Russell, Bertrand (2004) [1928]. *Sceptical Essays*. Routledge.
- Saul, Jennifer (2013). Scepticism and Implicit Bias. *Disputatio* 5 (37): 243-263.
- Siegel, Harvey (1997). *Rationality Redeemed?* Routledge.
- Sperber, Dan and Hugo Mercier (2017). *The Enigma of Reason*. Harvard University Press.
- Stroud, Barry (1984). *The Significance of Philosophical Scepticism*. Oxford University Press.
- Tavris, Carol, and Elliot Aronson (2008). *Mistakes were made (but not by me): Why we justify foolish beliefs, bad decisions, and hurtful acts*. Houghton Mifflin Harcourt.

Unger, Peter (1975). *Ignorance: A Case for Scepticism*. Oxford University Press.

Vavova, Katia (2018). Irrelevant Influences. *Philosophy and Phenomenological Research* 96 (1): 134-152.

Williams, Michael (2001). Contextualism, Externalism and Epistemic Standards. *Philosophical Studies* 103: 1-23.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.

Wilson, Timothy (2002). *Strangers to Ourselves*. Harvard University Press.

Wilson, Timothy, and Nancy Brekke (1994). Mental contamination and mental correction: unwanted influences on judgments and evaluations. *Psychological bulletin* 116 (1): 117.