



Cognitive bias, situationism, and virtue reliabilism

Steven Bland¹

Received: 17 July 2018 / Accepted: 19 November 2018 / Published online: 23 November 2018
© Springer Nature B.V. 2018

Abstract

Mark Alfano claims that the heuristics and biases literature supports inferential cognitive situationism, i.e., the view that most of our inferential beliefs are arrived at and retained by means of unreliable heuristics rather than intellectual virtues. If true, this would present virtue reliabilists with an unpleasant choice: they can either accept inferential skepticism, or modify or abandon reliabilism. Alfano thinks that the latter course of action is most plausible, and several reliabilists seem to agree. I argue that this is not the case. If situationism is true, then inferential non-skepticism is no more plausible than reliabilism. But inferential cognitive situationism is false. The heuristic-based inferences that facilitate successful perception and communication have proven remarkably accurate, and even the psychological research on inductive reasoning does not support Alfano's situationism. More generally, negative assessments of human reasoning tend to ignore the fact that the research on cognitive biases focuses primarily on the performance of individuals in isolation. Several studies suggest that we reason much more effectively when in critical dialogue with others, which highlights the fact that our epistemic performance depends not only on the inner workings of our cognitive processes, but on the environments in which they operate.

Keywords Cognitive bias · Virtue epistemology · Skepticism · Epistemic situationism

Mark Alfano's epistemic situationism is supposed to pose an empirical challenge to virtue epistemologies. Against responsibilist theories, he uses the situationist literature in social psychology to argue against the existence of stable epistemic traits (Alfano 2012), and against reliabilist theories, he uses the heuristics and biases literature to argue that our inferential faculties are insufficiently reliable to generate much in the way of knowledge (Alfano 2014). In light of recent replication studies, however, Alfano now recommends suspending judgement on the situationist challenge

✉ Steven Bland
sbland2@uwo.ca

¹ Department of Philosophy, Huron University College, 1349 Western Road, London, ON, Canada

to responsibilism, while insisting that the threat to reliabilism remains (Alfano 2017, p. 58).

Alfano's challenge to inferential reliabilism (IR) is that it is incompatible with a joint commitment to inferential non-scepticism (INS) and inferential cognitive situationism (ICS), neither of which should be abandoned. He claims that INS should be retained because it is an overwhelmingly popular position within the philosophical community, and ICS should be accepted because it is supported by a vast and growing empirical literature. ICS is the view that we acquire and retain most of our inferential beliefs by means of unreliable heuristics. Together with a commitment to INS, this implies that inferential knowledge cannot require reliable inferential faculties, contra IR.

Several important virtue epistemologists have made conciliatory responses to Alfano's challenge, effectively weakening their reliabilist theories of knowledge to accommodate INS and ICS. I will argue that this is a mistaken strategy and that instead we should focus our critical attention on Alfano's defense of INS and ICS. Indeed, it seems that his arguments for these two doctrines are inconsistent, such that he can be pressed with the following dilemma: if there is convincing empirical evidence for situationism, then the rationale for non-scepticism is unsuccessful, and if there isn't convincing evidence for situationism, then his challenge to reliabilism fails.

I will argue that the second horn of the dilemma should be embraced. There is evidence that heuristic-based inferences are generally reliable, and that they can actually outperform optimal deliberative reasoning at some cognitive tasks. These facts have motivated some philosophers to adopt an *ecological* conception of the epistemic virtues, according to which the virtue of an intellectual faculty consists in its being used in conditions in which it tends to facilitate valuable epistemic ends. The heuristics that have such a cognitive niche must be counted among these epistemic virtues. Alfano anticipates this move and objects that we overwhelmingly use heuristics outside of these conditions, so the threat of inferential skepticism remains. Against this claim, I offer evidence that reliable heuristics are indispensable to our abilities to perceive our physical environments and communicate within our social environments, and that Alfano's dire conclusions about our inductive reasoning are unwarranted.

Finally, I entertain an argument for situationism based on the broader literature on cognitive bias. The fact that human reasoning is replete with cognitive biases that routinely contaminate our judgements may be leveraged to motivate situationism. However, this is an empirical fact about how we reason *in isolation*; experiments looking for biased cognition generally focus on how individuals perform at various cognitive tasks. Yet much of our reasoning takes place in deliberation with others, and in many of these social environments our cognition tends to be significantly less biased. Moreover, some of our cognitive biases yield epistemic benefits for both individuals and groups in dialogical contexts. This is one more reason to reject epistemic situationism and adopt an ecological conception of the virtues.

1 The situationist threat

Virtue reliabilism is the view that knowledge is true belief that is acquired and retained through one or more truth-conducive intellectual virtue(s). Alfano claims that this the-

ory must be abandoned if we wish to maintain a non-skeptical position that recognizes the role of heuristics in the generation of inferential beliefs.

Heuristics are rules of thumb that function as cognitive shortcuts. They are often characterized as System 1 (or Type 1) processes, which is to say that they are automatic, involuntary, intuitive, and efficient forms of cognition. Alfano focuses on the *availability* and *representativeness* heuristics. The availability heuristic uses the ease with which a type of event comes to mind as an index of its probability: the more easily tokens come to mind, the more likely we think events of that type are to occur. The representativeness heuristic uses the perceived similarity between tokens and stereotypes as an index of their probabilities: the more representative the token, the more likely we deem it to be.

Heuristics yield outcomes much more efficiently than the deliberate reasoning of System 2 (or Type 2 cognition). Their drawback is that their outcomes are also more likely to be contaminated by biases. Tversky and Kahneman (1973) found that people think seven-letter words ending in ‘ing’ are more likely to occur in a section of text than seven letter words with ‘n’ in the sixth position. Of course, this cannot be the case, since every seven-letter word ending in ‘ing’ is a word with ‘n’ in the sixth position. They hypothesize that subjects arrived at these conclusions because they found it much easier to think of words of the first kind than words of the second. In short, their use of the availability heuristic led to inconsistent beliefs. In another well known experiment, Tversky and Kahneman (2002) presented subjects with the following description: “Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.” They were then asked to rank the statements below in order of their probabilities:

1. Linda is a teacher in elementary school.
2. Linda works in a bookstore and takes yoga classes.
3. Linda is active in the feminist movement.
4. Linda is a psychiatric worker.
5. Linda is a member of the League of Women Voters.
6. Linda is a bank teller.
7. Linda is an insurance salesperson.
8. Linda is a bank teller and is active in the feminist movement.

Tversky and Kahneman found that more than 80% of their subjects ranked (8) as being more probable than (6), thus committing the conjunction fallacy, i.e., assigning a higher probability to a conjunction than to one of its conjuncts. They explain this result by invoking the representativeness heuristic: because Linda’s description is more representative of a bank teller who is active in the feminist movement than it is of a bank teller, subjects generally thought of the former as being more probable than the latter. Anyone reasoning in conformity with the principles of probability theory could not have arrived at this result.

The groundbreaking work of Tversky and Kahneman in the 1970s spurred an entire movement within psychology, now known as the heuristics and biases approach. Its results have led some psychologists and philosophers to the pessimistic conclusion

that we deploy unreliable heuristics more often than sound reasoning. Alfano calls this view *inferential cognitive situationism*:

(*inferential cognitive situationism*) People acquire and retain most of their inferential beliefs through heuristics rather than intellectual virtues.

He then points out that this position cannot be consistently held with two other widespread commitments in epistemology (Alfano 2014, p. 109):

(*inferential non-skepticism*) Most people know quite a bit inferentially.

(*inferential reliabilism*) Inferential knowledge is true belief acquired and retained through inferential reliabilist intellectual virtue.

If most people have extensive inferential knowledge, but they acquire this knowledge by means of unreliable heuristics, then reliabilist accounts of knowledge must be mistaken. If reliabilism is true, then we can possess extensive inferential knowledge only if situationism is false. And if reliabilism and situationism are both true, then we are left with inferential skepticism. This is known as Alfano's *inconsistent triad*.

To escape the triad, we must give up one of the three commitments. The question is: which one? Alfano insists that "...if one of the three propositions must go, it's unlikely to be *inferential non-skepticism*" because it is "near orthodoxy" among epistemologists, and philosophers more generally (Ibid.). In support of this contention, he cites the fact that 81.6% of philosophers and 84.3% of epistemologists rejected skepticism in a recent PhilPapers survey. Furthermore, he claims that the heuristics and biases literature provides robust empirical support for ICS. Therefore, he concludes that IR is the tenet that should be given up.

Alfano's argument has drawn a number of responses from important virtue epistemologists, many of which are conciliatory. John Turri offers abilism as an alternative to virtue reliabilism (Turri 2017). On this view, "knowledge is true belief manifesting cognitive ability", and persons possess a cognitive ability to detect the truth when their use of the faculty produces true beliefs at a rate exceeding chance. The move from reliabilism to abilism effectively lowers the threshold for knowledge, so that unreliable belief-forming processes, such as the availability and representativeness heuristics, can produce knowledge, as long as they generate true beliefs at a rate greater than chance.¹ Abilism, then, can accommodate both the orthodox position of INS and the empirically well founded position of ICS.

Duncan Pritchard has made a similar move in response to Alfano's situationist challenge. He and J. Adam Carter make the following distinction between *robust* and *modest* virtue epistemologies:

On the robust view, rationally grounded knowledge results only when the agent's cognitive success (i.e., her true belief) is *primarily attributable* to her exercise of (relevant) cognitive ability. In contrast, according to the modest view, all that is required is that the agent's cognitive success is *significantly attributable* to her exercise of cognitive ability. That is, modest virtue epistemology allows that the agent's cognitive success need not be primarily attributable to her exercise of

¹ Alfano notes that Sosa (2017) also reacts to the situationist threat by lowering the bar for knowledge, while remaining committed to reliabilism (Alfano 2017, n. 11).

cognitive ability, and so enables factors outside of the subject’s manifestation of cognitive agency, such as epistemically friendly features of her physical or social environment, to play an explanatory role in her cognitive success, including explanatory roles that will be (once discovered) surprising (Carter and Pritchard 2017, p. 176).

When someone uses a heuristic to generate a true belief, this success is partly the result of the believer’s epistemic agency, and partly the result of a friendly epistemic environment; this prevents the accomplishment from being *primarily* attributable to the exercise of her cognitive ability, but not to its being *significantly* attributable thereto. Consequently, modest virtue epistemologies are also compatible with both INS and ICS.

Alfano correctly reports that,

Within this dialectic, then, following Turri or Pritchard should be seen as a concession to epistemic situationism. Robust virtue epistemology is abandoned in favor of a weaker theory of knowledge. It’s a concession that I am happy to accept, but some virtue epistemologists may find it too much to stomach (Alfano 2017, p. 56).²

My purpose in the following two sections is not to comment on these weakened theories of knowledge, but to motivate alternative responses that target the components of the inconsistent triad that Alfano, Turri, and Pritchard leave untouched: INS and ICS. In particular, I argue that Alfano’s rationale for INS is inconsistent with the truth of ICS, and that ICS is not well supported by empirical evidence. This being the case, epistemic situationism does not pose a threat to virtue reliabilism, and should not be addressed by weakening the theory.

2 Against inferential non-skepticism

As mentioned in the previous section, Alfano uses the results of a PhilPapers survey to support his claim that INS is “near orthodoxy” among philosophers. However, the survey question does not concern *inferential* skepticism, but a much more radical, *Cartesian* form of skepticism. Radical forms of skepticism, such as those of the Pyrrhonian and Cartesian variety, are committed to the claims that: (1) we know nothing about the domain in question, and (2) there is nothing we can do to improve our epistemic standing.³ Epistemologists overwhelmingly reject radical forms of skepticism, for if (1) and (2) are true, then epistemology cannot meet any of its apologist or ameliorative aims (respectively). Radical skepticism, then, is an attack on the very projects in which most epistemologists are engaged. Inferential skepticism, on the other hand, constitutes no such attack. As Alfano formulates it, inferential skepticism is the view that

² Alfano then proceeds to voice his own reservations concerning this reaction to epistemic situationism.

³ Different forms of radical skepticism concern different domains: while Pyrrhonian skeptics deny the possibility of any knowledge whatsoever, Cartesian skeptics deny knowledge of the external world.

most people know less than quite a bit inferentially.⁴ This leaves open the possibility that epistemologists can determine what it is that we *do* know, and design strategies to *improve* our epistemic standing. So philosophers need not reject inferential skepticism to salvage the projects of traditional epistemology. And it's not clear that philosophers do reject this attenuated form of skepticism; our extensive training in formal logic and informal argumentation suggests that we are dubious of peoples' normal capacities to make, reconstruct, and evaluate inferences. Consequently, without any further evidence, there is no reason to think that INS is near orthodoxy among philosophers.

For the sake of argument, though, let's suppose that most philosophers do endorse INS. Might this not be because they tacitly *reject* ICS? One possible rationale for thinking that we do know quite a bit inferentially is the supposition that we acquire and retain most of our inferential beliefs through *reliable* cognitive processes. Indeed, this is often touted as a reason for endorsing reliabilism. Consider, for example, the problem of forgotten evidence. Reliabilists claim that internalists cannot countenance forgotten evidence as a source of justification. And yet, many of our true beliefs were adopted on the basis of evidence that we can no longer recall. According to reliabilism, these beliefs do constitute knowledge if the processes that generated them are sufficiently reliable. This argument—that reliabilism resists a form of skepticism that internalism cannot—can be persuasive only if most philosophers believe that our cognitive faculties, including inferential ones, are generally reliable. If this is the case, then we have the same reason to reject ICS that Alfano gives to retain INS: they are both orthodox positions among epistemologists. Of course, Alfano may reply that there is empirical evidence that speaks against the former course of action, but this will be dispelled in the next section.

Let's go even one step further and suppose that ICS is true. I submit that this would invalidate Alfano's rationale for maintaining INS, for in arriving at this position philosophers rely on their *inferential* faculties. Furthermore, the fact that philosophers find non-skepticism *intuitively* obvious, such that it is typically a position that is argued *from* rather than argued *for*, suggests that it is the result of heuristic-based inferences rather than careful deliberation.⁵ If heuristic-based inferences are generally unreliable, then the fact that INS is a nearly orthodox position among philosophers (if it is one) is not the definitive piece of evidence that Alfano takes it to be. Moreover, there are good reasons for thinking that our acceptance of this position could be the result of the overconfidence effect and knowledge illusion rather than sound inferential practice: philosophers, like everyone else, are likely to mistake some of their unjustified/untrue beliefs for knowledge (overconfidence effect), and to attribute to themselves knowledge that is widely distributed throughout their communities rather

⁴ This formulation of inferential skepticism is highly implausible. Given a weak principle of epistemic closure, we can be said to know a tremendous amount inferentially. I know, for example, that because it is the year 2018, it is not the year 2017, 2016, 2015, etc. A more plausible formulation of inferential skepticism, and one in keeping with Alfano's concerns about our use of heuristics, would exclude these sorts of deductively closed trivial truths. This is a much *weaker* doctrine than any form of radical skepticism. I am thankful to Jon Marsh for pointing this out to me.

⁵ Alternatively, one might want to classify this belief as the result of intuition *rather than* inference. This re-classification does nothing to address my objection, however, since the heuristics and biases literature that Alfano relies on to make his case for ICS may be used to call the reliability of this intuition into doubt, as I do in the remainder of this paragraph.

than possessed by any one person (knowledge illusion). So, Alfano is faced with the following dilemma: if ICS is true, then we should ignore the fact that INS is widely believed among philosophers (if it is), and if it is false, then it should be the doctrine that gets rejected. In the next section, I will argue for taking the second horn of the dilemma.

3 Against inferential cognitive situationism

After reviewing the literature on the availability and representativeness heuristics, Alfano offers the following argument for ICS:

The robustness of the representativeness heuristic throws a pall of doubt over the notion that most people possess the intellectual virtues related to even rudimentary deductive and inductive reasoning. The process used to arrive at beliefs about likelihood in no way resembles sound inferential practice; rather, people follow a heuristic that treats representativeness as an index of probability. My claim is that the same holds true for other heuristics, and that this creates trouble not only for the cognitive virtues related to deductive and inductive reasoning, but for many of the other cognitive virtues related to inference, such [as] abduction (Alfano 2014, pp. 114–115).

I will reconstruct his argument as follows:

- (MA1) The availability heuristic is an unreliable source of true beliefs.
- (MA2) The representativeness heuristic is an unreliable source of true beliefs.
- (MA3) Our inferential practices rely principally on heuristics.
- (MA4) Therefore, our inferential practices are generally unreliable.

Alfano's crucial supposition is that heuristics are insufficiently reliable to constitute the intellectual virtues required for knowledge. This is made explicit in his formulation of ICS: People acquire and retain most of their inferential beliefs through heuristics *rather than* intellectual virtues. On this view, the intellectual virtues belong exclusively to the more deliberative and effortful cognitive processes of System 2: sound logical and mathematical reasoning, Bayesian inference, reasoning from counterfactuals, etc. Alfano thus thinks that cognition involves an *accuracy-efficiency trade-off*: the more efficient a cognitive process is, the less likely it is to deliver accurate information.

A number of facts speak loudly against Alfano's supposition. First, as Kahneman repeatedly emphasizes, heuristics are *generally reliable* sources of accurate beliefs: "The heuristic answers are not random, and they are often approximately correct" (Kahneman 2011, p. 416). Using the representativeness heuristic in the absence of relevant base-rates and reliable information about the individual(s) in question is a good strategy, at least when the operative stereotype is accurate.⁶ Our assumption that people who act friendly are friendly facilitates many more true predictions than false ones. The problem that Tversky and Kahneman (1973) identify is not with the application of the representativeness heuristic, but with its *over-application*, i.e., its use when

⁶ For an important collection of papers on stereotype accuracy, see Lee et al. (1995).

base-rates are known and/or reliable information is available. Second, heuristics can outperform ideal epistemic rules. Gerd Gigerenzer illustrates this point by considering two strategies for estimating the path of a projectile. Baseball players accomplish this feat routinely, but not the way that physicists do; instead of calculating the path of a baseball by solving a set of differential equations, players use the *gaze heuristic*: “fix your gaze on the ball, start running, and adjust your running speed so that the image of the ball rises at a constant rate” (Gigerenzer 2007, p. 10). By relying on this heuristic rather than calculus, Gigerenzer argues, ballplayers can solve the problem of projectile motion more efficiently *and more accurately*, since they are more likely to miscalculate than to misapply the simple rule. The remarkable success of statistical prediction rules is another case in point. It is tempting to think that complex problems require complex solutions, yet when it comes to making predictions in these domains, simple rules have routinely outperformed expert judgement.⁷ For example, Carroll et al. (1988) show that a statistical rule for predicting criminal recidivism based on individuals’ criminal and prison records yields more accurate results than the predictions of criminologists, which take several more (apparently irrelevant) factors into account. There is no accuracy-efficiency trade-off in these cases: more information and greater reflection lead to poorer results. Finally, several epistemologists have correctly pointed out that cognitive efficiency is both a pragmatic and an *epistemic* virtue: “It is beneficial to one’s narrowly conceived cognitive success (achieving true beliefs), as well as overall intellectual flourishing, to utilize efficient methods for acquiring and maintaining true beliefs” (Fairweather and Montemayor 2017, p. 42). Focussing inordinate cognitive resources on a narrow range of problems yields diminishing or negative returns: we update relatively few beliefs; ignore evidence pertaining to related problems; find too much evidence in favour of false hypotheses; etc. Our beliefs are more likely to be comprehensive and accurate when we think efficiently; since heuristics are an indispensable means of achieving cognitive efficiency, they are also means of attaining these epistemic ends.

Alfano’s supposition that sound reasoning alone can yield what virtue reliabilists would call inferential knowledge is mistaken. This is not to say that heuristics are in and of themselves intellectual virtues, but that they can be *used* virtuously to solve problems to which they’re well suited. When attempting to determine the likelihood of a terrorist attack, we should use statistical reasoning rather than the availability heuristic; when estimating the path of a projectile will land, we should use the gaze heuristic rather than differential calculus. Properly evaluating epistemic behaviour isn’t just a matter of comparing operative cognitive processes with ideal epistemic norms; we must also take into account the relevant features of the *environments* in which cognitive processes operate, to determine whether or not the latter are properly responsive to the former. What’s needed, then, is an *ecological* conception of epistemic virtue.⁸ Adam Morton provides one such conception, according to which the virtues

⁷ For a classic formulation of this finding, see Meehl (1954). For a synopsis of more recent evidence, together with a philosophical analysis of its implications for epistemology, see Bishop and Trout (2002).

⁸ This move parallels Gigerenzer’s move to an *ecological conception of rationality*, which “...refers to the study of how cognitive strategies exploit the representation and structure of information in the environment to make reasonable judgments and decisions” (Gigerenzer 2000, p. 57).

are "...sensitivities to features of the environment linked to sensitivities of the person's own capacities" (Morton 2012, p. 60).

The move to an ecological virtue theory does not by itself undermine ICS. Alfano admits that heuristics are truth-conducive in *some* circumstances, but claims that we overwhelmingly use them *outside* of those circumstances (Alfano 2014, p. 116). In other words, he denies that we manifest the sensitivities required for much of what we take ourselves to know inferentially: many of our inferential beliefs are false because they're produced by cognitive processes that are unreliable in the vast majority of the circumstances in which we use them. In the remainder of this paper, I will argue that this view is demonstrably false.

First it should be noted that because Alfano does not specify what he means by 'inference', his doctrine of ICS is ambiguous. There are several philosophical accounts of inference, some more inclusive than others. For example, according to the *reckoning model*, a subject infers a conclusion on the basis of information when s/he reckons that the former is rationally supported by the latter. On the *response model*, a conclusion need only be a response to an informational state in order to constitute an inference.⁹ I have no interest in defending any particular conception of inference here, but wish only to point out that the implications of ICS depend on which conception one adopts: the more inclusive the conception, the more radical a thesis ICS is. However, on any reasonable conception of inference, ICS is empirically false. If we construe 'inference' broadly, then our ability to accurately perceive our surroundings and successfully communicate belies ICS; if we construe it more narrowly, as I suspect Alfano does, then ICS is undermined by empirical research into inductive reasoning.

3.1 The broad reading of ICS

It has long been known that the contents of perceptual states and beliefs outstrip the information presented to our sense organs. Psychologists, beginning with Helmholtz, have argued that this is because perception involves "unconscious inferences". Steven Pinker makes this point specifically about vision:

Vision has evolved to convert...ill-posed problems into solvable ones by adding premises: assumptions about how the world we evolved in is, on average, put together. For example...the human visual system "assumes" that matter is cohesive, surfaces are uniformly colored, and objects don't go out of their way to line up in confusing arrangements. When the current world resembles the average ancestral environment, we see the world as it is. When we land in an exotic world where the assumptions are violated—because of a chain of unlucky coincidences or because a sneaky psychologist concocted the world to violate the assumptions—we fall prey to an illusion (Pinker 1997, pp. 212–213).

The worlds in which our assumptions are violated are exotic because they are unusual. We assume that objects remain the same size when they recede from us, even though they appear smaller. This assumption of size invariance is generally true and beneficially deployed by our visual systems to 'correct' the images that appear on our

⁹ For a clear delineation of the reckoning and response theories of inference, see Siegel (2017, Ch. 5).

retinas. When presented with cleverly constructed two-dimensional images, we can be fooled by distance cues into thinking that one object is larger than another when they are in fact the same size. However, these sorts of visual illusions serve more to reveal how our visual systems consistently get things right, than to induce worry about the possibility that our visual beliefs are consistently mistaken.

These unconscious inferences don't involve reckoning, but they are responses to information, so they would be included in more permissive theories of inference:

By and large, those interested in the epistemic role of perceptual experience have not concluded that experiences epistemically depend on the inferences psychologists describe. But if the process of combining stored generalizations with incoming information from a particular situation operates over epistemically powerful (or even epistemically appraisable) states, then it has all the hallmarks of...inferential responses... (Siegel 2017, p. 106).

Alfano explicitly categorizes perceptual knowledge as non-inferential, which suggests that he subscribes to a theory of inference more like reckoning than response. The problem is that these perceptual assumptions seem to function just as cognitive heuristics do: “I believe that intuitive judgments work the same way as these perceptual bets. When given insufficient information, the brain makes things up based on assumptions about the world” (Gigerenzer 2007, pp. 42–43). If Alfano is going to count heuristic judgements as inferential, then he should count perceptual judgements as inferential as well, in which case the manifest reliability of perception counts as evidence against ICS.

Communication is also made possible by our robustly reliable inferential abilities. Fairweather and Montemayor argue: “If Alfano insists that inference must be something rule based, formal, and regimented, then one can hardly think of a type of inferential process that satisfies these constraints better than knowledge of syntax” (2017, p. 46). Children are remarkably proficient at distinguishing well formed from ill formed grammatical structures, despite their lack of formal training and exposure to unambiguous information. This success is best explained by their ability to *accurately infer* complicated formal rules of syntax from impoverished and messy evidence. As is the case with perception, these inferences are made effortlessly and without any reflection whatsoever. More generally, when learning a language we engage in what Davidson (1986) calls *radical interpretation*: the process of interpreting speakers' utterances without knowing what they believe or what their expressions mean. To accomplish this feat, Davidson argues, we must be able to *accurately infer* the causes of speakers' utterances: by inferring that a speaker was prompted to say ‘gavagai’ by the sudden appearance of a rabbit, I come to know that the speaker believes that she sees a rabbit, and that the word ‘gavagai’ means rabbit in her language. If speakers or interpreters consistently misidentified the things that are talked about, radical interpretation would be impossible. The fact that we do learn languages is thus a testament to the reliability of our inferential capacities. So too is the process of ordinary interpretation, i.e., the process we engage in when communicating with others in a common language. Interpreters must bring to bear on this process not only their knowledge of the language—its syntax, semantics, and pragmatics—but norms conversational implicature, salient contextual features, and significant background knowledge. To know the literal

meaning of a sentence like “You really know what you’re talking about” requires a host of inferences, but to know the speaker’s meaning requires many more inferences still. And yet, we perform these inferences quickly, effortlessly, and involuntarily; they too are products of System 1. Moreover, these inferences are *overwhelmingly accurate*; while miscommunication does occur, it is a relatively rare occurrence.

Negotiating our way through the physical world and communicating with one another require substantial and reliable inferential successes, most of which are accomplished by System 1. We often don’t notice these routine successes, in part because we expect them; cases of systemic inferential failure are much more interesting precisely because they’re so surprising. And because they are more interesting, they are more likely to figure in general appraisals of our inferential abilities. We might thus conclude that Alfano’s situationism is itself the product of the salience effect and availability heuristic, rather than sound inferential practice. Alternatively, he could be operating with a narrower notion of inference that excludes the abilities discussed in this section, despite their structural similarities to the heuristic-based reasoning that he focuses on. In any case, the heuristics and biases literature does not even support Alfano’s situationism about inductive reasoning, as I will argue in the next section.

3.2 The narrow reading of ICS

Discussing his own reaction to Kahneman and Tversky’s work on the representativeness heuristic, Stephen J. Gould writes (1992, p. 469):

I am particularly fond of [the Linda] example, because I know that the [conjunction] is least probable, yet a little homunculus in my head continues to jump up and down, shouting at me, “but she can’t just be a bank teller; read the description.”...Why do we consistently make this simple logical error? Tversky and Kahneman argue, correctly I think, that our minds are not built (for whatever reason) to work by the rules of probability.

It is this kind of thinking that leads Alfano to conclude: “The robustness of the representativeness heuristic throws a pall of doubt over the notion that most people possess the intellectual virtues related to even rudimentary deductive and inductive reasoning”. Yet, we should not conclude from Kahneman and Tversky’s results “that our minds are not built to work by the rules of probability”, i.e., that they generally rely on heuristics rather than virtuous reasoning.¹⁰ Gould’s own description of his struggle reveals why this is so: he struggles against his homunculus because he *knows* that it is committing the conjunction fallacy. To know this, he must reason in conformity with the rules of probability theory. Thus, the Linda case does not show that individ-

¹⁰ In fairness to Alfano and Gould, Kahneman and Tversky themselves drew this conclusion from their early work:

In making predictions and judgments under uncertainty, people do not appear to follow the calculus of chance or the statistical theory of prediction. Instead, they rely on a limited number of heuristics which sometimes yield reasonable judgments and sometimes lead to severe and systematic errors (Kahneman and Tversky 1973, p. 273).

uals lack the capacity or inclination to reason virtuously; rather it shows that intuition can sometimes *interfere* with our doing so. Kahneman explains: “The [Linda] problem...sets up a conflict between the intuition of representativeness and the logic of probability” (Kahneman 2011, p. 157). In the absence of this conflict, our judgements generally obey the laws of probability, including the conjunction rule: we know immediately that it’s more likely that Mark has hair than that he has blonde hair, and that Jane is more likely a teacher than a teacher who walks to school. Alfano and Gould’s pessimistic conclusion can be drawn only if intuition interferes with our probabilistic reasoning *most of the time*, and the results Kahneman and Tversky’s carefully designed experiments certainly don’t support this generalization.

Gigerenzer argues that even in the Linda experiment, subjects’ intuitions are not interfering with their probabilistic reasoning. He claims that people interpret probabilities as frequencies. Since subjects in the Linda experiment are asked for the probability of a *single event*—the probability that Linda is a bank teller; a bank teller who is active in the feminist movement; etc.—and single events are *not* frequencies, they rightly refrain from probabilistic reasoning. Instead of asking subjects to rank the probabilities of statements about Linda, Hertwig and Gigerenzer (1999) present the task as follows:

There are a hundred persons who fit the description above (i.e., Linda’s). How many of them are
bank tellers?
bank tellers and active in the feminist movement?

When the task is presented in this way, more than 90% of subjects did *not* commit the conjunction fallacy. Gigerenzer and Hoffrage (1995) also found that subjects generally don’t neglect base rates when dealing with frequencies rather than probabilities. From results such as these, Gigerenzer (2000, Ch. 12) draws the controversial conclusion that many of the cognitive illusions that figure prominently in the heuristics and biases literature simply disappear once we understand that people are uncompromising frequentists. A more modest conclusion is that our minds are not built to work by the rules of *mathematical* probability, but they are built to handle problems framed in terms of frequencies. This is enough to undermine Alfano and Gould’s radical pessimism.

Alfano follows the heuristics and biases literature in focussing exclusively on the distorting influence of non-statistical heuristics to our probabilistic/statistical thinking. But Nisbett et al. (2002) argue that this narrow focus obscures the fact that peoples’ inductive thinking also makes use of various *statistical* heuristics:

[The heuristics and biases literature] indicates that nonstatistical heuristics play an important role in inductive reasoning. But it does not establish that other heuristics, based on statistical concepts, are absent from people’s judgmental repertoire. Indeed, if one begins to look for cases of good statistical intuitions in everyday problems, it is not hard to find some plausible candidates (Nisbett et al. 2002, pp. 511–512).

They identify several common expressions whose use seems to indicate that people deploy sound statistical heuristics. The use of the expressions “beginner’s luck” and “nowhere to go but up/down” may encourage regressive predictions. Likewise, the

expressions “Don’t judge a book by its cover” and “All that glitters is not gold” may be used to overcome sample bias. Thus, once again, Alfano’s contrast between intellectual virtues and cognitive heuristics is faulty: in addition to non-statistical heuristics, such as availability and representativeness, we also use statistical heuristics to mitigate bias and facilitate truth-conducive inferences. The latter are *both* heuristics *and* intellectual virtues.

Nisbett et al. note that statistical heuristics are not used consistently, but are more likely to be deployed when sample spaces and sampling processes are clear, the role of chance factors in producing events are recognized, and there are cultural prescriptions to reason statistically (pp. 513–516). The situationist may argue that these conditions are met too infrequently to make our statistical reasoning a generally reliable source of accurate beliefs. This point is fair enough when it’s made about *most* people. But ICS is a thesis about the epistemic practices of *all* people. As a result, it fails to recognize the significant variability in the quality of peoples’ probabilistic/statistical reasoning. This variability is due, in large part, to peoples’ uneven access to statistical training. Nisbett et al. (1987) report that such training significantly improves peoples’ statistical reasoning, especially when it emphasizes how statistical principles, such as regression to the mean, apply to everyday circumstances. A study by Lehman et al. (1988), found that two years of graduate training in psychology resulted in an 80% improvement in subjects’ abilities to apply appropriate statistical rules to both scientific and everyday problems. Results like these speak against the situationist’s general claim that people acquire most of their probabilistic/statistical beliefs through unreliable heuristics rather than intellectual virtues. Even if this is true of most people—as I’ve argued above, it’s not clear that it is—there is reason to think that it is *not* true of those people who have undergone prolonged statistical training that focuses, at least in part, on everyday problems.

This response can be construed as advocating for a kind of epistemological elitism: a select group of people possess the intellectual virtues required for statistical/probabilistic knowledge, and everyone else does not. Alfano says that such a repudiation of virtue egalitarianism is tantamount to “admitting partial defeat”:

I don’t find this response appealing in the case of virtue ethics, and it seems even less appealing for virtue epistemology, for it would entail a great deal of skepticism. Yes, some people have knowledge, but they’re an elite epistemic minority. This flies in the face of the Moorean platitude of non-skepticism (Alfano 2014, p. 115).

I am not making this general claim. In fact, when it comes to the reliability of the inferential processes involved in perception, communication, and social cognition, I think that everyone is more or less on an equal epistemological footing. Furthermore, as I’ve indicated, I think we are on an equally *good* epistemological footing, such that we know much of we take ourselves to know. But it seems clear that other epistemic virtues, including sound probabilistic-statistical reasoning, are distributed more unevenly in the population. And it’s hardly more controversial to claim that effective training plays a significant role in this uneven distribution. This is not an admission of defeat. Someone who claims that the intellectual virtues required to generate knowledge in

theoretical physics are unevenly distributed because the training required to develop these virtues is unevenly distributed cannot be credibly accused of skepticism.

The heuristics and biases literature has shown that our judgements about statistical matters systematically depart from the strictures of mathematical probability theory. This may be the result of intuitions interfering with our probabilistic reasoning, or it may be because our minds deal better with frequencies than with probabilities, or perhaps some combination of these explanations is correct. In any case, the pessimism of Alfano and Gould is unwarranted. And the fact that people deploy sound statistical heuristics, and can be trained to do so more often, is grounds for optimism. All of this is to say that ICS is not supported by psychological research on patterns of statistical inference.

The literature on heuristics does not exhaust the empirical work on cognitive biases, however. Though Alfano focuses on the former, his case for ICS might be significantly bolstered by drawing on the likes of: hindsight bias; overconfidence effect; bias blindspot; outcome bias; authority bias; omission bias; self-serving bias; cognitive dissonance; confabulation; and the like. If our judgements are routinely contaminated by such biases, then it's *prima facie* plausible that our inferential processes are much less reliable than we think they are, and insufficiently reliable to avoid a troubling form of skepticism given a commitment to virtue reliabilism.¹¹

In the following section, I argue that this more general argument for ICS fails as well. While individuals do manifest cognitive biases in experimental conditions, these conditions seldom resemble the commonplace dialogic contexts in which we reason with others. And several experiments reveal that when we are in these contexts, our biases tend to be notably mitigated. Furthermore, there is a plausible line of argument that suggests that our being subject to cognitive biases when reasoning in groups can be an epistemic *benefit*. I will explore this line of argument by focussing on Mercier and Sperber's (2011, 2017) interactionist analysis of the role of confirmation bias in group cognition.

4 Reasoning better together

Confirmation bias, or myside bias, is the tendency to preferentially seek out, remember, and interpret evidence in ways that confirm one's pre-existing beliefs. Its existence is supported by decades of empirical research, and can be used to explain the adoption and persistence of many irrational beliefs. Nickerson reports:

Most commentators, by far, have seen the confirmation bias as a human failing, a tendency that is at once pervasive and irrational. It is not difficult to make a case for this position. The bias can contribute to delusions of many sorts, to the development and survival of superstitions, and to a variety of undesirable states of mind, including paranoia and depression. It can be exploited to great

¹¹ This line of argument constitutes what Carter and Pritchard (2017) call *bias-driven skepticism*. They find it not only in Alfano's work on situationism, but in Saul's claim that "...what we know about implicit biases shows us that we have very good reason to believe that we cannot properly trust our knowledge-seeking faculties" (2013, p. 243).

advantage by seers, soothsayers, and fortune tellers, and indeed anyone with an inclination to press unsubstantiated claims. One can also imagine it playing a significant role in the perpetuation of animosities and strife between people with conflicting views of the world (Nickerson 1998, p. 205).

More generally, confirmation bias interferes with an individual's capacity to distinguish true beliefs from false ones by restricting the relevant evidence at their conscious disposal. If this tendency is as pervasive as Nickerson suggests, then it may pose a threat capable of motivating ICS.¹²

Even more distressingly, confirmation bias is difficult to mitigate. One of the more successful debiasing strategies implemented by psychologists is *consider the opposite*, which involves prompting subjects to think about conditions in which their beliefs would be false (Fischhoff 1982; Pronin et al. 2002; Wilson et al. 2002; Sedikides and Gregg 2007). While this intervention has been successful in experimental contexts, Kenyon and Beaulac argue that it has significant limitations:

The problem is that the strategy is extremely difficult to implement as a self-deployed skill. Existing biases and attentional limits can easily make themselves felt as an unwillingness or inability to generate plausible alternative scenarios (O'Brien 2009, pp. 329–330); and even a willingness to do so is no guarantee that the generation and consideration of alternatives will be sufficiently disciplined or constrained to actually lead to a less distorted judgment (Tetlock 2005, p. 199). Absent the sort of facilitation or guidance by assistants that tends to characterize the experimental contexts in which “consider the opposite” is an effective strategy, there is little reason to expect it to be employed with regularity by individual agents in normal contexts, nor to work well when it is employed (Kenyon and Beaulac 2014, p. 347).

One major obstacle to successfully implementing this strategy is our overconfidence about the quality of our judgements and intellectual abilities (Hoffrage 2004). This overconfidence can lead to the belief that we are immune to biases that influence the judgements of others, a phenomenon known as *bias blindspot* (Pronin et al. 2002). And this tendency can be reinforced by our ignoring or discounting evidence of our biases and other intellectual shortcomings. In this way confirmation bias is self-reinforcing, such that we often don't see the need to deploy the consider the opposite strategy, or we believe that we've done so more effectively than we actually have.

These empirical results call into question not only the reliability of our reasoning faculties, but what Mercier and Sperber (2011, 2017) call the *intellectualist approach* to understanding reason. According to this approach, reason evolved to serve the function of improving the accuracy of individuals' beliefs and the rationality of their decisions. If overconfidence and confirmation bias routinely blind us to our cognitive

¹² Two cautionary points are worth emphasizing. First, it remains to be empirically established that confirmation bias is sufficiently ubiquitous to pose a threat to our inferential cognition generally. Second, confirmation bias has a possible upside: while it sometimes prevents us from abandoning false beliefs, it can also decrease our chances of abandoning *true* beliefs. Thus, the existence of confirmation bias can be used to bolster ICS only if there are *independent* grounds for thinking that our inferential processes produce a significant number of false beliefs. Even if both of these claims can be established, however, there are reasons to be dubious of the situationist's pessimistic conclusion, as I will argue below.

shortcomings, then reason is ill equipped to serve this ameliorative function. Furthermore, Mercier and Sperber survey an empirical literature that suggests that we are lazy and careless when evaluating reasons for our own beliefs, and demanding and vigilant when evaluating the reasons of others. Once again, these results are out of step with the intellectualist model. On the other hand, they fit very well with the *interactionist approach* to reason, according to which "...the function of reason is to produce and evaluate justifications and arguments in dialogue with others" (Mercier and Sperber 2017, p. 203). Since we use reason to justify ourselves and convince others, it makes sense that it would be biased in our favour. And because we cannot often anticipate what reasons someone else will find convincing before engaging them in discussion, we don't usually expend much effort in justifying ourselves ahead of time. On the other hand, it is in our best interest to effectively discriminate good reasons from bad ones when someone else is justifying themselves to us.

From the interactionist point of view, it is generally a *good* thing that individuals are stubbornly subject to confirmation bias, for when individuals reason together this leads to an *efficient division of cognitive labour*. Identifying beliefs that are well grounded and immune to counter-arguments is a difficult and time-consuming affair, as academics are well aware. It requires that we search a vast space of reasons, identify those that are relevant to our beliefs, evaluate their cogency, and respond accordingly. All of this is accomplished much more efficiently if we subject our beliefs to the critical attention of others rather than searching for first-rate reasons on our own. The feedback we get from this process can strategically direct our epistemic efforts by revealing which beliefs/reasons need updating and when our reasons are rationally effective. In doing so, it also mitigates the sorts of egocentric processing that are responsible for confirmation bias. In other words, the consider the opposite strategy is more effectively and efficiently implemented by *groups* than it can be individually.

More generally, if everyone is better at finding errors in other peoples' thinking than in their own, then each person's errors are more likely to be found by others than by themselves. Consequently, we should expect that most cognitive biases are less prevalent and less pronounced in deliberative groups than in individuals. And in fact Mercier and Sperber (2017, pp. 264–265) review a number of empirical studies that support this conclusion. It is also supported by Tetlock et al.'s *Good Judgement Project*, which found that teams of individuals forecasting social-political events consistently outperformed individual forecasters:

On average, when a forecaster did well enough in year 1 to become a superforecaster, and was put on a superforecaster team in year 2, that person became 50% more accurate. An analysis in year 3 got the same result. Given that these were collections of strangers tenuously connected in cyberspace, we found that result startling (Tetlock and Gardner 2015, p. 205).¹³

¹³ 'Superforecasters' is Tetlock's term for individuals who outperform the vast majority of forecasters. Tetlock found similar results more generally, i.e. with regular forecasters as well: "At the end of the [first] year, the results were unequivocal: on average, teams were 23% more accurate than individuals" (Tetlock and Gardner 2015, p. 201).

This marked improvement in forecasting accuracy can be explained by the fact that constructive, critical dialogue is an effective means of aggregating information and mitigating biases.

Most experiments aimed at uncovering cognitive biases take place in non-dialogic contexts. But from the fact that we are subject to biases in these circumstances, it does not follow that our cognition is generally biased or unreliable. Mercier and Sperber explain: “In our interactionist approach, the normal conditions for the use of reason are social, and more specifically dialogic. Outside of this environment, there is no guarantee that reasoning acts for the benefits of the reasoner” (2017, p. 247). The systematic flaws that psychologists have discovered in our cognition tend to be flaws of *solitary* cognition, many of which are mitigated in dialogic contexts. This view fits nicely with the *ecological* conception of epistemic virtues. Furthermore, some of the dispositions that degrade solitary cognition—such as confirmation bias—serve valuable epistemic ends when we engage in critical deliberations with others; they are components of what Smart (2018, p. 4171) calls *Madevillian intelligence*: “Cognitive and epistemic properties that are typically seen as shortcomings, limitations or biases at the individual level [that] can, on occasion, play a positive functional role in supporting the emergence of intelligent behavior at the collective level”.¹⁴ If, as Mercier and Sperber suggest, dialogic environments are the *normal* conditions for the use of reason, then the literature on cognitive biases, on its own, provides little support for pessimism about the reliability of our intellectual faculties and dispositions. Alfano’s situationism is sustainable only if he can show that the vast majority of our inferential reasoning takes place in solitary conditions where it is more likely to be biased. He has not endeavored to establish this, and it is far from clear that he could. Given the amount of time we spend communicating our beliefs to others (implicitly and explicitly), it is plausible to suppose that we routinely reap the cognitive benefits of group deliberation.

On the other hand, the early findings in the field of group cognition provide grounds for no more than *cautious* optimism. There are, after all, many ways in which group deliberation can *amplify* individual biases. According to the interactionist model, the engine of effective group thinking is *critical feedback*; so in these contexts the majority of individuals must be motivated to, and capable of, criticizing the views under consideration. When there is general agreement within a group, deliberation often leads to overconfidence and group polarization (Baron et al. 1996). When brainstorming takes place in ‘supportive’ atmospheres that don’t tolerate criticism, a variety of biases, including confirmation bias, play a larger role in most participants’ reasoning (Stasser and Titus 2003). While online platforms and social media have made it easier to cultivate these types of uncritical atmospheres, it is generally difficult to carry on deliberations in forums where most participants are either in agreement or silent. Thus, there’s little reason to believe that our thinking generally takes place in social environments that are hostile to the manifestation of individual and collectivist epistemic virtues, though there is a danger of this becoming the case.

In optimal dialogical circumstances, critical feedback must not only be freely and aptly offered, but openly and appropriately *received*: criticisms that fall on deaf ears can have no influence on our beliefs. And we must have the good sense to know how

¹⁴ They are also what Morton (2012) calls, more broadly, *paradoxical virtues*.

to *react* to criticism, i.e., when to update our beliefs in light of it. In other words, we must manifest intellectual virtues *in collective contexts* that we tend not to manifest in solitary contexts, such as: open-mindedness, humility, and reflectiveness. Groups of deliberating individuals must be actively open-minded, even when their members generally are not. Tetlock has found that this can be the case:

But what makes a team more or less actively open-minded? You might think it's the individuals on the team. Put high-AOM people in a team and you'll get a high-AOM team; put lower-AOM people in a team and you'll get a lower-AOM team. Not so, as it turns out. Teams were not merely the sum of their parts. How that group thinks collectively is an emergent property of the group itself, a property of communication patterns among group members, not just the thought processes inside each member. A group of open-minded people who don't care about one another will be less than the sum of its open-minded parts. A group of opinionated people who engage one another in pursuit of the truth will be more than the sum of its opinionated parts (Tetlock and Gardner 2015, pp. 207–208).

There are two ways of interpreting these results. Tetlock's interpretation seems to be that a group's AOM doesn't depend on the AOM of its members; rather, it's an emergent property of deliberative groups with well aligned incentives. The other interpretation is that some individuals can lack AOM when reasoning in isolation, but manifest it when reasoning in well motivated groups. In any case, these types of findings reveal the limits of assessing the reliability of human cognition by appealing only to research that focuses on how individuals reason in isolation.

The empirical fact that individuals manifest a multitude of cognitive biases in experimental contexts by itself does not serve as convincing evidence for ICS. There are several reasons for this. First, the artificial conditions of most experiments don't match the dialogic conditions in which we do much of our reasoning. Second, some of the biases that we manifest in isolation are epistemically beneficial when found throughout groups of deliberating individuals. Finally, many of our biases are drastically mitigated when we engage others in critical dialogue with the common aim of approaching the truth.

5 Conclusion

The heuristics and biases literature is a rich source of epistemological insights, but philosophers should be cautious when drawing general conclusions from this research. It does reveal that human beings suffer from cognitive shortcomings that we're largely unaware of: our thinking about abstract probabilities can be surprisingly inaccurate; expert predictions about complex phenomena regularly fail; our construal of what constitutes cogent evidence is often myopic. But these shortcomings shouldn't blind us to our routine cognitive successes any more than perceptual illusions should cause us to distrust our senses. Moreover, they reveal some of the ways in which we can *improve* our thinking: statistical problems are more likely to be solved when framed in terms of frequencies and addressed by those with statistical training; in many complex domains, we are better off relying on simple statistical rules rather than expert testimony; we

are more likely to consider alternative points of view, and the reasons for them, when in critical dialogue with others. The epistemological moral to be drawn from these results is not that our inferential processes are insufficiently reliable to yield much in the way of knowledge, but that they can be unreliable in conditions that were rare in the ancestral environments in which they evolved. This conclusion supports an ecological view of the epistemic virtues rather than a situationist rejection or reconceptualization of them.

Acknowledgements I am grateful to an audience at the University of Glasgow for their discussion of an earlier version of this paper. I owe a special note of thanks to Jon Marsh and two of this journal's referees for their insightful, detailed, and constructive comments.

References

- Alfano, M. (2012). Extending the situationist challenge to responsibilist virtue epistemology. *Philosophical Quarterly*, 62(247), 223–249.
- Alfano, M. (2014). Extending the situationist challenge to reliabilism about inference. In A. Fairweather (Ed.), *Virtue epistemology naturalized* (pp. 103–122). Dordrecht: Springer.
- Alfano, M. (2017). Epistemic situationism: An extended prolepsis. In A. Fairweather & M. Alfano (Eds.), *Epistemic situationism* (pp. 44–61). Oxford: Oxford University Press.
- Baron, R. S., Hoppe, S. I., Kao, C. F., Brunzman, B., Linneweh, B., & Rogers, D. (1996). Social corroboration and opinion extremity. *Journal of Experimental Social Psychology*, 32, 537–560.
- Bishop, M., & Trout, J. D. (2002). 50 Years of successful predictive modeling should be enough: lessons for philosophy of science. *Philosophy of Science: PSA 2000 Symposium Papers* 69, 69, 197–208.
- Carroll, J., Winer, R., Coates, D., Galegher, J., & Alibrio, J. (1988). Evaluation, diagnosis, and prediction in parole decision-making. *Law and Society Review*, 17, 199–228.
- Carter, J. A., & Pritchard, D. (2017). Cognitive bias, scepticism and understanding. In S. R. Grimm, C. Baumberger, & S. Ammon (Eds.), *Explaining understanding: New perspectives from epistemology and philosophy of science* (pp. 272–292). New York: Routledge.
- Davidson, D. (1986). A coherence theory of truth and knowledge. In E. LePore (Ed.), *Truth and interpretation: Perspectives on the philosophy of Donald Davidson* (pp. 307–319). Oxford: Basil Blackwell.
- Fairweather, A., & Montemayor, C. (2017). *Knowledge, dexterity, and attention: A theory of epistemic agency*. Cambridge: Cambridge University Press.
- Fischhoff, B. (1982). Debiasing. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgement under uncertainty: Heuristics and biases* (pp. 422–444). Cambridge: Cambridge University Press.
- Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the real world*. Oxford: Oxford University Press.
- Gigerenzer, G. (2007). *Gut feelings: The intelligence of the unconscious*. London: Penguin Books.
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102, 684–704.
- Gould, S. J. (1992). *Bully for brontosaurus: Further reflections in natural history*. New York: Penguin Books.
- Hertwig, R., & Gigerenzer, G. (1999). The “conjunction fallacy” revisited: How intelligent inferences look like reasoning error. *Journal of Behavioral Decision Making*, 12, 275–305.
- Hoffrage, U. (2004). Overconfidence. In R. Pohl (Ed.), *Cognitive illusions: A handbook on fallacies and biases in thinking, judgment and memory* (pp. 235–254). Hove: Psychology Press.
- Kahneman, D. (2011). *Thinking fast and slow*. London: Penguin Books.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237–251.
- Kenyon, T., & Beaulac, G. (2014). Critical thinking education and debiasing. *Informal Logic*, 34(4), 341–363.
- Lee, Y., Jussim, L. J., & McCauley, C. R. (1995). *Stereotype accuracy: Toward appreciating group differences*. Washington, D.C.: American Psychological Association.
- Lehman, D. R., Lempert, R. O., & Nisbett, R. (1988). The effects of graduate training on reasoning: Formal discipline and thinking about everyday-life events. *American Psychologist*, 43(6), 431–442.

- Meehl, P. (1954). *Clinical versus statistical prediction: A theoretical analysis and a review of the evidence*. Minneapolis: University of Minnesota Press.
- Mercier, H., & Sperber, D. (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, 34(2), 57–74.
- Mercier, H., & Sperber, D. (2017). *The enigma of reason*. Cambridge: Harvard University Press.
- Morton, A. (2012). *Bounded thinking: Intellectual virtues for limited agents*. Oxford: Oxford University Press.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomena in many guises. *Review of General Psychology*, 2, 175–220.
- Nisbett, R., Fong, G. T., Lehman, D. R., & Cheng, P. W. (1987). Teaching reasoning. *Science*, 238(4827), 625–631.
- Nisbett, R., Krantz, D. H., Jepson, C., & Kunda, Z. (2002). The use of statistical heuristics in everyday inductive reasoning. In T. Gilovich, D. W. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 510–533). Cambridge: Cambridge University Press.
- O'Brien, B. (2009). Prime suspect: An examination of factors that aggravate and counteract confirmation bias in criminal investigations. *Psychology, Public Policy, and Law*, 15(4), 315–334.
- Pinker, S. (1997). *How the mind works*. New York: W. W. Norton & Company.
- Pronin, E., Lin, D., & Ross, L. (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin*, 28, 369–381.
- Saul, J. (2013). Scepticism and implicit bias. *Disputatio*, 5(37), 243–263.
- Sedikides, C., & Gregg, A. P. (2007). They why's the limit: Curtailing self-enhancement with explanatory introspection. *Journal of Personality*, 75, 783–824.
- Siegel, S. (2017). *The rationality of perception*. Oxford: Oxford University Press.
- Smart, P. R. (2018). Mandevillian intelligence. *Synthese*, 195, 4169–4200.
- Sosa, E. (2017). Virtue theory against situationism. In A. Fairweather & M. Alfano (Eds.), *Epistemic situationism* (pp. 116–134). Oxford: Oxford University Press.
- Stasser, G., & Titus, W. (2003). Hidden profiles: A brief history. *Psychological Inquiry*, 14, 204–313.
- Tetlock, P. (2005). *Expert political judgment*. Princeton: Princeton University Press.
- Tetlock, P., & Gardner, D. (2015). *Superforecasting: The art and science of prediction*. Toronto: Signal.
- Turri, J. (2017). Epistemic situationism and cognitive ability. In A. Fairweather & M. Alfano (Eds.), *Epistemic situationism* (pp. 158–167). Oxford: Oxford University Press.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5, 207–232.
- Tversky, A., & Kahneman, D. (2002). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. In T. Gilovich, D. W. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 19–48). Cambridge: Cambridge University Press.
- Wilson, T. D., Centerbar, D. B., & Brekke, N. (2002). Mental contamination and the debiasing problem. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 185–200). Cambridge: Cambridge University Press.