



Knowledge in real-world contexts: not glamorous, but indispensable

Patricia Rich¹

Received: 10 February 2023 / Accepted: 1 November 2023 / Published online: 30 November 2023
© The Author(s) 2023

Abstract

During the past several decades, many epistemologists have argued for and contributed to a paradigm shift according to which knowledge is central to assertion, action, and interaction. This general position stands in sharp contrast to several recently developed accounts regarding specific epistemic contexts. These specific accounts resist applying traditional epistemic norms, including strong knowledge norms, to real-world situations of interest. In particular, I consider recent arguments about the epistemic standards for scientific pronouncements, expert testimony in a political context, and interactive reasoning. I argue, firstly, that knowledge does have a crucial role to play in each case, contrary to appearances. Clarifying the role of knowledge fills gaps in our understanding left open by the existing accounts. Secondly, I show that combining the insights from the knowledge-centric approach and from the more specific accounts provides a new perspective on the open problem of developing an account of knowledge-based decisions. Specifically, I argue that the biggest open problem regarding knowledge-based decisions is not how we respond to uncertainty, but rather how we assess the relevance of our many disparate pieces of knowledge, choosing which to integrate, and how. A strong case emerges for a procedural solution to this part of the problem of how to make knowledge-based decisions.

Keywords Knowledge · Assertion · Action · Reasoning · Science · Deliberation

1 Introduction

1.1 Overview

The view that knowledge is central to assertion, action, and interaction—which I will call KCAA—has become prominent in Epistemology in recent years. KCAA includes two key ideas, as I construe it here:

✉ Patricia Rich
patricia.rich@uni-bayreuth.de

¹ Department of Philosophy, University of Bayreuth, Bayreuth, Germany

K- NORM: Knowledge sets the normative standard for assertion, action, etc.

K- FUNCTION: The concept of knowledge serves a fundamental social function, helping us to identify those people in our society whose information we can appropriately act on.

K- NORM originated from the argument that knowledge is a primitive, the “unexplained explainer” (see Williamson 2000, pg.10), while Hannon (2018) defends K- FUNCTION as part of the project of defining knowledge in terms of its functional role. Setting aside the question of whether knowledge is primitive or how to approach its definition, however, the two ideas are compatible. Many people find both to be attractive or at least very interesting, and we will see that it makes sense to discuss them together here.

This paper is motivated by two problems for KCAA. The first problem is that knowledge does not play the prominent role in many contemporary, contextual accounts of real-life assertion, action, and epistemic interaction that we would expect, if KCAA were correct. For those who find these accounts attractive, this raises the question of why knowledge is not central to them, and whether they are compatible with K- NORM and K- FUNCTION.

The second problem is that, according to K- NORM, actions should be based on knowledge; if that is true, though, then we would like to *characterize* (rational) knowledge-based actions in some detail. This task is the focus of the present special issue as well as a previous topical collection (see Heil et al. 2022), but it is clear that there is a lot more work left to be done. For example, there is disagreement about whether a formal, knowledge-based decision theory is required (Fassio and Gao, 2021), and if so, what that theory should look like (see, e.g., Rich 2021, Goldschmidt 2023, for recent discussion).

This paper attempts to address both of these problems simultaneously by synthesizing several strains of literature—especially relatively new ones—which have so far not played a role in discussions of KCAA. I will argue that not only is KCAA more compatible with the context-specific accounts than it might appear, but that they provide valuable resources for KCAA, and vice versa. These specific accounts can help the development of KCAA especially by confronting it with a variety of real-world problems, while KCAA can inform them especially through its forceful normative claims.

Before outlining the paper’s argument in more detail, I explain why we should pay attention to KCAA in the first place.

1.2 Why think knowledge is central to assertion, action, and interaction?

A growing number of epistemologists take knowledge (as opposed to something else like belief or justification) to be truly central to our epistemic lives, and in particular to assertion, action, and epistemic interaction (the view I refer to as KCAA). One important impetus for KCAA is the observation that it best matches human language use. As Hannon explains:

The word ‘know’ is remarkable for a number of reasons. It is one of the 10 most commonly used verbs in English ... [It] seems to find a comfortable meaning-equivalent in every human language. ... [It is] one of a very small number of words that are allegedly culturally universal. This all suggests that knowledge is deeply important to human life (Hannon, 2018, pg. 1).

We frequently tell each other what we and others know and don’t know, presumably because this is valuable information. With the possible exception of some highly inured epistemologists, we do not even frequently tell each other what we believe or what we find to be justified, let alone our credences or degrees of belief in different propositions. In fact, in many human languages, speaking of “beliefs” or “credences” tends to have a religious connotation which may suggest irrationality or a lack of justification.¹

These observations are especially interesting given that so much of formal epistemology and decision theory has focused on beliefs, probabilistic credences, and so forth. At the very least, the role of knowledge in decision-making has clearly been under-explored, and this deficit in our understanding is now (rightly) being corrected. By way of motivation, I will briefly elaborate on the ideas included in KCAA which are relevant to this paper’s discussion.

According to K- NORM, knowledge sets the normative standard for assertion, belief, and action (Williamson, 2000) (though in this paper we do not care about belief).

Examples like the following show how knowledge norms explain our everyday expectations regarding action and assertion:

Abdul wants to travel from Germany to Denmark for the upcoming holiday weekend. The railways have been damaged by storms, and if the main line isn’t operational, the journey will be more trouble than it’s worth. His friend Bella tells him that everything will be fine with the trains, so he goes. The trip indeed goes smoothly, but Abdul learns from a railway employee that the storm damage has not been fully repaired; the repairs were only paused and some schedules adjusted to accommodate holiday travelers. Furthermore, this backup plan had been announced after his discussion with Bella. Although he is able to enjoy the Danish beaches, Abdul complains to Bella when he returns. He says, “You shouldn’t have told me there wouldn’t be problems with the trains; you did not know that at the time!”

This everyday exchange suggests that agents’ assertions are expected to constitute knowledge; it was not enough that Bella believed the journey would go smoothly, nor even that it turned out to be true. The story is similar regarding the standard for action:

Chuck has the day off and wants to visit his old and frail grandmother. The COVID-19 pandemic is underway. Chuck feels a bit unwell, but he thinks that it is just a cold, and not Covid. He decides not to take a rapid test, which he finds to be painful and inconvenient. When his sister finds out a week later, she is furious: “You should not have visited Grandma when you didn’t know that

¹ So far, all philosophers that I have asked have confirmed this to be true of their non-English native languages.

you only had a cold!” In fact none of them had contracted the virus, but Chuck’s sister remains upset.

Again, this unexceptional exchange supports the thesis that a true belief is not enough to justify an action like Chuck’s; only knowledge that he would not greatly endanger her would have justified his decision to visit his grandmother.

There is of course not complete consensus regarding K- NORM; debate over the components is ongoing.² The purpose here is not to place this specific view beyond doubt, but only to show that there is substantial motivation for it. There are *prima facie* strong reasons to expect knowledge to play a central role in the most enlightening analyses of real-world assertion and action.

K- FUNCTION comes from a research program initiated by Craig (1991) and further developed by Hannon (2018), whose motivation is expressed in the passage quoted above. In his book of the same title, Hannon develops an extended answer to the question “what’s the point of knowledge?” — given that it clearly does something important for us.³

This account of the point of knowledge builds on the increasingly commonplace (and crucial) observation that humans live in (among other things, epistemic) communities; we in fact depend on others for most of our epistemic resources (our information, knowledge, arguments, concepts, theories, etc.) (Hardwig, 1985, 1991). A problem that arises in this setting is that agents need to determine who in their community is a good source of information about a given topic—who they should trust and whose information they should potentially act on—and who is instead relatively ignorant or even likely to misinform. Craig and Hannon’s basic thesis is that the concept of knowledge serves the important function of flagging the “reliable informants” about a given topic. That is—as Hannon spells it out—we label as knowers those whose information we take to be reliable enough for our general purposes within the community (Hannon, 2018, Ch. 2).⁴

As with K- NORM, K- FUNCTION fits well with our use of language, as the following example shows:

A group of out-of-town visitors wants to know whether they should go and attempt a well-known rock climbing route in the area. The weather has been bad recently, and the group wants to check whether the challenging route is reachable and reasonably safe. When they inquire about this with their host, their host says,

² For critical discussion of the knowledge-first paradigm of which K- NORM is a part, see Greenough and Pritchard (2009) and McGlynn (2014). There is a lot of literature on the knowledge norm of action, but see, e.g., Hawthorne and Stanley (2008), Schiffer (2007), Fantl and McGrath (2009), Brown (2008), Anderson (2015), and Gao (2017). Especially relevant to the relationship between knowledge and decisions is the recent topical collection of Synthese on the topic (Heil et al., 2022). There is also some empirical support for knowledge norms for action and assertion (see, e.g., Turri 2013, Turri et al. 2017), but not all empirical tests have supported the same conclusions about assertion (Kneer, 2018; Marsili and Wiegmann, 2021).

³ Craig aims to give a genealogical account, whereas Hannon aims to make their ideas less controversial by framing his approach as “function-first,” hence not necessarily telling us how our concept of knowledge in fact came about. Their core claims are otherwise in agreement.

⁴ Elgin (2020) discusses an important problem with such a practice, namely that the process of flagging reliable informants is likely to exhibit epistemic injustice.

“hmm, I don’t know about the conditions. You should go ask Daji; she climbs there frequently and will know what you can expect.”

By saying essentially “I don’t know, but Daji does,” the host communicates the fact that if the climbing group needs reliable information regarding the likely state of their candidate destination, then Daji is the one in a position to provide this information. While this is just one made-up example, when combined with the observation that we use the word “know” very frequently, it shows how it seems to be extremely important for people to be able to flag knowers and non-knowers, to facilitate reasoning and acting on a dependable information base.

1.3 Knowledge in real-world contexts

I have of course left out many details and discussion regarding KCAA and the particular claims that it involves, but the upshot is that KCAA is an intuitive and well-motivated position. For the purposes of this paper, I won’t argue for anything stronger than this; I’ll simply adopt KCAA for the sake of argument. Adopting KCAA means, in turn, that we should expect the newest and best analyses of real-life epistemic activities to reflect the central role ostensibly played by knowledge. In other words, knowledge should also be central to our *explanations and judgments* of people’s assertions, actions, and epistemic interactions—beyond the simple examples meant to illustrate the point.

The focal problem of this paper, however, is that this is not what we observe. Instead, prominent, plausible, and indeed highly attractive and insightful accounts from contemporary epistemology more broadly may leave knowledge out of the picture, or even directly reject its central role. The purpose of this paper is to analyze this state of affairs, ultimately reconciling KCAA with important developments elsewhere in epistemology, providing more coherence and depth to our developing understanding of how epistemic communities function. This integrated picture will carry important lessons regarding how knowledge is used and transformed into choices, in particular.

The remainder of this paper is organized as follows: Sect. 2 presents three very worthwhile arguments in contemporary epistemology, each of which on its own appears compelling, but which looks more problematic or puzzling from the KCAA perspective. The first argument pertains to the standard for scientific pronouncements (Dang and Bright, 2021), the second to the standard for contributions to political deliberation (Peter, 2021), and the third to the purpose and procedures of reasoning (Mercier and Sperber, 2011, 2017). For each argument, I highlight the apparent tension or contradiction with knowledge-centric tenets, as described above. Section 3 looks more carefully at each of the three arguments presented in Sect. 2, and argues that knowledge plays a crucial role in each setting, contrary to appearances. Section 4 integrates key points from the preceding discussion in order to argue for a different perspective on knowledge-based decisions. Specifically, I argue that the unappreciated challenge in characterizing knowledge-based decisions is that we have a plethora of knowledge and can easily misuse it, and that it is worth pursuing a process-based approach to the problem of determining what knowledge is relevant to a given decision. Section 5 briefly concludes.

2 Three contrary stories

2.1 Science

The first argument I discuss is presented in Dang and Bright (2021). The authors express their conclusion starkly as follows: “We argue that the main results of scientific papers may appropriately be published even if they are false, unjustified, and not believed to be true or justified by their author” (Dang and Bright, 2021, pg. 1). In the terms of KCAA, they argue that scientific pronouncements (of the kind intended for consumption by other scientists; we turn to science’s contributions to public discourse later) need not be knowledge, nor anything close to it. While Dang and Bright explicitly do not commit to the premise that the scientific pronouncements in question qualify as assertions, it certainly appears that their conclusion contradicts the knowledge norm of assertion. Whether this is really the case will be considered later.

For now, we just need to understand the main reasons for their conclusion, so that we can see that they are compelling and very much in line with contemporary thinking in epistemology and the philosophy of science. Dang and Bright make two key points.

First, the scientific community makes progress partly due to a division of labor between individual scientists, who pursue different methods, theories, and projects in general, and who assume different positions regarding unsettled questions (see Solomon 1992; Hardwig 1991; Kitcher 1990, 1995, for seminal works on the division of cognitive labor). Specifically to the point that (at least some) scientists cannot know the conclusions they put forth, a division of labor over competing research hypotheses, theories, or even paradigms is now seen to be essential: the community cannot know in advance of years or decades of research which theories are best and which projects are most worthwhile, so it is good that it can cover its bases by having various possibilities explored (for a detailed historical example, see Zollman (2010)). Necessarily, then, some scientists will spend their time developing and defending conclusions which turn out to be wrong, and such endings will often not be terribly surprising.

Second, Dang and Bright appeal to pessimistic meta-induction arguments which suggest that “almost all scientific public avowals turn out to be false” (2021, pg. 10); current scientists have no particular reason to think that history will eventually show their conclusions to be true, when their predecessors’ conclusions have been overturned.

These two points are important, and I will not raise objections to them. In particular, it seems pretty clear that science can only make progress (at least at a reasonable speed) if scientists are willing to put out conclusions of which they are not certain. A true conclusion may need to be accepted and worked with for a while before it can be confirmed. Some conclusions which look well supported will turn out to be false. Scientists themselves may recognize that the evidence is not completely sufficient to support a radical new claim, but nonetheless argue for and share it because it would be extremely important if it turned out to be true, because it provides an important counterpoint to the existing alternative, or for other similarly sensible reasons. But all of this means that scientists must state conclusions that they do not know; a knowledge norm here would make science much too conservative and hence block progress.

Hence, Dang and Bright reject K- NORM. SCIENTISTS REGULARLY ASSERT CONCLUSIONS THAT THEY DO NOT KNOW, AND IT IS GOOD THAT THEY DO SO. FURTHERMORE, SCIENTISTS ARE WELL AWARE THAT THEY DON'T KNOW THEIR CONCLUSIONS; DANG AND BRIGHT PROVIDE SOME HISTORICAL EVIDENCE THAT SCIENTISTS SOMETIMES DO NOT EVEN *believe* THEIR OWN CONCLUSIONS.

I do not think this historical evidence is conclusive, but in any event, based on Dang and Bright's argument, there seem to be very good reasons to think that a strong knowledge norm of assertion does not reflect the de facto standards within the scientific community, and furthermore that instituting such a norm would interfere with important mechanisms of progress by increasing caution and conformity.

2.2 Politics

The above argument pertains to scientific assertions which are meant to stay within the scientific community, and which primarily influence "epistemic actions" such as the research activities of other scientists. These assertions then arguably have relatively low stakes; certainly the stakes are much higher when a scientist's assertion will be used to determine actions which more directly influence people's well-being. This is the case, for example, when a scientist's expert testimony influences the public discourse and ultimately how we respond to practical problems such as pandemics, climate change, and war. K- NORM clearly indicates that these decisions should be based on knowledge.

Moreover, it has been argued that as the stakes rise—as when many human lives are on the line—the standards for the basis of the decision also rise (Williamson 2000, pg. 99, Fantl and McGrath 2002). Mere knowledge might not be enough; higher stakes may require higher-order knowledge (Schulz, 2017) or stronger knowledge (Schulz, 2021) as a basis for action. Alternatively, whether an agent knows a proposition may depend on the context (Hawthorne, 2004; Stanley, 2005; Fantl and McGrath, 2009); in high-stakes contexts, the agent may need to be in a stronger epistemic position in order to know a proposition that would support their action. There is quite a wide variety of views regarding the relationship between stakes, epistemic status, and action, but one way or another, people seem to agree that when facing a consequential choice, as political choices often are, the epistemic standards are *at least* as high as they are for mundane choice problems, and likely higher.

As part of her work on political legitimacy, Peter (2021) considers the hypothesis that it is indeed knowledge which sets the standard for the admission of a piece of testimony into political deliberation. Very much in line with knowledge-centric thinking, she points out that a knowledge norm avoids the *recklessness problem* by ensuring that testimony is well supported by evidence. Contrary to K- NORM, however, Peter comes to the conclusion that a knowledge norm for political deliberation is too strong. Let's examine her reasons next.⁵

Essentially, Peter's objection to the knowledge norm in this case echoes a familiar concern about the knowledge norm of action. The concern is that we are simply faced with too much uncertainty; we know too little for our knowledge to form a sufficient

⁵ For additional discussion of the norms for scientific advising, see, e.g., Blessohl and Sarikaya (2022).

basis for action. We don't know how likely it is that a person experiences long-term consequences from a viral infection, nor do we know the extent of heat waves to be expected in a particular place and time-frame given a possible course of climate policy. Surely, then, our choices need to select these unknowns, and not just the knowns.

Peter expresses her argument thusly:

There are too many uncertainties in the political context to allow for a meaningful restriction of well-ordered political deliberation to what is known. Even the best scientific advice - the kind of advice we would want political decisions to be based on - tends not to consist of what is known but reflects a temporary broad consensus among scientists about what is justifiably believed in this regard. More generally, we typically neither know all relevant details of the situation we're in nor what the future holds, but political decisions need to be made and assessed anyway (2021, pg. 401).

As has been extensively discussed in the philosophy of science literature (Frigg and Hartmann, 2020), much scientific work involves building models to understand phenomena (such as the spread of infections or changes in ocean temperatures) and to make projections. Importantly, modeling assumptions must be made, and it can happen that these assumptions turn out to be false in such a way that a model's results are problematically influenced. We therefore cannot usually say that the projection provided by a particular model can be known to accurately indicate what will really happen. Nonetheless, the results of our best scientific models should surely be admitted into deliberation.⁶

2.3 Reasoning

The final argument to be discussed is developed by Mercier and Sperber (2010, 2011, 2017), in support of their argumentative theory of reasoning. This argument is different from the preceding two in that a central role for knowledge is not discussed and rejected; instead, knowledge is conspicuously absent from their picture. It should be noted that while some parts of Mercier and Sperber's account are quite controversial, the most controversial aspects (such as the modular mind thesis) are detachable from the argument of interest here, and at any rate the account as a whole is important enough to merit serious attention.⁷

Mercier and Sperber problematize the apparent status of human reason as a "flawed superpower": on the one hand humans seem to owe their great success as a species to our ability to reason, but on the other hand there is a mountain of literature on all of the ways in which reason seems not to work correctly. The authors argue that this analysis should make us suspicious, since human reasoning is an adaptation that we should expect to be tailored to our needs. Therefore, if reason appears not to function correctly, we should re-evaluate what its function really is. Their core conclusion is

⁶ In fact, Peter points to the widespread use of models in science to reject a *truth* norm for political deliberation.

⁷ For a good idea of the critical reception, see the many commentaries following their expository article (Mercier and Sperber, 2011).

that humans “produce reasons in order to justify our thoughts and actions to others and to produce arguments to convince others to think and act as we suggest” (Mercier and Sperber 2017, pg. 7).

Critically, then, reason evolved for use in an argumentative, social context. This allows us to make sense, for example, of what Mercier and Sperber refer to as “myside bias,” the phenomenon of people being very attentive to reasons in support of views that they already hold or which work in their favor, and relatively good at arguing for positions they endorse, but less attentive and capable when it comes to contrary reasons and argumentation. When people reason with others, myside bias produces a division of labor and allows people to be shown by others where their reasoning goes awry. When people reason alone—outside of the context in which reason evolved to be used—myside bias can lead people to become more and more committed to non-meritorious conclusions.

So far, knowledge simply doesn’t appear in this account. The account also seems very much at odds with KCAA. We can see how both K- NORM and K- FUNCTION seem to fit poorly with the account by considering the two components of interactive reasoning, as Mercier and Sperber describe them.

First, agents offer their own reasons. Mercier and Sperber characterize people as being very lax in giving reasons. When an agent asserts something in an argumentative context, they are not cautious to only put forward well-supported premises; the fact that a premise supports their position is the key criterion. This doesn’t seem to fit with K- NORM, though.

Second, agents (consider whether to) accept the reasons offered by someone else. In contrast to agents’ lax standards for their own reasons, Mercier and Sperber characterize people as vigilant here: an agent carefully considers whether to accept the reason proffered by the other, because although they want the benefits of good information they do not want to be open to manipulation. This line of reasoning thus shares K- FUNCTION’s premise that getting information from others is extremely important in human societies. There is also the question of whom to get what information from. Yet Craig and Hannon argue that we solve this problem by flagging the reliable informants as “knowers,” whereas Mercier and Sperber’s picture looks more competitive and even anarchic, with individuals left to decide for themselves whether to accept each premise and obliged to be skeptical. Hence, the account also seems to be at odds with K- FUNCTION.

Interestingly, Mercier and Sperber’s account covers all of human reasoning, including reasoning in scientific and political contexts. Descriptively, they provide many examples which seem to demonstrate that human reasoning works as they claim. We need to worry, then, that Dang and Bright’s and Peter’s scientists will also interact in this self-oriented manner. From a theoretical perspective, this reveals that those accounts also indirectly put pressure on K- FUNCTION; in both cases, the status of reliable informant (which we presume scientists to have with respect to their own work) seems to be disconnected from the status of a knower. From a practical perspective, a worry grows: Knowledge norms and identifying knowers are supposed to ensure that the information passed around and acted upon are reliable and of high quality. But then where have these safety devices gone, and shouldn’t we worry if they go missing?

3 Reconciliation: hidden knowledge

3.1 Science

In this section, I will re-examine each of the three arguments discussed in Sect. 2, and explain how the important points of these arguments can be reconciled with KCAA. The norms for assertions in science are discussed first. It should be noted, though, that the domains of science and political deliberation are connected, especially insofar as we are especially concerned with the contribution of scientists to political discourse. Furthermore, argumentation takes place in both of these domains (and elsewhere). Hence, my responses to the three arguments are not intended to be kept separate, and indeed they will merge to an extent.

Dang and Bright make several important points which must be accommodated: scientists need to divide labor to explore various alternatives, including those with initially dubious merits; if scientists are too conservative, then progress will be hampered; and this means that scientists must sometimes be willing to go out on a limb, and wager claims beyond what they can know. As with everything else in life, however, moderation is key. If the standards for assertion in science are too low, then progress will also be hampered. After all, we still think that scientists should tend to work on the projects which are truly best justified and most promising. And how could the scientific community be expected to identify and eventually converge on worthwhile theories, projects, etc. if the journals and other communication channels were swamped with junk that nobody really took seriously?

The key to the middle ground here, which actual science arguably occupies, lies in the distinction between different types of scientific claims which is implicit in Dang and Bright's argument. Dang and Bright do not state that scientists in general need not know anything they say; instead, they argue that scientists' "main conclusions" in particular need not be known. As I interpret their discussion, we can think of these main conclusions as the general takeaways that scientists identify based on pieces of research, and which often appear in the titles, abstracts, introductions, and conclusions of their articles.

Critically, however, these "main conclusions" represent only a tiny fraction of the statements appearing in a scientific article or presentation. I argue, furthermore, that they are far from the most important statements. The backbone of a piece of scientific research is instead the numerous more specific claims that stand behind the general conclusions. These more specific statements include descriptions of the state of the literature and the open problems to be addressed, reports of the evidence used or gathered, reports of research protocols, descriptions of the content of the theories used, descriptions of the models employed and their rationales, and claims about the consequences of the models, i.e., what is true within a model or modeling framework or what is observed when a simulation model is run.

These more specific assertions are what ultimately move our collective knowledge base forward, and what other scientists carefully inspect, combine, and build on. Therefore, I argue, these specific assertions are expected to be known, in accordance

with K- NORM.⁸ This makes sense: We know from history that we really need some scientists to go out on a limb and argue that theory A is better than theory B (as a main conclusion), even when the total evidence may make this unclear or even unlikely. An argument that we similarly need scientists to go out on a limb and declare that they performed experiment C when they really performed D (a nuts-and-bolts claim), or that their simulations revealed some outcome E when they revealed not E but F, is not forthcoming. Indeed, we would call this fraud.

It's true that choices and interpretation are needed, to an extent, throughout the scientific process, and it may not always be completely clear whether something should be seen as a subjective conclusion or a (relatively) objective claim about the world. Nonetheless, actual practice seems to reflect my distinction between the normative standards in place for general conclusions and for the nuts-and-bolts of research. Generally speaking, scientists are permitted to speculate a bit when arguing for the general consequences of their work; they may even be expected to do so, as the practice may have practical benefits, as discussed below. Such speculation is not held to the same standards as the rest of the work; we might even say that these speculations are not treated as assertions. A scientist who does solid work but uses it to defend a position which turns out to be wrong is not in any kind of trouble with the community.

In contrast, scientists do seem to be held to something like a knowledge norm when it comes to the more specific claims which collectively form the body of their research. A scientist who reports evidence they don't believe to be true is taken to act dishonestly and fraudulently. A scientist who does not check their work carefully and so reports unjustified "results" from modeling work is taken to task and seen to lack the necessary integrity. It is similarly unacceptable to misrepresent the state of research on a topic. Scientists who are found to make non-knowledge assertions of these kinds are held responsible for violating community standards.⁹ We see this borne out in the standards for retractions or corrigenda of published scientific work: It is common for journals to publish a corrigendum if an error is found in a paper's formal argumentation, for example if there is a problem with a proof. In extreme cases of false or misleading nuts-and-bolts research, as when experimental practices are not accurately reported or data use is grossly inappropriate, work can be retracted (Wikipedia contributors, 2023). Hardly anyone would find it necessary to formally correct or retract a methodologically solid paper, on the sole grounds that the general conclusions turned out to be wrong (e.g., the work was used to defend a theory that was later agreed to be wrong).¹⁰

⁸ As noted above, I adopt KCAA in this paper for the sake of argument, and aim to show that it is compatible with arguments like Dang and Bright's. This means that I argue that these specific assertions should be known *in contrast with* being potentially unknown, unjustified, and unbelievably, as Dang and Bright claim for the "main conclusions." I cannot provide an argument here that the nuts-and-bolts claims of scientific papers should satisfy a knowledge norm *in contrast with* a closer alternative such as a certainty (Beddor, 2020) or surety norm (Goodman and Holguín, 2022). See also Lackey (2007), McKinnon (2013), and Simion (2016).

⁹ We're ignoring here the fact that due to power imbalances, biases, and so forth, individuals may not be punished equally for violating scientific standards. Standards may also be interpreted and implemented in a biased manner. The in-principle standards themselves are what matters here.

¹⁰ Thanks to an anonymous referee for pointing to the conditions for retraction as supporting my argument.

It is then an interesting question why there would be a different norm for “main conclusions” than for the body of research. A plausible answer is that the two kinds of statements have different functions, which support different standards. Specific claims about why, how, and what research was done provide the concrete basis for further work, providing concrete information about the phenomena under study given our current understanding. If this cannot be relied on, then it is difficult to see how progress could be possible.¹¹ In contrast, general conclusions seem to have a more psychological function. They may help to catch the interest of possible interested parties, unleashing creativity or providing inspiration; efficiently communicate what an article or talk is about, possibly making the research easier to digest and understand; signal what sub-community the scientists behind it belong to, helping others to see what other work this work might connect to, and how; help other scientists to understand which options are considered live and what the general perspective of the community on general questions is; and ultimately motivate other scientists to pursue a particular project. Science without general conclusions could be too boring to attract much effort, or simply too hard to approach for cognitively limited beings. In effect, main conclusions may have a heuristic function, and it is plausible that they couldn’t serve their purpose (e.g., rendering research more accessible) if they were held to a high epistemic standard.

Darwin seems to have had a similar view on the different purposes and standards for different kinds of claims within science. Smaldino (2022) paraphrases Darwin as follows:

What [Darwin] is saying ... is that we shouldn’t worry too much about false theories, because academics are competitive and love to take each other down a peg by demonstrating logical inconsistencies in one another’s theories. ... However, any coherent explanation must rely on a firm foundation of facts. If our facts are false, we end up wasting our time arguing about how best to explain something that isn’t even true (2022, pg. 20).

While our conception of science has grown beyond a simple distinction between empirical facts and logically constructed theories, Darwin’s point remains fundamentally sound and supports my claim that there are different standards for different kinds of scientific claims. Implicitly, Darwin also acknowledges that this is because the different kinds of claims are used in different ways. While theory building is essential to science, I have suggested that the “main conclusions”—as related to, though distinct from, theories—serve a heuristic and communicative purpose rather than a deep scientific one.

Where does this leave us with regard to K- NORM? We already saw that K- NORM is not challenged by scientists’ apparent standards for the nuts-and-bolts claims that comprise most of their published work; Dang and Bright haven’t given us an argument that knowledge is not the relevant norm for those assertions. They have given us a convincing argument that knowledge is not the norm governing *all* scientific pronouncements,

¹¹ This is not to deny that specific scientific claims are theory-laden or relative to a paradigm that may not be in place forever. I take for granted that we can have (scientific) knowledge, even if our beliefs are not eternally and universally valid.

in particular because it doesn't govern "main conclusions." We could say, then, that K- NORM is not universally valid; it has at least some exceptions. A more elegant solution, however, would be to retain K- NORM, but not categorize the "main conclusions" as assertions.¹² This would follow anyhow if the knowledge norm were a *constitutive* norm for assertion, as (Williamson, 2000, Ch. 11) argues; then we would reason from the fact that those particular pronouncements are not (expected to be) known to the conclusion that they are not assertions. Dang and Bright themselves are less interested in the question of whether these claims count as assertions (which they do not insist on) than with the standards that apply to them. On that point we are in agreement, and in any event, their argument is not as dangerous for K- NORM as it at first appeared.

3.2 Politics

Political deliberation presents an importantly different context from internal scientific argumentation. Peter recognizes this, writing, "I take it as a given that certain speculative claims can be validly asserted in a scientific context, but not necessarily in other contexts, e.g., in a context of policy-planning" (Peter, 2021, pg. 398). Since political deliberation is typically aimed directly at decision-making, norms for both assertion and action are relevant, and the stakes are (in general) correspondingly higher since errors are more likely to have serious short-term and long-term consequences for many peoples' lives. At the same time, however, the kind of heuristic function served by general conclusions could be even more important than in the scientific context. While scientists may use general conclusions to orient themselves but still necessarily dig deeper into the details of relevant research work, there is plausibly less capacity for communicating the details in a political context; parties to the deliberation will often lack the necessary expertise to engage with the full details of scientific research and the state of knowledge, and will certainly lack the time to do so regarding every relevant research domain. Hence, it appears simultaneously more valuable to apply a knowledge norm for admission into political deliberation, and less practical to restrict contributions in this way.

In light of the former consideration, I will argue in defense of knowledge as the epistemic standard for contributions to political deliberation (hence defending K-NORM).¹³ I will focus on contributions from scientists providing expert testimony, since Peter's relevant arguments are based on considerations about scientific testimony. Even relatively minor political decisions (such as slight changes to the regulations regarding agricultural subsidies) can end up having significant consequences for the individuals whose lives are directly affected. The most active political deliberation,

¹² Other solutions may be possible. For example, a referee points out that constructive empiricists might say that "main conclusions" aim at empirical adequacy rather than truth (see Van Fraassen 1980). In that case, plausibly, K- NORM might apply to the belief that the conclusion one articulates is empirically adequate, rather than to the conclusion itself. This would have the benefit of providing a clear standard for main conclusions.

¹³ Here again, I defend knowledge over much weaker epistemic standards, rather than knowledge over close alternatives such as a certainty norm.

moreover, tends to concern major decisions like how to respond to problems such as climate change; such decisions can be expected to have significant consequences for most people. Thus, lowering the epistemic standard for contributions must be seen as potentially very costly in terms of the expected decision quality, and cannot be taken lightly. As noted above, several authors have defended the very intuitive claim that as the stakes rise, the epistemic standard for the basis of action becomes correspondingly higher. From this perspective, making (especially major) political decisions on the basis of less sound, unknown premises is unacceptably irresponsible.

Peter's claim, however, was not that it is undesirable that contributions to deliberation constitute knowledge, but that such a standard cannot be met in practice. The task, then, is to show that the knowledge standard is not in fact too high. The crucial claim in support of this is that insisting that we act on knowledge does not—contrary to the general perception—mean ignoring the uncertainty and complexity that real-world agents face; instead, it means taking it seriously, by not over-simplifying.

To see this, let us compare different kinds of contributions that scientists might make to political discourse, in light of specific concerns one might have about insisting that these contributions constitute knowledge. First, take Peter's concern (highlighted above) that scientific testimony often "reflects a temporary broad consensus... about what is justifiably believed." Now, it could be the case that if the scientific community generally agrees that we are well justified in believing x , then the community should simply be treated as knowing x (and is then entitled to act as if x were true; let's accept that knowledge is fallible and we must not all be skeptics). If not, then there are presumably specific reasons why the community is hesitant to treat x as knowledge. Perhaps x is the best hypothesis at the time, but it is clear that the evidential basis is not all that strong. Perhaps x would be the case if the present situation were like the past ones in all relevant respects, but it's still not clear whether this case might not be importantly different. If the community is not held to a knowledge norm, then perhaps they are allowed to simply assert " x ." If, instead, they are held to a knowledge norm, then they don't simply keep quiet; they assert something like "we largely agree that x is most likely to be true according to everything we know (which is quite a lot!)." However, it is always possible that this case turns out to be special, and so we will keep an eye on the situation to see whether y doesn't turn out to be the case instead." The latter contribution both respects the knowledge norm and more honestly reflects the real uncertainty involved in the situation instead of sweeping it under the rug. It arguably captures how scientists often communicate with the public.

Why would we prefer the former kind of contribution to the latter? The main reason that I can see would be a concern that the more nuanced contribution is somehow too complicated to be integrated into a decision-making process. More information is harder to process, and perhaps more likely to be ignored rather than dealt with. I acknowledge that the audience may prefer simple contributions, but deny that this makes simpler, non-knowledge contributions preferable to more nuanced, known ones; if y is really a live possibility for the future, then it is good that society be aware of this and have the option of either ignoring it for now or incorporating it into their choices. How cautious society should be in this regard is arguably a question that society should make *via* deliberation, and not one that scientists should make for them by hiding

existing uncertainty.¹⁴ Some of the public may be overwhelmed by the uncertainty, but then it is the job of their representatives to deal with it, to communicate what is important, and to act in their constituents' interests.

Second, let's take Peter's concern that much scientific theorizing proceeds via modeling. Modeling results (a) are dependent on underlying modeling assumptions, and (b) typically cannot provide precise and reliable information about what will occur in the future. We can again compare different kinds of contributions that scientists might make, and consider whether the knowledge-constituting contributions are better. If (a) is the main concern, would we prefer that scientists tell us simply "x," or that "our best models all suggest that x, this result is fairly robust to different modeling assumptions, and so x is strongly supported although not absolutely certain"? If the main concern is (b), would we prefer "sea levels should rise by x centimeters during the next ten years" or "sea levels will almost surely rise between y and z centimeters in the next ten years, with x centimeters being the most likely increase"? As before, the former contributions are simpler, easier to generate and easier to digest, but do not count as knowledge. The latter contributions (by hypothesis) do count as knowledge, or at least represent the kinds of statements scientists would make if they aimed to satisfy a knowledge norm. These contributions can therefore be trusted and relied on in a way that the former cannot. They do not ignore real uncertainty and complexity. As a result, they provide a better foundation for decision-making. If x is very likely but not certain, then we may not want to lock ourselves into a course of action that will be catastrophic if x turns out not to be the case. If scientists know that sea levels may well rise more than x centimeters, then surely the public should make an informed choice about whether to prepare for the possibility of a greater rise than expected.¹⁵

Still, there is a large literature questioning how it is possible to learn about the real world via idealized, abstract models (see Reiss 2012, for a useful systematization). A main source of skepticism about modeling in science is essentially (a) above: given that models rely on assumptions which may be false—as well as assumptions we *know* to be false—it is hard to see how they could deliver knowledge about the world.

This is an important issue that should be taken seriously, but not, I would argue, by lowering the epistemic standard for important decisions. If we really think that scientific models cannot produce knowledge of a certain type, then conclusions based on these models should not be passed along as if they were a solid basis for policy when they are not. In this case, it would be important to determine what kinds of knowledge can actually be produced, and how, even if this means that conclusions cannot be drawn as quickly or presented with as much certainty. If we think that there

¹⁴ Obviously, scientists do not further debates about things like pandemics by pointing out far-fetched possibilities, such as the apparent non-zero probability of "bizarre 'quantum' events" like a dropped marble passing through the floor to land beneath the house (Hawthorne and Lasonen-Aarnio 2009, pg. 94). However, epistemologists tend to agree (via one argument or another) that we can have (and communicate) everyday knowledge despite the existence of various undermining skeptical hypotheses which we strictly speaking cannot rule out. We should apply our ordinary pragmatic standards for knowledge and assertion to the scientific communication under consideration.

¹⁵ In this vein, Samuels (1989) presents a nice argument that scientists (specifically, economists) should be much more "diffident and restrained" in giving policy recommendations.

are particular ways that models can be used to produce particular types of knowledge, then scientists should use models in those ways and then pass along the knowledge which can actually be attained. There are proposals along these lines in the literature. For example, Alexandrova and Northcott have argued that models can be used to develop causal hypotheses, which when tested and verified can constitute knowledge (2009). They make the forceful point that when it really matters, it would be foolish to simply trust a formal model with no genuine empirical confirmation, and we often don't (or do so to our peril). Grüne-Yanoff (2009) has argued that models on their own can prove *modal* hypotheses, i.e., show us what is possible or what necessarily follows from what. Acting on the basis of such relatively modest model implications is intuitively sensible and responsible. Much of the discussion about models is motivated by concerns that much too much stock is put into model results in themselves. Given this, there are better reasons for retaining a high standard for contributions from scientists for policy-making purposes than for relaxing it to accommodate uncertain modeling results.

As a closing observation, I would point out that real-world scientists themselves seem to prefer a higher standard for contributions that leaves room for the uncertainty which we really face. It has become a topic of public discussion, especially during the COVID-19 pandemic, that scientists often try to communicate both their best hypotheses and the uncertainty surrounding them to the public, but their qualifications and nuances are often lost or hidden behind attention-grabbing headlines.¹⁶ This is frustrating to scientists who try to communicate important understanding to the public, and problematic for policy-makers because catchy slogans and other oversimplifications are not in fact a good basis for important choices. It could be the case that people no longer have the time, energy, or attention spans to properly attend to the subtleties of our state of knowledge on important issues. I would argue, however, that the proper normative response is to develop decision-making processes that facilitate proper engagement with the relevant knowledge base, and not to base our decisions on easily digestible, but superficial and misleading, grounds. I return to this point in Sect. 4.

3.3 Reasoning

Unlike Dang and Bright's and Peter's arguments, Mercier and Sperber do not directly argue against KCAA; instead, they seem to simply tell a very different kind of story, as explained above. I will close this section by arguing that this appearance is deceiving. Specifically, I will offer three proposals for how knowledge may be playing an important role in an argumentative account of reasoning, even if this role has not been the part of the story Mercier and Sperber sought to highlight.

¹⁶ Nuance can also potentially be exploited by malevolent actors to spread conspiracy theories and the like, or more generally, interpreted by some of the public to mean that scientists don't know what's going on, and anything could be true. I think that this is a serious practical problem, but one that should be dealt with through science education, improving digital competence and robustness to fake news, and so forth. Basing decisions on over-simplified versions of our best science seems to me to be a problematic overreaction. Thanks to *suppressed for review* for pressing this issue.

The first proposal is that in an argumentative context in which agents exchange reasons, the agents are in fact expected to know their reasons (i.e., K-NORM holds). Agents are characterized as lax in giving reasons, but this need not mean that the reasons aren't known; it can mean that the agent does not apply a high standard for the relevance of the provided reason or for its ability to justify the conclusion the agent is trying to defend. This proposal seems to be perfectly compatible with the argumentative theory of reasoning.

The ease of combining a knowledge norm of assertion with the argumentative theory can be illustrated by considering a realistic argumentative exchange which is very much typical of Mercier and Sperber's examples, but with a knowledge-centric twist:

Suppose Esther and Frana are trying to decide what to do on Saturday. We can well imagine the following exchange taking place:

Esther: We haven't been to the movies in so long! Let's go see that new movie people keep recommending.

Frana: But the sun is also out for the first time in ages! Why not take advantage and have a nice picnic?

Esther: We will have good weather for a picnic next weekend, too. I'm really itching to see this film.

Frana: You don't know that the weather will stay like this! You probably haven't checked the forecast for tomorrow, let alone a week from now.

The agents exchange self-serving reasons, trying to convince the other to go along with their preferred plan rather than objectively considering the best course of action. They serve their interests, however, by highlighting some pieces of information rather than others—we can even imagine that their friend Gengis has a birthday the next day, and it would make a lot of sense to prepare the food for the party, but neither of them mentions this because they don't feel like doing it. An agent who asserts something that they clearly don't know will still be challenged, as we see when Esther tries to get away with postponing the picnic on the invented grounds that next week will work just as well. The knowledge norm of assertion suggests that pointing out that a proffered reason is unknown should be sufficient grounds for rejecting it.

One of the few places in *The Enigma of Reason* where knowledge plays a substantial role is Chapter 16, where Mercier and Sperber discuss the universality of reason. The discussion here supports the hypothesis that reasons are held to a knowledge standard. The authors report on experiments in which completely unschooled individuals were provided with simple logic problems, including the following (from experiments by Luria):

In the Far North, where there is snow, all bears are white. Novaya Zemlya is in the Far North. What color are bears there? (Mercier and Sperber, 2017, pg. 278)

Analyzing the transcripts from Luria's experiments, Mercier and Sperber argue that the unschooled individuals understand what the answer to this problem is supposed to be, but refuse to give this answer, essentially, because they are unfamiliar or not on board with the idea of claiming a conclusion which they wouldn't claim in the

real world. In other words, the participants seem not to be willing to assert the logical conclusion because they do not *know* it to be true. This interpretation is supported by the transcript excerpts that Mercier and Sperber provide (2017, pgs. 280–281):

Young Uzbek: From your words it means that bears there are white.

Older man: What the cock knows how to do, he does. What I know, I say, and nothing beyond that!

If we accept that reasoning has a fundamentally argumentative function and is meant to be used in a social context, we still require an account of which reasons are permitted and which are not.¹⁷ This example further supports the idea that a knowledge norm of assertion should be part of this account.

My second proposal is that Mercier and Sperber underestimate the importance of cooperative argumentation (a point also made by Dutilh Novaes (2018)). Here, it is important to note that while they argue for an argumentative function of reasoning, it remains unclear what the function of argumentation is. At some points, the authors provide examples where agents seem to improve with respect to a commonly shared goal through argumentation, as when a group of reasoners is able to find the answer to a logic problem that none could identify alone, or when scientists improve their theories by discussing them with others. Often, however, the argumentative theory is expressed in a more competitive way, with agents seeking to convince others to adopt their own views. This has been a common theme in discussion of Mercier and Sperber's work (2011, Peer Commentary). Quite plausibly, argumentation is an important tool in both kinds of cases; we can argue together to figure out, for example, which insurance plan best meets our needs or how to go about fixing a broken appliance, but we can also argue to determine which of us gets to eat the last piece of pizza or who has to clean up while the other relaxes. One context is cooperative and the other competitive, but in both cases we use argumentation to make a decision (plausibly, again, with knowledge as input).

Mercier and Sperber, however, say that people with common interests and mutual trust would have little or no use for justification and arguments (2017, pg. 334). This makes sense if we imagine argumentation being used to determine who gets their way or enforces their opinion, and involving attempts at manipulating others through reasoning. If we have a common goal and mutual trust, then indeed there is no point in trying to manipulate one another. Argumentation, however, is still useful insofar as the mutually-beneficial course of action or the best viewpoint is not completely obvious. This is clearly the case for many interesting real-world contexts. If our common aim is to try to identify a yet-unknown truth, then argumentation is a valuable tool for reasons Mercier and Sperber themselves discuss at length (e.g., we can point out flaws in one another's reasoning and take advantage of different perspectives or evidence bases). The same applies if we need to identify a good course of action. If we consider cases such as trying to anticipate the consequences of climate change or plan a response to it, then it is easy to see that argumentation is not necessary *only*

¹⁷ There is no need to start from scratch here; see Dutilh Novaes (2022) for an overview of the literature on argumentation, and also Norman (1998) for an example of an account of proper reason-giving.

because people have partially different interests and may not always trust one another. Even setting these complications aside, argumentation is required because it is simply very hard to anticipate the specific consequences of possible courses of action and to plan accordingly.

Exchanging reasons is basically the only method we have to try to collect a very diverse array of complicated information and goals, and to transform this into a unified perspective or plan. Critically, then, we do want the parties to the discussion to be exchanging items of knowledge or something similarly reliable (as argued previously regarding political deliberation). When there are convergent interests, putting forth shakier premises—to say nothing of deliberate misinformation—can generally be expected to do more harm than good. The problems to be solved are still hard—and argumentation is needed—because the *relevance* of each item of knowledge, and how it fits into the bigger picture, need to be determined. So, in line with the first proposal above, agents are expected to know their reasons, and in a cooperative reasoning context argumentation serves to sift and integrate knowledge in an acceptable way. That is, in a cooperative reasoning context, both K- NORM and K- FUNCTION are well motivated.

My third proposal is that K- FUNCTION—the role of reliable informants (knowers)—can and should be integrated into an argumentative theory of reasoning. Reliable informants are arguably a bit of a double-edged sword, from this perspective (a point which I reiterate in Sect. 4.4 below). On the one hand, more and better knowledge can enable better decision-making. Especially in more cooperative contexts, the group will be better off if this knowledge can be exploited. This is why expert advice tends to be sought in important choice situations. As previously observed, however, knowledge on its own does not suffice for choice. For one, the relevance question has to be settled: which knowledge do we act on? Secondly, choices are the product of both knowledge and values, and the integration of knowledge must reflect goals and values which are typically not just a matter of expert testimony.

On the other hand, therefore, people who know more have more argumentative power and more ability to steer the group towards their preferred conclusion or course of action. This is not an unmitigated boon given that they (a) are subject to the same cognitive biases as everyone else and reason best in an argumentative context, per the theory, and (b) typically have domain-restricted status as a reliable informant about some, but not all, topics (for example, the expert on biodiversity need not be any kind of expert when it comes to questions of moral value). The worry here is similar to one discussed by Mercier and Sperber; they provide colorful historical examples to show that better reasoners are also better at finding ways to justify poor and even abhorrent conclusions (see esp. 2017, Ch. 13).

It is important, then, to integrate the distinction between knowers and non-knowers (emphasized by Hannon) with the argumentative theory. The argumentative context is supposed to provide a corrective to myside bias because reasoning can be challenged. Yet challenging someone else's reasoning will be harder, or even impossible, when that person knows much more about the topic at hand. Knowledge asymmetries, then, can amount to power asymmetries, giving the knower an argumentative advantage that can work to the detriment of others, or even to all. This is especially worrying in high-stakes contexts with substantial knowledge asymmetries, as in political deliberation

with contributions from experts. This brings a new dimension to Peter's problem of what and how scientists should contribute to political discourse—even if scientists only assert what they know, their argumentative advantage could be detrimental to decision quality. This needs to be balanced against the benefits of their knowledge.

4 Consequences for knowledge-based decisions

4.1 The three arguments and KCAA

So far, we have examined three insightful contemporary arguments about the functioning of our epistemic lives and communities, pertaining to science, politics, and the role of reasoning more generally. Each of these arguments fits, *prima facie*, poorly with the tenets of KCAA, namely that knowledge provides the normative standard for assertions and actions (K- NORM) and that knowers play a central role in epistemic communities as reliable informants (K- FUNCTION). I have argued, in each case, that the irrelevance or unsuitability of knowledge is merely apparent; items of knowledge may not play the glamorous role that the authors have seen fit to emphasize, but knowledge can still be seen as playing an indispensable role in each of these contexts, enabling scientific progress, informed political decisions, and the exchange of legitimate reasons in argumentation behind the scenes. This addresses the first problem targeted by the paper, regarding the role of knowledge in our epistemic lives, and specifically in the question of why knowledge does not seem to play the role we would expect it to in contemporary accounts. In short, the proponent of KCAA can argue that knowledge does play a key role; this role simply has not been made explicit or spelled out in detail. Or in other words, KCAA is more compatible with the argumentation of Dang and Bright, Peter, and Mercier and Sperber than it first appeared.

The preceding discussion also helps us to address the second target problem, however: how should we characterize rational, knowledge-based decisions? By highlighting diverse aspects of real-life knowledge generation and usage, the arguments about science, politics, and reasoning provide a new perspective on knowledge-based decisions. Specifically, I draw out four consequences for knowledge-based decisions on the basis of the preceding discussion. These pertain to the role of uncertainty, the relevance of complexity, the “myside” bias of individuals, and the asymmetry between knowers and non-knowers. An important upshot will be that we need to theorize about the *process* of transforming knowledge into choice, and not just the outcome of fixed decision problems.

4.2 The role of uncertainty

Discussions of the knowledge norm for action often reflect a perception that the uncertainty we face presents a problem for basing decisions on knowledge. In the context of this paper, this shows most clearly in Peter's argument against the knowledge norm for contributions to political deliberation; even the current best scientific consensus is uncertain, hence scientists must be able to make contributions which are uncertain,

since it is better to use such contributions in the decision-making process than to forego them. Similarly, on Dang and Bright's account, scientists must put forward uncertain conclusions because science could only proceed at a snail's pace if we had to wait for certainty.

One way or another, it seems clear that decisions must reflect our uncertainty, in particular in cases of non-trivial decision-making. Contrary to common perception, however, the conflict between basing decisions on knowledge and respecting the underlying uncertainty is not substantial. Rich (2021) spells out an argument for this as a way of defending the knowledge norm of action; so, this point is not entirely new and can be made independently of the arguments in this paper. The political decision-making context discussed here provides a new way of seeing the point and its importance, though, and it is relevant to our subsequent conclusions as well.

I have already argued, in response to Peter, that scientists' contributions to deliberation should reflect the underlying uncertainty, even though this will make their contributions more complex, since the decisions being made are important and decision quality will generally suffer if we erase nuance and details and pretend to know what we do not know. For example, we will be better equipped to decide how to prepare for rising sea levels if we are told the range of possibilities indicated by the models and the scientists' judgments of these levels' likelihoods, than if we are given a simplistic item of non-knowledge like a specific magnitude of increase. The decision-makers are free to use this complex knowledge in various ways, for example by simplifying it to suit their needs and goals. To put it differently, to say that a decision is "based on" knowledge is to speak loosely, and the most sensible way to interpret the injunction to "base decisions on knowledge," on my view, leaves a lot of freedom to the decision-maker to transform the knowledge into a choice in different ways, for example by (deliberately) focusing on the most likely sea level rise or on a worst-case scenario and choosing a policy that works best for that chosen case.

This position fits well with the literature on decision-making for hard cases with severe uncertainty, for example given climate uncertainty or the consequences of genetic modification. Nobody here seems to advocate dealing with severe uncertainty in a decision-making context by replacing scientists' complex understanding of the possibilities with highly simplified generalizations—certainly not before representing the choice problem in a way that reflects the real underlying uncertainty. Take for example what Mitchell writes (2007, pgs. 61–62):

Policy makers would like neat, certain answers to questions of risk so that an easily enforceable policy can be made. However, fixed probability assignments cannot reflect our scientific knowledge in these situations. We cannot pretend that there is certainty when there is not—and we cannot hold out for certainty when it is not going to be found. ... [T]o make our policy depend on [certainty] is a mistake.

Again, the crucial point here is that basing a decision on knowledge doesn't mean ignoring uncertainty; instead, it means taking it seriously, basing a choice on complex-but-known propositions rather than simple-but-false ones.

This leaves open exactly how to characterize our uncertain knowledge-based choices. There are many possibilities, and different characterizations may be suited to different purposes. For some purposes, we might characterize qualitative, reasons-based choice. Then, for example, we could represent agents as having ordinary propositional knowledge (Williamson, 2005) or “probabilistic knowledge” à la Moss (2017) (allowing knowledge to be a graded attitude). For other purposes, we might be better served by a formal decision theory interpreted such that a knowledge standard would be applied to some components. Variations thereof have been proposed, for example, by Levi (1980), Weatherson (2012), Hawthorne and Stanley (2008), Schulz (2017), and Rich (2021); Goldschmidt (2023) discusses how we might evaluate such proposals and determine their contexts of application. I think that K- NORM, which we are here taking as a premise, *prima facie* leaves room for all of these characterizations.

Considering high-stakes choices beset by serious uncertainty, however—like many big policy choices—suggests that (at least for those kinds of choices) the simplest of those characterizations, or those in which knowledge plays the smallest role, will not be the most appropriate. Hence, we may need to turn to a characterization such as Schulz (2017) in which higher-order knowledge is used in higher stakes situations (which fits well, for example, with the call for scientists to consider the broader implications of error (Douglas, 2009)). Or, if we are in a situation of radical uncertainty—like those described by Mitchell above—then we may need to apply a theory of choice that doesn’t force us to oversimplify by using fixed probabilities. For example, we might apply a formal decision theory that accommodates imprecision (Rich, 2021). No matter what, a careful consideration of what counts as known and unknown can’t be avoided; uncertainty is always there, but that is no excuse to be careless.

4.3 The relevance of complexity

The second important lesson for those seeking an account of knowledge-based decisions has to do with complexity and its consequences. Real-world decision-making is complex (in part due to the uncertainty emphasized above). We don’t have the capacity to consider everything we know and don’t know about the world, our options, and their consequences all at once. Because of this, we need not only normative theories of decision-making for ideal agents, but normative theories that apply to bounded agents like us.¹⁸

One might have thought that a normative theory for bounded agents is only a challenge for adherents to highly formalized accounts of rational decision-making like expected utility theory. Based on the foregoing discussion, I will argue that this is not the case, though: The upshot of this section is that we have scarce short- or medium-

¹⁸ Weirich (2004) is an example of someone who has taken up this challenge for the expected utility approach to rational decisions. He explains the challenge thusly (Weirich, 2004, pg. vii):

Normative decision theory ... invokes idealizations to put aside problems arising from our cognitive limits and mistakes framing our decision problems. As decision theory matures, it pays more attention to its idealizations and works to remove them. Its ultimate goal is a general theory that evaluates and provides direction for our decisions, decisions made by nonideal agents in nonideal circumstances.

Weirich discusses other ways of de-idealizing decision theory in Ch. 1, before providing his own theory.

term prospects for a fully general, frame-independent, outcome-based characterization of knowledge-based decisions such that if we look at a real decision situation that a real agent faces, we can uniquely identify the action the agent would take if they base their choice properly on their knowledge. Both the agent and we—the modelers—are bounded in ways that prevent this. I will explain why this is the case and argue that we should respond by stepping back to investigate the process by which knowledge-based choices are made, in particular how those choices are characterized or framed.

4.3.1 The problem of relevance

Why can't we simply make a general theory of knowledge-based decisions that characterizes the decision a real agent should make, given their knowledge? An important problem that stands in the way is known as the *problem of relevance*, which came out of the discussion of the notorious frame problem in AI (for characterizations of the different problems, see Samuels 2010).

The previous discussion of the argumentative theory of reasoning hinted at the difficulty. There, I claimed that trying to get one's way in an argument typically takes the form of selectively presenting items of knowledge, rather than presenting items of non-knowledge and hoping not to be called out. In my example, Esther and Frana had selected weather and the appeal of movies and picnics as relevant to their activity choice, while ignoring their obligations for their friend's birthday party. Can we provide a general theory of knowledge-based decisions that will tell us which items of knowledge Esther and Frana *should* use in making their choice? A large literature in cognitive science, inspired by AI research, basically says that we cannot.

Glymour (1987, pg. 65) phrases the general problem thusly:

Frame Problem Instance: Given an enormous amount of stuff, and some task to be done using some of the stuff, what is the *relevant stuff* for the task?

In our case of interest, the stuff is an agent's knowledge, and the task is decision-making. It is widely acknowledged that solving this problem of relevance is extremely hard (if not impossible), in particular if the agent in question is as complex as a human, if the solution to the problem is to be computationally tractable, and especially if the relevant stuff is to be identified either correctly or well enough to enable human-level performance on the task (the literature on the frame problem is enormous, but for present purposes see especially Glymour 1987, Samuels 2010).

Given the consequence elaborated in Sect. 4.2---that knowledge is compatible with uncertainty---this relevance problem becomes more acute because we see that people don't have a paucity of knowledge, but rather a plethora. Especially through the discussion of Peter's and Mercier and Sperber's works, we have also seen that the relevant items of knowledge are not pre-determined or given, but have to be selected for the purposes of decision-making. Which items are selected goes a long way to determining what choice is ultimately made. But we know from the literature on the various frame and relevance problems that it's probably theoretically impossible to find the relevant knowledge in an optimal way, and even if humans solve the problem in a non-optimal way, we have no clue how they do it. To get some further intuition as to why the problem is so hard, note first that any item of knowledge can *in principle* be

relevant to any choice, but inspecting each item of knowledge to check its relevance to a given choice problem would be computationally intractable.

The typical approach when discussing choice problems—whether these are described formally or informally, knowledge-based or not—involves simply presuming that we know what is relevant, and describing the problem in terms of those features. Occasionally, this strategy is made explicit. For example, Savage (1954, pgs. 82–83) points out “the practical necessity of confining attention to, or isolating, relatively simple situations in almost all applications of the theory of decision developed in this book.” Savage then gives a formal treatment of how “small world” decision problems relate to the “grand world” and in particular of the conditions for the small worlds to be satisfactory. Bradley (2017, Ch. 1) argues that a decision problem may be legitimately represented in multiple ways, but that our theory should ideally yield the same verdict for any of these framings; but this requires, for one, that the most relevant factors for the decision-maker are included.¹⁹

This might be a fine approach for an abstract normative theory for ideal agents, but not for bounded agents (or their bounded modelers). The problem of relevance tells us that in practice, we cannot verify that there is no further knowledge of the agent which, if we incorporated it into our representation of the decision problem, it would change our verdict about what the agent should do. We may be able to show in a particular case that incorporating some specific further knowledge *would* change the choice, and thus that it should be included. The fact that we can’t point to such knowledge, however, doesn’t prove that it does not exist. Again, potentially any item of knowledge could be relevant,²⁰ and we cannot check them all. This means that if we want to apply any (informal or formal) characterization of the normative relationship between choice inputs and choice outputs to a real-world choice problem, there will come a point at which we need to assume that the inputs we are considering are sufficiently complete, although we cannot be sure.

As noted, this kind of assumption is standard, and for many purposes we need not worry too much about it. Clearly, though, this is a significant assumption, and we are not telling the whole story. This is especially problematic when we as theorists are not in a position to just assert which knowledge we think a decision should be based on. Many of the most interesting and important choice problems are like this, though. Why should GDP be relevant to a choice about tax policy, but not the happiness of the population, or the income distribution of a different population, or the stress levels of squirrels? Why should Esther and Frana choose how to spend their afternoon on the basis of weather and film schedules, but not on the basis of their friend’s birthday? Why should the politicians decide how to respond to sea level rises by considering

¹⁹ Thanks to an anonymous reviewer for pointing out the relevance of these discussions.

²⁰ A nice example to illustrate this is given by Samuels (2010, pg. 5): “on the face of it, the current cost of tea in China has little to do with whether or not my brother’s baby in England will cry on Saturday morning. But suppose that I believe my brother has stocks invested in Chinese tea, that he reads the business section of the newspaper every Saturday morning, and that on reading bad financial news he tends to fly into a rage.” It is not hard to imagine a decision problem in which Samuels has (perhaps probabilistic) knowledge of Chinese tea prices and faces a decision problem for which this knowledge makes a difference. It is just as easy to see that if the tea prices aren’t present in his decision framing, we are typically not going to be in a position to detect this important omission. See Samuels (2010, Ft. 12) for the origins of the example.

some set of factors (maybe flood damage to city centers and costs of measures) but not other factors (maybe impact on plant and animal life or the opportunity cost of using the necessary material to build flood walls)? When we think about these kinds of examples, and others inspired by the accounts of the authors discussed in this paper, it becomes clear that deciding relevance is fundamental to deciding how to act. Relevance is thus subject to negotiation when groups need to make decisions—a point brought out by our discussion of Mercier and Sperber's theory. Once we decide the relevance question, the rational choice often becomes a trivial matter; I can always frame a choice problem such that my preferred option comes out as uniquely rational, if you let me get away with it. The heart of a real-world decision problem is often in what we pay attention to and what we do not.

For proponents of KCAA, then, who are tasked with characterizing rational, knowledge-based choices, considering messier real-life choice problems as the preceding discussion has encouraged us to do makes salient an important gap in our theorizing so far, namely adequate decision framing and the relevance problem. To be clear, the problem of relevance is not only relevant to KCAA advocates, or only for advocates of a formal knowledge-based decision theory. I think, for example, that economists using decision theory should also pay more attention to the impact of relevance assumptions on ultimate judgments (as Samuels (1989) argues nicely, in somewhat different terms). Philosophers may be in an especially good position to start theorizing about this, though, now that we can see how fundamental judgments of relevance can be to determining choices and outcomes. One might have thought that requiring decisions to be based on knowledge would provide a sufficient safeguard on decision quality, but the discussion of Mercier and Sperber's work in particular shows that it does not.

4.3.2 Addressing relevance

We should look for ways to systematically address relevance, then. It will greatly enrich our understanding of decision-making. A promising approach could involve characterizing normatively acceptable decision-making *processes*. As I will point out, there is a fair amount of existing literature to draw on if we take this approach, in particular if decision-making takes place in a social setting. In Esther and Frana's case, if they decide to have a picnic, we might criticize this choice on the grounds that they could no longer justify it if someone else were bring up an obvious item of knowledge that both of them know, namely that they had promised to prepare food for Gengis' birthday party and they cannot both picnic and cook. By taking such an approach, we may be able to say more about the characteristics of (good) knowledge-based decisions.

This perspective on knowledge-based choices also fits well with other new lines of research. Consider, for example, Bermúdez (2020) research on decision framing. As already explained, choices must be based on some subset of the agent(s)'s knowledge which has been deemed relevant. Another way to express this, invoking a more decision-theoretic perspective, is that choices are made relative to particular frames. Different ways of framing the same problem can lead to different decisions. While frame sensitivity has long been seen as a hallmark of irrationality, Bermúdez (2020)

offers a detailed and compelling argument that frame sensitivity can be perfectly rational, and moreover that frame changes can be used for good. For example, a person may be better able to withstand the temptation to eat junk food if they frame eating broccoli in a positive way, including being resolute and improving health, rather than in a negative way (like foregoing a treat). Similarly, the right frame could support mutually beneficial cooperation over a less advantageous alternative. So, while often skipped in presentations and applications of decision theory, selecting a frame is not only a prerequisite to choice, but a step which can go a long way to determining the choice that gets made (or that gets dubbed rational)—and the quality of that choice from a wider perspective—and hence an extremely important part of decision-making.

It is worth emphasizing that it is possible to fruitfully theorize about decision-making in a way that includes the framing step, and even to make normative judgments about the framing process—despite the fact that even we theorists cannot identify the unique best frame for any real choice problem. In addition to Bermúdez’s work, a further strain of research demonstrates how progress can be made.

Proponents of “Conviction Narrative Theory” (Tuckett and Nikolic, 2017) integrate various lines of research to put forward a theory according to which people use narratives to make choice in typical real-world cases. In line with our discussion so far, they point to six challenges of everyday choice: radical uncertainty, fuzzy evaluation, social embeddedness, imagination (of possible futures), commitment (to a course of action over time), and sense-making (of the present) (Johnson et al., 2020). These challenges all reflect the substantial gap between a neatly framed decision—possibly in table form with actions, events, and outcome values specified—and a messy real-world choice for which all of these must be determined by the decision-maker.

As messy as real-world choice is, proponents of Conviction Narrative Theory have put together a fairly detailed picture of how, descriptively, choice proceeds in typical difficult cases. As they explain the basic idea (Johnson et al., 2020, pg. 3-4),

we use narratives to make sense of the past, imagine the future, commit to action, and share these judgments and choices with others. ... Governments debate whether a virus is more like flu or plague; these narratives yield very different explanations of the situation, hence predictions about the future, hence emotional reactions to particular options. The couple can interpret their fights as signaling differences in fundamental values or resulting from temporary stresses; either narrative can explain the fights, portending either a dark or rosy future. The toaster CEO might consider her company ossified, complacent, or innovative; these narratives have different implications about the risks and benefits of new ventures, motivating different decisions. In each case, the decision-maker’s first task is to understand the current situation, which informs how they imagine a particular choice would go, which is deemed desirable or undesirable based on how the decision-maker would feel in that imagined future.

The details of this theory are not relevant to this paper’s main argument. What I want to argue here is, first, that an account of knowledge-based decisions with bearing on real choices must address the relevance problem, and secondly, that looking at the process by which relevant items of knowledge are selected is the most viable way to get

at relevance. The example of Conviction Narrative Theory shows that theorizing about this part of the choice process is possible, and that there is more existing descriptive theory available than one might have thought.

What the example also shows, though, is that the normative aspect of the choice framing process is under-studied. According to Conviction Narrative Theory, descriptively, learning and cultural evolution act on narratives. Still, narratives can be bad and misleading, and little is said about how to avoid bad narratives or aim at better ones. The authors say that when optimization is meaningless (as in difficult choice problems), we also can't sharply divide choices into the rational and irrational. This may be partly the nature of the beast, but we should nonetheless have a normative account of decision-making (for bounded agents) which covers this aspect of choice. A knowledge norm for inputs could help to get us started, providing more normative teeth than the mere appeal to learning. The same goes for procedural norms about how that knowledge may and may not be used.²¹

I have now argued that an account of knowledge-based decisions has less need of defending the requirement that *only* knowledge be used in choice than it has of explaining *which* knowledge should be used, and *how*. I think we will be able to make significant progress towards such an account by directing more attention to the framing part of decision-making.

4.4 Myside bias and decisions

Mercier and Sperber (2017) drive home the point that individuals exhibit “myside bias,” easily finding reasons to support their own beliefs or preferred actions. Decisions may be poor, then, if the agents simply rationalize actions which are somehow convenient or otherwise appealing to them (such as lazy or unfair actions). Troublingly, when agents anticipate having to justify their actions to others, they tend to choose actions which are easy to justify but possibly worse than what they otherwise would have chosen. Again, this can lead to poor but rationalizable decisions.

These scenarios present a challenge for a normative account of knowledge-based decisions. The reason is that, as alluded to above, there is not in general a unique action that we can show to be supported by the agent's total knowledge. As I explained, agents are swimming in knowledge, and decision-making requires selecting the items of knowledge which are taken to be relevant to the problem at hand. This means, however, that the agent may well be able to produce items of knowledge which support any given action in any given situation, even actions which are lazy, silly, unjust, or unrewarding. What kind of theory can help us to catch and condemn such actions?

The above discussion suggests that an account appealing to the process of decision-making, and possibly to hypothetical argumentation with other agents regarding the choice, may be the most promising. There will always be additional individual items of knowledge which could have been used to support a different action, and appealing to the action recommended by the whole of the agent's knowledge is also a non-starter (because we can't characterize it in a principled way). An action that other

²¹ I have defended a process-based approach to normative questions in *citations suppressed for review*.

agents can be expected to criticize, and which the agent couldn't defend well given their own knowledge, is reasonably not a good decision, whether based on knowledge or not. This suggests that it would be worth trying to integrate an argumentative approach to reasoning (or a social approach to judgment and decision-making) with knowledge-centric ideas. The social, argumentative approach would benefit from a clear explication of the role of knowledge, while the project of providing an account of knowledge-based decisions would benefit from the existing understanding of biased reasoning and the impact of other agents on decisions.

This is in line with some of Peter (2021) arguments. After rejecting substantive epistemic norms such as the knowledge norm, she argues for procedural norms of political deliberation. Specifically, she suggests that a "responsiveness norm" and an "epistemic justice norm" should be spelled out as part of the landscape of norms for political deliberation (which aims at decision-making). These norms would require that people respond to the reasons given by others and that their own contributions not be a result of identity prejudice. Peter's discussion could be a useful starting point for an account of how knowledge should be collected and used for decision-making purposes, in a way that puts some limits on agents' biases.

Regardless of how knowledge-based (normative) decisions are characterized, though, the problems of motivated reasoning and bias should be taken seriously.²² One might have thought that the requirement to base decisions on knowledge would have provided protection against these problems (and a way to criticize problematic biases), but we can see that any such protection is limited.

4.5 The asymmetry between knowers and non-knowers

The final lesson for an account of knowledge-based decisions builds on the third. Knowledge can both help to improve decision quality, for example by providing more specific information about the consequences of the available actions, and worsen decision quality, by providing more ways for a decision-maker to justify an appealing but poor choice. This is already an issue for individual choosers; Mercier and Sperber emphasize the enhanced ability of good reasoners to justify problematic conclusions, and knowers are surely similarly positioned.

The difficulty becomes even more interesting in a social or argumentative context. I pointed this out above, but elaborate here. Due to the asymmetry between knowers and non-knowers about a topic, a knower is not only a reliable informant (as K-FUNCTION says), but especially well situated to manipulate the non-knower. For example, suppose we are arguing about what to order for dinner, and I know more about nutrition than you do. I may well be able to get you to agree to my preferred dishes by highlighting their beneficial properties and pointing out the negative qualities of your favored options. This may be the case even if I know of further, unmentioned reasons which could

²² Carter and McKenna (2020) address the epistemic status of beliefs formed by motivated reasoning, which as they point out, has been ignored by "mainstream epistemology." The authors address the concern that motivated reasoning prevents agents from getting knowledge, but the problem here is that motivated reasoning can also take place *using* knowledge that the agents have already secured.

make your favored options look as good or better than mine. If we then order the dishes that I prefer and skip the ones you prefer, we have made a knowledge-based choice, specifically based on my knowledge of nutrition. We may nonetheless want to say that, normatively, we made the wrong choice, in this case as a result of my ability to use my greater knowledge to manipulate the choice process.

This connection between Mercier and Sperber's theory and K- FUNCTION shows the need, again, for the argumentative theory of reasoning to incorporate knowledge (being vigilant against falsehoods doesn't help if truths are being used to manipulate others) and for the knowledge-centric approach to consider social and strategic aspects of reasoning and choice (such as the less desirable capabilities of knowers). In particular, any *normative* account of knowledge-based choices, whether in an individual or a group setting, must avoid allowing knowledge to be used in unfair, dishonest, or otherwise problematic ways.

This is a real danger in cases of practical interest, such as political deliberation. In such cases, scientists have relevant expertise, and we rightly want to incorporate their knowledge into the decision process. Yet we must also guard against the possibility that scientists highlight items of knowledge that support their personally preferred policy, and that in the absence of any kind of check or balance to their relatively greater power, non-scientists' preferences are improperly subverted in the same way that your dinner preferences were subverted by my nutritional arguments. Again, a normative account that succeeds at this will probably focus on aspects of the process of making choices, along the lines Peter suggests for the political context, as mentioned above.

5 Conclusion

This paper was motivated by two interconnected problems. First is the problem of the strength of KCAA as an approach to problems in epistemology, broadly construed. The approach is compelling in many respects, but appears to be at odds with neighboring literature which should be taken seriously. Second is the problem of the knowledge norm of action specifically: even if we accept K- NORM, there are many open questions regarding how to characterize rational, knowledge-based choices.

The paper aimed to address both of these problems by synthesizing key aspects of KCAA with insights from three contemporary accounts that at first appeared at odds with it. I have argued that significant consensus between all of these strains of research can be found, and moreover, that all can be enriched and further developed through such a synthesis. All the underlying accounts get at different aspects of the basic and critical process of producing, sharing, and applying knowledge. As such, they supplement and complement one another.

Importantly, the three contemporary accounts confronted with knowledge-centric ideas also tend to be more focused on real-world epistemic activities, and more descriptively grounded, than is the average epistemology research. As a result, they allow us to shine the spotlight on particular aspects of knowledge-based choice that otherwise receive little attention. In brief, they show the importance of the uncertainty reflected in complex real-world knowledge, the need to determine what knowledge is relevant,

the bias of agents in judging relevance, and the resultant potential for relative expertise to be used in problematic ways. Collectively, these considerations show that we need additional theorizing about the process by which (knowledge-based) choices are framed and made.

Acknowledgements The author thanks participants of the Oberseminar in Theoretical Philosophy at the TU Dresden for comments on an earlier version of the manuscript, especially Anna-Maria Asunta Eder, Norbert Engemaier, Rico Hauswald, Roman Heil, Andy Mueller, Moritz Schulz, Sergiu Spatan, and Sebastián Sánchez Martínez. She also thanks two referees of this journal for helpful feedback and suggestions and the audience of the LSE Popper Seminar for discussion. Valuable inspiration for the arguments came from Fabienne Peter during her Wittgenstein lectures in Bayreuth and from Iris van Rooij during discussions of other topics.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations The research did not involve generating data.

Competing interests The authors declare no competing interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alexandrova, A., & Northcott, R. (2009). Progress in economics: Lessons from the spectrum auctions. *Oxford Handbook of Philosophy of Economics*, 306–337.
- Anderson, C. (2015). On the intimate relationship of knowledge and action. *Episteme*, 12(3), 343–353.
- Beddor, B. (2020). New work for certainty. *Philosophers' Imprint*, 20(8)
- Bermúdez, J. L. (2020). Frame it again: New tools for rational decision-making. Cambridge University Press
- Blessenohl, S., & Sarikaya, D. (2022). Science advice: Making credences accurate. *Synthese*, 200(2), 174.
- Bradley, R. (2017). *Decision theory with a human face*. Cambridge University Press.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167–189.
- Carter, J. A., & McKenna, R. (2020). Skepticism motivated: On the skeptical import of motivated reasoning. *Canadian Journal of Philosophy*, 50(6), 702–718.
- Craig, E. (1991). Knowledge and the State of Nature: An Essay in Conceptual Synthesis. Clarendon Press
- Dang, H., & Bright, L. K. (2021). Scientific conclusions need not be accurate, justified, or believed by their authors. *Synthese*, 199(3), 8187–8203.
- Douglas, H. (2009). *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.
- Dutilh Novaes, C. (2018). The enduring enigma of reason. *Mind & Language*, 33(5), 513–524.
- Dutilh Novaes, C. (2022). Argument and Argumentation. In Zalta, E. N. and Nodelman, U., editors, The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, Fall 2022 edition
- Elgin, C. Z. (2020). The mark of a good informant. *Acta Analytica*, 35(3), 319–331.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *The Philosophical Review*, 111(1), 67–94.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.

- Fassio, D., & Gao, J. (2021). Do we really need a knowledge-based decision theory? *Synthese*, 199(3), 7031–7059.
- Frigg, R., & Hartmann, S. (2020). Models in Science. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2020 edition
- Gao, J. (2017). Rational action without knowledge (and vice versa). *Synthese*, 197, 1901–1917.
- Glymour, C. (1987). Android epistemology and the frame problem. In Z. Pylyshyn (Ed.), *The Robot's Dilemma: The Frame Problem in Artificial Intelligence* (pp. 63–75). Ablex.
- Goldschmidt, Z. (2023). Foundations for knowledge-based decision theories. *Australasian Journal of Philosophy*. forthcoming
- Goodman, J., & Holguín, B. (2022). Thinking and being sure. *Philosophy and Phenomenological Research*
- Greenough, P., & Pritchard, D. (Eds.). (2009). *Williamson on Knowledge*. Oxford: Oxford University Press.
- Grüne-Yanoff, T. (2009). Learning from minimal economic models. *Erkenntnis*, 70(1), 81–99.
- Hannon, M. (2018). *What's the Point of Knowledge?: A Function-first Epistemology*. Oxford: Oxford University Press.
- Hardwig, J. (1985). *Epistemic dependence*. *The Journal of philosophy*, 82(7), 335–349.
- Hardwig, J. (1991). The role of trust in knowledge. *The Journal of Philosophy*, 88(12), 693–708.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., & Lasonen-Aarnio, M. (2009). Knowledge and objective chance. In P. Greenough & D. Pritchard (Eds.), *Williamson on Knowledge* (pp. 92–108). Oxford: Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105(10), 571–590.
- Heil, R., Koscholke, J., Rich, P., & Schulz, M. (2022). Knowledge and decision: Introduction to the synthese topical collection. *Synthese*, 200
- Johnson, S. G., Bilovich, A., & Tuckett, D. (2020). Conviction narrative theory: A theory of choice under radical uncertainty. *Behavioral and Brain Sciences*, pages 1–47.
- Kitcher, P. (1990). The division of cognitive labor. *The Journal of Philosophy*, 87(1), 5–22.
- Kitcher, P. (1995). *The Advancement of Science: Science without Legend, Objectivity without Illusions*. New York: Oxford University Press.
- Kneer, M. (2018). The norm of assertion: Empirical data. *Cognition*, 177, 165–171.
- Lackey, J. (2007). Norms of assertion. *Noûs*, 41(4), 594–626.
- Levi, I. (1980). *The enterprise of knowledge: An essay on knowledge, credal probability, and chance*. MIT press
- Marsili, N., & Wiegmann, A. (2021). Should I say that? an experimental investigation of the norm of assertion. *Cognition*, 212, 104657.
- McGlynn, A. (2014). *Knowledge First?* Hampshire: Palgrave Macmillan.
- McKinnon, R. (2013). The supportive reasons norm of assertion. *American Philosophical Quarterly*, 50(2), 121–135.
- Mercier, H. (2010). The social origins of folk epistemology. *Review of Philosophy and Psychology*, 1(4), 499–514.
- 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*. Harvard University Press.
- Mitchell, S. D. (2007). The import of uncertainty. *The Pluralist*, 2(1), 58–71.
- Moss, S. (2017). *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Norman, A. (1998). Towards a logic of resolution-oriented dialogue. *Dialogue and Universalism*, 8(11/12), 155–167.
- Peter, F. (2021). Epistemic norms of political deliberation. In Hannon, M. and de Ridder, J., editors, *The Routledge Handbook of Political Epistemology*, Routledge, 395–406
- Reiss, J. (2012). The explanation paradox. *Journal of Economic Methodology*, 19(1), 43–62.
- Rich, P. (2021). The key to the knowledge norm of action is ambiguity. *Synthese*, 199, 9669–9698.
- Samuels, R. (2010). Classical computationalism and the many problems of cognitive relevance. *Studies in History and Philosophy of Science Part A*, 41(3), 280–293.
- Samuels, W. J. (1989). The methodology of economics and the case for policy diffidence and restraint. *Review of Social Economy*, 47(2), 113–133.
- Savage, L. J. (1954). *The Foundations of Statistics*. New York: Dover Publications Inc.
- Schiffer, S. (2007). Interest-relative invariantism. *Philosophy and Phenomenological Research*, 75(1), 188–195.
- Schulz, M. (2017). Decisions and higher-order knowledge. *Noûs*, 51(3), 463–483.

- Schulz, M. (2021). Practical reasoning and degrees of outright belief. *Synthese*, 199(3), 8069–8090.
- Simion, M. (2016). Assertion: Knowledge is enough. *Synthese*, 193(10), 3041–3056.
- Smaldino, P. E. (2022). Five models of science, illustrating how selection shapes methods. In A. De Block (Ed.), *Ramsey, G* (pp. 19–39). The Dynamics of Science: Computational Frontiers in History and Philosophy of Science, University of Pittsburgh Press.
- Solomon, M. (1992). Scientific rationality and human reasoning. *Philosophy of Science*, 59(3), 439–455.
- Stanley, J. (2005). *Knowledge and practical interests*. Clarendon Press.
- Tuckett, D., & Nikolic, M. (2017). The role of conviction and narrative in decision-making under radical uncertainty. *Theory & Psychology*, 27(4), 501–523.
- Turri, J. (2013). The test of truth: An experimental investigation of the norm of assertion. *Cognition*, 129(2), 279–291.
- Turri, J., Friedman, O., & Keefner, A. (2017). Knowledge central: a central role for knowledge attributions in social evaluations. *Quarterly Journal of Experimental Psychology*, 70(3), 504–515.
- Van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Oxford University Press.
- Weatherson, B. (2012). Knowledge, bets and interests. In J. Brown & M. Gerken (Eds.), *Knowledge Ascriptions* (pp. 75–103). Oxford: Oxford University Press.
- Weirich, P. (2004). Realistic decision theory: Rules for nonideal agents in nonideal circumstances. Oxford University Press
- Wikipedia contributors (2023). Retraction in academic publishing — Wikipedia, the free encyclopedia. https://en.wikipedia.org/w/index.php?title=Retraction_in_academic_publishing&oldid=1157470440. [Online; accessed 16-August-2023]
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly*, 55(219), 213–235.
- Zollman, K. J. S. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72(1), 17–35.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.