

# Questions Should Have Answers

Michael Deigan  
*Rutgers University*

Some questions appear to be unanswerable for us because they seem to have too many possible answers that we cannot rule out. How long will humanity continue to exist? A billion years? A million? A thousand? Less? I wonder. But I have no hope of finding out, unless the end is truly near.

Other questions appear to be unanswerable for us because we seem unable to come up with *any* possible answer that *isn't* ruled out. Many philosophical questions are like this. Is the liar sentence true? No answer I can think of can be right. Why is there anything at all, facts included? Again, every potential answer I can think of seems to fail. Non-philosophical questions can be like this, too. Where is my umbrella, given that I've checked all the places where I think I could have left it?

There's a puzzling difference between these two kinds of apparent unanswerability, one that I think has interesting implications for thinking about norms of rational inquiry. The difference is that the latter kind of unanswerability, but not the former, seems to demand cognitive revision.

I may feel some disappointment that my curiosity on the matter of humanity's fate will not be satisfied, and it isn't advisable to dwell on the question too much, given the hopelessness of settling it.<sup>1</sup> Nevertheless, there doesn't seem to be anything particularly wrong with continuing to wonder this question, idle though this wondering may be. Wondering an apparently unanswerable question of this kind does not itself seem to constitute or even indicate any kind of rational failing, nor does it seem to demand that I revise any of my mental states.

---

\*For very helpful discussions, thanks to Keith DeRose, Daniel Ferguson, Daniel Greco, Jason Stanley, Zoltán Szabó, and Nadine Theiler.

<sup>1</sup>It may well make sense, though, to dwell on related questions about probabilities, for decision-making purposes.

In contrast, continuing to wonder apparently unanswerable questions of the second kind—wondering while having ruled out all the answers one can think of—seems to make fitting a special kind of puzzlement and discomfort which goes beyond simply not being able to know what you want to know. One feels rational pressure to extricate oneself from this puzzled state, either by dismissing the question or else coming up with an answer that at least *could* be right, given one's evidence. I feel some relief when I temporarily convince myself that we can reject the question about why there is anything, on the grounds that it falsely presupposes that there is such an explanation. I feel a similar relief when I think of some place where my umbrella could be, even if it is some place I won't be able to check.

What explains the contrast? This paper attempts an explanation, in two stages.

First stage: I posit and defend the following norm of rationality.<sup>2</sup>

### **QUESTIONS SHOULD HAVE ANSWERS (QSHA)**

Don't both

(i) wonder Q

and

(ii) reject all answers to Q you can conceive.

Defending this will involve showing that QSHA gets a range of cases right, including some cases that we can't get right by appeal to simpler and more familiar norms, as well as answering a concern that QSHA is implausibly demanding. I'll do this in §§1–3, clarifying various aspects of QSHA along the way.

---

<sup>2</sup>I take this to be a pro tanto norm of structural rationality, though I will not take a stand here on whether such norms are non-derivatively normative (Broome 2013), or must derive their normativity from substantive rationality (Kolodny 2007), or should be understood in terms of a distinct kind of rational pressure (Fogal [forthcoming](#)). Additionally, my own view is that it is a norm of *epistemic* rationality, for reasons along the lines given in Friedman ([forthcoming](#)). See also Kelp (2020) for another recent defense of an inquiry-centric approach to epistemology. I don't wish to get mired in border disputes here, though, so will only argue that there is such a norm, not that we must think of it as epistemic.

Second stage: I propose in §4 that QSHA is partially explained by a more fundamental norm concerning what contents one should be able to think, which I speculate is itself traceable to a kind of explore-exploit tradeoff.

Besides resolving the initial puzzle, my account has broader implications. It implies that norms of rationality extend beyond actions and attitudes to mental states which aren't attitudes at all, but rather abilities to have attitudes towards particular contents. Norms of rationality govern our *concepts*, in other words. On this view, to rationally inquire is not just to manage one's beliefs and evidence gathering activities well, but also to manage one's concepts well.

## 1. QSHA Violations and Doxastic Norms

If one both wonders some question  $Q$  but also rejects all the answers to  $Q$  one can conceive, one thereby violates QSHA. In this section, we'll look at three kinds of situations which lead to QSHA-violations. In a *bad rejection* situation, one has rejected the true answer to a good question, in addition to rejecting incorrect answers. In a *bad question* situation, one is wondering a question that lacks a true answer, and so may have correctly taken all its would-be answers to be untrue. And finally, in an *inadequate hypotheses* situation, one may be wondering a good question and reject only incorrect answers, but simply be unable to come up with the correct answer.

In all of these cases, we'll see, something seems to be going awry. This awryness can be explained by appeal to QSHA. In some cases—the bad rejection (§1.1) and bad question (§1.2) situations—this awryness can plausibly be explained by appeal to familiar epistemic norms for belief. But I argue that in others—the inadequate hypotheses situations (§1.3)—these belief norms will not suffice, whereas QSHA does.

### 1.1. Bad Rejections

Suppose Al wonders where his keys are. That is, that Al stands in the wonder relation to the question content expressed by “where [Al's]

keys are”.<sup>3</sup> Such a question content can be usefully modelled as a set of all its would-be answers: the propositions which would *answer* or *resolve* the question.<sup>4</sup> Following semanticists, I will take this notion of answerhood in a somewhat special sense, which excludes responses like “I don’t know” and “Huh? Al doesn’t even have keys” from counting as would-be answers. In this case, would-be answers are propositions like [Al’s keys are in his left jacket pocket] and [Al’s keys are on his desk]. I will take this relation between questions and their would-be answers as primitive, but here’s a heuristic for identifying them applied to this case: the would-be answers are the propositions  $p$  such that, if Al knew  $p$ , Al would thereby know where his keys are.

What question content is expressed by a given interrogative clause is in part determined by context.<sup>5</sup> Normally, if all Al knows about the location of his keys is that they are in the Milky Way Galaxy, then he doesn’t know where his keys are. That information isn’t enough to resolve the question expressed by “where [Al’s] keys are” in most contexts. But we can imagine contexts in which this is enough. Suppose the Al we’re talking about is Al the Alien, who is trying to decide which galaxy to warp to in order to use a keyfinding device that teleports the keys to the spaceship once they are less than 1 million lightyears away. In that case, the question expressed would include the propositions [Al’s keys are in the Milky Way], [Al’s keys are in the Andromeda Galaxy], and so on. These propositions would resolve the rather coarse-grained *where*-question that would be expressed in this context. But let’s say instead that the Al we’re talking about is an absent-minded, human professor of philosophy. Then the question content expressed by “where [Al’s] keys are” would be something like {[Al’s keys are on the kitchen table], [Al’s keys are on his desk], [Al’s keys are on the floor

<sup>3</sup>Wondering is a paradigm question-directed attitude (Friedman 2013).

<sup>4</sup>This idea goes back to Hamblin (1958) and Hamblin (1973), and has been developed by Groenendijk and Stokhof (1984) and Ciardelli, Groenendijk, and Roelofsen (2019), and others.

It’s common to just to use ‘answers’ or ‘possible answers’ for what I call ‘would-be answers’. But this can be confusing, since ‘answer’ is sometimes taken to mean ‘true answer’, and ‘possible answer’ might suggest that the proposition is epistemically or metaphysically possible, which it need not be, on my view.

<sup>5</sup>See Boër and Lycan (1975), Hintikka (1978), Ginzburg (1995), Aloni (2001, Ch. 1), and van Rooij (2003).

next to his desk], [AI's keys are in AI's jacket's left pocket], . . . }. None of these will be entailed by galaxy-level information about AI's keys' whereabouts. The question expressed in this context is finer-grained.

How exactly a context determines the granularity of the question content expressed by an interrogative is an interesting issue, but not one we can pursue here. It is important for us to note, though, that an agent may be wonder-related to a question content that has a proposition  $p$  as a member even though the agent has not entertained  $p$ , or is even *able* to entertain  $p$ .<sup>6</sup> AI might wonder a question with [AI's keys are at the office] as a would-be answer even if this proposition has yet to cross his mind. And Jackson's Mary can wonder what seeing red is like even though she is unable to grasp any proposition that would resolve the question.

AI is wondering a question, which is one step along the way to violating QSHA. In order to get the rest of the way there, he needs to reject all the answers to that question that he can conceive. To *reject* a proposition is to believe that it is not true, either because it is false or because it presupposes something that isn't true. AI rejects that his keys are on the table because he has looked there, and he rejects that his keys are in Atlantis (though perhaps only implicitly, unless someone makes the suggestion) because he thinks there is no such place.

At any given time, there is a set of propositions that an agent is able to come up with as would-be answers to a given question; these are the answers that the agent *can conceive*. Often a question has would-be answers that an agent cannot entertain at a given time, either because they don't have the conceptual resources to think the proposition in question, or because the proposition just isn't coming to mind for some shallower reason, even when the agent is trying to come up with would-be answers that they have yet to consider. In order to violate QSHA, AI will have to continue to wonder the question, even after he has rejected all the would-be answers he can think of. It's easy enough to imagine how that might happen. One simple way is if AI has mistakenly rejected the correct answer.

AI starts out being able to think various would-be answers to the

---

<sup>6</sup>See Friedman (2013, pp. 160–163) for discussion.

question he is wondering that he has yet to reject. He begins his investigation with the usual suspects: pockets and surfaces he usually walks by after getting home. Some simple empirical investigation—looking at the surfaces, reaching into the pockets—leads him to reject several such would-be answers. [Al’s keys are on the kitchen table] and [Al’s keys are in his left jacket pocket] are now rejected, among others. Annoying, but not yet dire: there are various other places Al can still think of where his keys might be.

He now considers, for example, whether he might have left them at the office. No, he concludes, they can’t be anywhere outside the house, since he must have used them to let himself in. So Al now rejects every not-in-the-house answer he can think of on these grounds. This, it turns out, was his mistake. Al did not use them to let himself in, for he also forgot to lock the door when he left. His keys really are still in his office. Having rejected this possibility, however, Al continues the search in the house, checking all the places he can think of where his keys could be: floors, chairs, bookshelves, and so on, even the refrigerator. This is to no avail, of course.

Al racks his brain for answers he might have missed, but nothing plausible comes to mind. Various implausible answers come to mind—like that his keys have been stolen from inside the house, and so are now with the thief—but Al finds he already implicitly believes them to be false, and so already implicitly rejected them. At this point, Al is unable to think of any more unrejected would-be answers. But Al does not give up wondering the question. Indeed, the intensity of his wondering is only increased. Al is now very puzzled about where his keys could be. So Al is in a state of both wondering a question and rejecting all the answers to it that he can conceive. Al is violating QSHA.

As we’d expect if QSHA were a norm of rationality, it seems that there is something wrong with being in Al’s state. It is fitting for Al to feel a certain kind of discomfort about being in this state, a feeling that those of us who have struggled to find our keys know well. And, at least if there are no other more pressing concerns, there is reason for him to get himself out of this state.

The situation seems similar to those in which agents find themselves with inconsistent beliefs: there seems to be something wrong with being in such a state, and there is at least some rational pressure to get out of it. There might not be enough pressure for a rational agent to in fact do so, since there may be other more important concerns to attend to, or because it's not clear which belief to give up, but the pressure seems there nevertheless. That there seems to be something wrong with staying in an inconsistent belief state lends some support to the proposal that there is a norm of rationality against having inconsistent beliefs. Similarly, that there seems to be something wrong with Al's QSHA-violation lends some support to the proposal that QSHA is a genuine norm of rationality.

QSHA gets Al's case right, then, but is it necessary for getting Al's case right? Given that Al's predicament is in part due to his mistaken reasoning, perhaps we could explain what's wrong with being in Al's state in terms of familiar norms of epistemic rationality governing belief and reasoning. If so, then QSHA wouldn't be adding anything here, and so positing it to account for Al's case would be premature.

Here's a simple version of one such alternative account: first, we appeal to a truth norm of belief:

#### **NEGATIVE TRUTH NORM**

Don't believe  $p$  if  $p$  is not true.<sup>7</sup>

Al violates this norm as soon as he rejects that his keys are in his office. So we don't need QSHA to say that *something* is wrong with the state Al ends up in. However, the account shouldn't stop there. One important aspect of the case is that Al doesn't—and in some sense shouldn't—feel discomfort until well after he has formed the false belief, since at that point there were various answers which he still took to be possible. QSHA predicts this, since it's not until Al runs out of alternatives that QSHA is violated. The QSHA-less account needs to say something more about it.

---

<sup>7</sup>We might take this norm to be fundamental, or we might instead take it to be derived from a more general truth norm that requires one to believe  $p$  iff  $p$  is true (see Wedgwood (2002), Shah (2003), and Engel (2013), among others) or from a norm that requires one to believe  $p$  only if one knows  $p$  (see Williamson (2000, pp. 255–256), J. Sutton (2005), and others).

A natural suggestion is that it is only once he has ruled out the other possibilities that Al should realize he has mistakenly ruled out the true answer. At the time he formed the false belief, he did not have much, if any, evidence that he had done so. But by the time he has eliminated all the plausible alternatives, Al has strong evidence that has made a mistake *somewhere*, even if he can't easily tell where. It would be reasonable for Al to repeat certain aspects of his search, to see if he can discover where he went wrong.

We might just say that Al's having evidence that he has violated the Negative Truth Norm is what is responsible for the fitting discomfort he feels.<sup>8</sup> This seems unlikely to be the be the whole story, though, since often evidence that one has *some* false belief doesn't seem to bring any discomfort with it. I have excellent reason to think that I, like everyone else, have some false beliefs. But I do not feel about this anything like what Al should feel in his state, and do not think it would be appropriate for me to feel that way.

There are various factors in Al's case that we might appeal to in order to supplement the alternative account so that it can capture this contrast. First, Al has evidence that he has a false belief that prevents him from answering a question he is wondering. Second, he has evidence that he has a false belief that has some practical significance, since it's preventing him from getting what he wants: his keys. Finally, unlike my evidence about a general fallibility, Al has evidence that he has a false belief that may be relatively easy to identify and revise. I will not work these features into a full-fledged alternative theory of what is wrong with remaining in Al's QSHA-violating state. It seems

---

<sup>8</sup> To flesh this out, we might follow McHugh (2012), saying that the Negative Truth Norm is an evaluative rather than prescriptive norm, but one that gives rise to a prescriptive norm to believe only what is probable given one's evidence. Then we could take the discomfort to be a response to the violation of this prescriptive norm. Or we might follow Smithies (2012) and say that there's a principle linking justified  $\phi$ -ing and justified belief that  $\phi$ -ing is correct; then we could say that the evidence of violating the Negative Truth Norm makes one unjustified in thinking the belief is correct, and that the discomfort is a response to recognizing this unjustified belief. Alternatively, we might extend an idea from Williamson (forthcoming), taking the lack of evidence to be an *excuse* for failing to comply with the Negative Truth Norm. Then we could hold that the discomfort is fittingly felt when one fails to comply with a norm without excuse.

likely enough that some such theory could be developed in a way that sufficiently accounts for Al's case without any need to introduce QSHA as something that goes beyond ordinary doxastic norms.

Better support for introducing QSHA, then, must come from its handling of other sorts of cases. Let us, then, turn to another way QSHA can be violated.

## 1.2. Bad Questions

Al's problem was that he had rejected the correct answer to a good question. But sometimes the problem isn't with rejecting any of the answers, it's with the question one is wondering.

Consider Bea, a child who wonders why Santa brought fancier presents to the kids in rich families. Suppose she has gone through various would-be answers to this question—that Santa brings fancier presents to rich kids because they behave better, because he has been bribed by their parents, because he is taking marginal utility and hedonic adaptation effects into account, and so on—but has found good reasons to reject them all. And suppose she gets to the point where she can no longer think of any more answers which she hasn't already rejected. Bea is bothered by this, reasonably enough: it seems to her that her apparently good question has no good answer.

Bea is correct to reject the various answers she considers. It is indeed not true that Santa brings fancier presents to the kids in rich families because he has been bribed or whatever. And we can suppose Bea has conclusive reasons supporting her rejections. But now she is in a state which violates QSHA, since she is wondering a question, yet has ruled out all the answers to it she can conceive. And, like Al, Bea does seem to be in something of a bad state here, one she should feel puzzled about and try to get out of, as positing QSHA predicts.

We should again ask, though, whether something like QSHA is really needed to explain what's wrong with being in Bea's state. On the face of it, it seems that a straightforward appeal to familiar doxastic norms will not do the trick. Al's predicament arose from his bad beliefs about would-be answers, but Bea's seems due to her wondering a bad

question, a question without any true answer, rather than to her beliefs. And QSHA gets us the attractive result that Bea's wondering is itself contributing to the violation of a rational norm.

But aren't there still false beliefs we can trace the problem to, like Bea's beliefs that Santa exists and is responsible for some of the presents some kids get? Perhaps familiar doxastic norms can account for Bea's case by impugning these beliefs. There are a few ways we might pursue such a strategy. One is to deny that there's anything irrational about the wondering itself, and take the wondering to be merely a sign of a bad belief. On another view, one admits that the wondering itself is bad, but takes its badness to be wholly derived from the badness of the beliefs in the false presupposition. One would need to introduce a norm for how non-doxastic states inherit badness from bad beliefs, but it need not look anything like QSHA. Finally, one might take the wondering to have non-derivative badness, but due to a simpler norm which says to not wonder questions with false presuppositions.

As in Al's case, though, it doesn't seem simply to be the false belief that makes Bea's discomfort fitting. Contrast her case with Cecil's. Like Bea, Cecil wonders why Santa brought the rich kids fancier presents. Unlike Bea, Cecil has not yet done much investigating, so can think of various would-be answers which he has yet to reject. By his own lights, this question may well turn out to be answerable in a straightforward way. In some respects, Cecil's position compares epistemically unfavorably to Bea's—Bea has correctly rejected a variety of answers that Cecil has yet to investigate. But in another respect, it seems that Bea is in a worse state than Cecil is. She's continuing to wonder a question which, by her own lights, cannot have an answer. Something about this state seems incoherent, and seems to call more urgently for change than does Cecil's state. This despite the fact that both of them have the same false presupposition underlying what they're wondering. It seems, then, not just to be the false presupposition that's doing the work.

We might still be able to make sense of this without anything like QSHA, though. Perhaps we could take Bea's failed investigation to constitute evidence that what she wonders has a false presupposition,

or evidence that she has some false belief that is responsible for inability to find a good answer to her question. Cecil, by contrast, doesn't have such evidence yet. So again, QSHA gets the case right, but it's still not clear that positing it gets us anything we couldn't get with independently motivated and more familiar doxastic norms. A better test case for QSHA would be one in which it is violated by someone who is wondering a good question and who also hasn't rejected the correct answer.

### 1.3. Inadequate Hypotheses

A question content, recall, is composed of propositions that would resolve the question. But not every proposition is available to every thinker to consider at all times. Some propositions come to one whenever one needs them, but others will be available only in certain specific circumstances, and many may be totally unavailable, beyond one's capacity to grasp.<sup>9</sup> Sometimes one wonders a question with would-be answers that are unavailable to one at the time one wonders. When the true answer is among these, one's access to hypotheses is inadequate for answering the question, and one may end up violating QSHA without wondering a bad question or having incorrectly rejected a true answer.<sup>10</sup>

Consider, for example, a different way a case similar to Al's might go. Suppose that Deidre did use her keys to unlock her door, but forgot to take them out of the lock. She might then correctly rule out all of the places her keys could be inside the house, and also correctly rule out

---

<sup>9</sup>One useful way to model these restrictions is with fragmentation-based models, as in Elga and Rayo ([forthcoming](#)), Yalcin (2018), and Pérez Carballo (2016). Another way is with unawareness models used by economists and computer scientists; see Schipper (2015) for a useful overview.

<sup>10</sup>This is a special case of what Friedman (2013, p. 160) calls *abductive ignorance*, wherein someone is ignorant of what the would-be answers to a question are. Sylvain Bromberger (1965, p. 27) also discusses these kinds of cases. Here's his example:

- (1) Why do tea kettles emit a humming noise just before the water begins to boil?

Like Bromberger, "I can think of nothing, I can imagine nothing, I can conjure up nothing, I can invent nothing, I can remember nothing, that can survive confrontation with what I take to be the conditions on the right answer" (Bromberger 1965, p. 28).

every place she can think of outside of the house on the grounds that she remembers using the keys to open the door. If she were to think of [Deidre's keys are hanging from the lock of the front door], she would immediately take this to be a plausible would-be answer and go check. But for 20 minutes, suppose, she does not think of it, and—try as she might—only thinks of places she has already ruled out. For a time, then, Deidre is violating QSHA, and as predicted, appropriately feels the same kind of discomfort that Al felt in his predicament, the same pressure to get out of this state.

The fittingness of Deidre's discomfort here, though, cannot be explained by the Negative Truth Norm, or anything straightforwardly derivable from it. At the time in question, we can assume, she has no relevant false beliefs. She doesn't believe that her keys aren't outside the house, for example, just that they aren't in any of the particular outside places she has yet thought of, which is true. We can also assume that she doesn't have any strong misleading evidence that she has some relevant false belief. Deidre has a very clear memory of using her keys to unlock the door, and has been very careful to rule out locations in the house only when thoroughly searched. We can add that Deidre has found herself in similar scenarios on many occasions, and her keys have always been in a place that she just hadn't yet thought of, rather than somewhere she had incorrectly ruled out. Deidre, with good evidence, realizes that she just hasn't yet thought of the right would-be answer. Nevertheless, it is fitting for her to feel a special kind of pressure to escape her QSHA-violation. This supports positing QSHA.

One might propose instead that we explain Deidre's fitting discomfort by appeal to a positive doxastic norm, like the following.

#### **POSITIVE TRUTH NORM**

If p is true, believe p.

There are, of course, many truths that Deidre does not believe at the time in question, so she does indeed violate this norm. But she also is violating this well before she comes to violate QSHA, so this account will need to say what the relevant change was. I do not see a good way of doing this. Deidre has no more evidence than before that she

lacks some true belief. She knew all along that there's a truth about where her keys are that she does not yet believe. She also already recognized that this is a practically relevant truth, and a truth that is within her investigative reach. Running out of candidate hypotheses doesn't change any of this.

Moreover, though her keys not being in the places she has looked may raise the evidential probability of [Deidre's keys are hanging from the lock], they do not make it so likely that Deidre must believe it on pain of violating any plausible evidentialist norm. As far as she knows, there are many other places her keys could be that she has yet to think of. When she finally thinks of [Deidre's keys are hanging from the lock], she needs to go check before she should believe that it's true. And her QSHA-violation-induced discomfort is temporarily relieved as she makes her way to the door, before she has formed the belief that her keys are there. That Deidre is failing to believe the truth of [Deidre's keys are hanging from the lock] seems to be a dead end in the search for a relevant positive doxastic norm violation.

A more promising alternative for this account to look to is Deidre's failure to believe what is expressed by (2).

(2) Deidre's keys *might* be hanging from the lock.

It is controversial whether a sentence like (2) expresses a proposition. Nevertheless, it does seem to express something that one can believe. And at the time in question, it's plausible to say that Deidre does not believe what it expresses.<sup>11</sup> And so it's plausible that she would be failing to meet the requirement of the Positive Truth Norm by not believing it.<sup>12</sup> But again, this would have been so at the beginning of the investigation, when Deidre should not have felt any puzzlement or discomfort. What changed?

Could we say that Deidre's failure to find her keys in any place she can think of provides her with evidence for (2)? I don't think so. It's

---

<sup>11</sup>C.f. Yalcin (2011, pp. 313ff.).

<sup>12</sup>See Moss (2018, esp. Ch. 8) for an account that takes sentences like (2) not to express truth-conditional contents of a standard sort, but nevertheless allows belief (and failure of belief) in what they express to be subject to knowledge (and truth) norms.

not clear how to make sense of the idea of evidence for *might*-claims, and with some views about evidence and what these claims express, it cannot be done. But we do sometimes talk about something's being evidence for a *might*-claim, as when we cite the weather report as evidence for our belief that it might rain. So I am inclined to grant the alternative account that the idea makes sense.

The problem is that this doesn't seem to be what's going on in Deidre's case. If Deidre had started out thinking that her keys *must* be inside, so they can't be hanging from the lock, then I think it's plausible that her failed search of the inside spots gives her evidence that they might not be inside after all. And perhaps this would give her evidence that they might be at the office, that they might be in her car, that they might be hanging from the lock, and so on. In that case, the change that occurs when Deidre starts to violate QSHA could be that she now has evidence for (2) and other *might*-claims that she does not believe.

But this is not what happens in Deidre's case. It's true that Deidre starts by considering some restricted set of possibilities, but that's not because she thinks that her keys must be in one of those places. All along, she thinks they might be in some place that she won't easily think of, since this has happened to her before when she's lost her keys. She doesn't fail to believe (2) because she thinks there's not enough evidence for it, but simply because she hasn't come up with the proposition [Deidre's keys are hanging from the lock]. If someone were to say something along the lines of (2), Deidre would immediately grant it, not because that person's testimony gives her evidence that it is true, but simply because it drew her attention to the possibility.<sup>13</sup> A random word generator producing the word 'door' could have the same effect, even without Deidre having occult background beliefs about the insights of the generator. Insofar as Deidre had evidence for believing (2), she had that evidence all along.

I conclude that neither a positive doxastic norm like the Positive Truth Norm nor any plausible evidential variant of it can explain why it is fitting for Deidre to feel special discomfort when she comes to violate

---

<sup>13</sup>See Ciardelli, Groenendijk, and Roelofsen (2011) and Bledin and Rawlins (2020) on the role of *might*-claims in changing attentional states.

QSHA, or the rational pressure she feels to get out of this state. It seems, then, that no standard doxastic norm will explain what is going wrong in Deidre's case, but QSHA can. The need for QSHA seems even clearer in the next kind of case I want to discuss.

Sometimes, as in Deidre's case, one is unable to access adequate hypotheses for relatively shallow psychological reasons. A simple reminder would have sufficed for her to come up with a new, unrejected hypothesis. But in other cases the inability to come up with adequate hypotheses can be traced to a deeper kind of problem: one's conceptual limitations. Sometimes one lacks the concepts needed to grasp some of the would-be answers to a question one wonders. And sometimes these are the only would-be answers to the question that one has not already eliminated. QSHA-violations in cases like these often motivate conceptual innovation.<sup>14</sup>

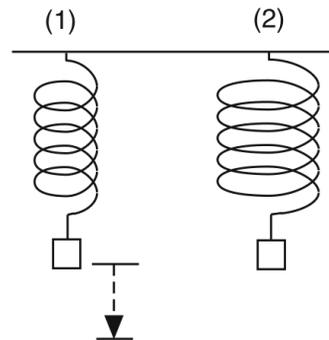
Consider the following case, from a study by John Clement (2008, p. 26). The subjects of the study reason aloud as they try to solve various puzzles. One such problem was this one:<sup>15</sup>

A weight is hung on a spring. The original spring is replaced with a spring:

- Made of the same kind of wire,
- With the same number of coils,
- But with coils that are twice as wide in diameter.

Will the spring stretch from its natural length, more, less or the same amount under the same weight? (Assume the mass of the spring is negligible compared to the mass of the weight.)

Why do you think so?



The subject S2 starts with the correct hunch that the wider spring would stretch further. But he does not know why it would do so, and so he wonders about this. He considers various answers to do with how bending relates to length, but correctly rejects any such explanation on the basis of some creative thought experiments. Soon enough, he has rejected every answer he can come up with. This leaves S2 puzzled,

<sup>14</sup>Work across several fields suggests that inability to answer some question with an original stock of concepts motivates the introduction of new ones (see Peirce (1908, §§3–4), Kuhn ([1962] 1996, Chs. 6–7), Stahl (1995), and Carey (2009, Ch. 11)).

<sup>15</sup>See also Nersessian (2008, Ch. 3) for further discussion.

without any answers to his question which he had not rejected, violating QSHA.

Naturally, this exercises S2. The interview is “peppered with expressions of frustration” as he tries to come up with a possible explanation (Clement 2008, p. 107). Indeed, he is so bothered by his failure to come up with a would-be answer to the question that he continues trying well beyond when the experimenter indicates that the allotted time for the study is over, and that S2 can stop. Unlike Deidre’s search, which would have involved a similar kind of frustration, S2 doesn’t need to just be reminded of a certain kind of would-be answer, but needs to think about properties of springs which he had previously had no idea of. In particular, he needs to think about the force from torsion stress (from twisting) on the springs, which is less for the wider spring than the narrower spring. Prior to the experiment, S2 may have lacked this concept. After a long, meandering struggle, S2 does manage to come up with thoughts along these lines, arriving at a would-be answer that he doesn’t see reasons to reject, and so is content enough to conclude the interview.

Evidently, S2 felt there was strong reason to get himself out of a QSHA-violating state. That there seems to be such reason—S2’s frustration was not pathological—can be explained by QSHA, but not by familiar doxastic norms. S2 did not reject a correct answer, nor had he reason to think he had done so. He was wondering a good question. And his failure to believe various truths didn’t change from the time before he came to violate QSHA and the time after. The problem seems simply to be his inability to conceive an adequate hypothesis. This can be explained by QSHA, but not by familiar doxastic norms.

We’ve now looked at three ways QSHA can be violated: through bad rejections, bad questions, and insufficient hypotheses. And while familiar doxastic norms can plausibly account for what’s going wrong with the first two ways of violating QSHA, I have argued that they are not able account for what’s going wrong with violations due to inadequate hypotheses.

## 2. QSHA Violations and Instrumental Norms

Perhaps we've been looking in the wrong place. Even if the cases QSHA can explain go beyond what can be captured by purely doxastic norms, we still may not need to posit QSHA. The puzzlement fittingly felt in the cases we've discussed is at least akin to the frustration one sometimes feels when one cannot achieve one's goals. We should explore, then, whether some principle of instrumental rationality can explain what we need explained.

In fact, as we'll see in a moment, it appears that QSHA itself may be derivable from an instrumental norm of inquiry—a *zetetic* norm—like the one proposed by Jane Friedman ([forthcoming](#)):

### **ZETETIC INSTRUMENTAL PRINCIPLE (ZIP)**

If one wants to figure out  $Q$ , then one ought to take the necessary means to figuring out  $Q$ .

If that's right, then QSHA may in some sense be a norm of rationality, but only as a special case of an already recognized instrumental norm. Any case it explains could already be explained by something we have independent reasons to accept.

I will argue, though, that the appearance the QSHA can be derived from a ZIP-like norm is deceptive, and that positing QSHA is needed to explain judgements about versions of our cases that instrumental norms cannot explain.

### 2.1. An Attempt to Derive QSHA

How might one go about trying to derive QSHA from an instrumental norm? As stated, QSHA is about wondering, rather than the wanting of instrumental norms. However, there's a plausible analysis of wondering that treats it as a kind of wanting:

WTK Theory of Wonder:  $S$  wonders  $Q$  iff  $S$  wants to know  $Q$ .

On this analysis, one wonders whether it's raining iff one wants to know whether it's raining, one wonders where one can buy an

Italian newspaper iff one wants to know where one can buy an Italian newspaper, and so on.<sup>16</sup>

Given the WTK Theory, it only takes a slight modification to Friedman's principle to get an instrumental principle that will apply to wondering:

ZIP\*: If one wants to know  $Q$ , then one ought to take the necessary means to know  $Q$ .

ZIP\* and the WTK Theory together imply that one ought to take the necessary means to know the answer to questions one wonders. But you cannot know the answer to  $Q$  if you reject all answers to  $Q$  you can conceive. Being able to think the true answer to a question and not rejecting that answer is a necessary means to knowing the answer to the question. So, if you wonder  $Q$ —if you want to know  $Q$ —then you had better not be in a state of rejecting all the answers to it that you can conceive. In other words: you'd better not violate QSHA. Thus it seems that with a plausible analysis of wondering, we can derive QSHA from a plausible zetetic instrumental principle.

## 2.2. Why the Derivation Fails

However, there are a couple of important problems with this attempt to derive QSHA from ZIP\*. First, it turns on an equivocation of two senses of 'want'. The second is that it gets certain cases wrong, through lumping together two different kinds of apparently unanswerable questions.

We should distinguish between what we might call all-things-considered wanting and some-things-considered wanting. Davis (1984) observes that if someone is asked to play tennis, they might be in a mental state which would allow them to truly say either (3-a) or (3-b).

<sup>16</sup>It is worth noting that Friedman herself rejects such an analysis (Friedman 2013, pp. 154–155), on grounds similar to those Whitcomb (2010) gives for rejecting a parallel analysis of curiosity. I am not persuaded, but for two reasons do not need to press the issue here. First, since instead of appealing to the WTK Theory, we could posit a version of ZIP that applies to wondering directly. Second, because I ultimately reject any straightforward ZIP-based explanation of QSHA. If this kind of explanation should be rejected on independent grounds, this would be additional support for my view.

- (3) a. I want to play, but I have to teach.  
 b. I don't want to play, since I have to teach.

The latter reports an all-things-considered desire, the former a some-things-considered desire. The example shows that one might want something some-things-considered without wanting it all-things-considered.

With this distinction in mind, let's return to the WTK Theory. It takes wondering to be a kind of wanting. But is it all-things-considered wanting or some-things-considered wanting? Answer: some-things-considered wanting, if the analysis is to remain plausible.

Suppose, for example, I wonder who the author of a paper I'm reviewing is, and find myself thinking of some plausible candidates. But I realize it would be better for me not to know—and indeed, better to not even narrow down the options—so as not to compromise the fairness of the review process.<sup>17</sup> I decide on this basis not to try to figure out the answer. In this case, it seems right to say there's a sense in which I want to know and a sense in which I don't want to know. We can report the situation either way:

- (4) a. I want to know, but I have to stay neutral.  
 b. I don't want to know, since I have to stay neutral.

I want to know some-things-considered, but not all-things-considered. But it doesn't seem right to say there's a sense in which I wonder and a sense in which I don't. In this case, it seems to me, I just do wonder who wrote the manuscript, even though I will make no effort to find out. Thus, it seems that if wondering is to be identified with a kind of wanting, it should be identified with some-things-considered wanting rather than all-things-considered wanting.

What about ZIP\*? To be plausible, it should be taken to be about all-things-considered wanting, rather than some-things-considered wanting. This is because instrumental principles about wanting in general need to be about all-things-considered wanting. It is a necessary step to playing tennis with my friend that I put on my tennis shoes.

---

<sup>17</sup>For a discussion of a variety of reasons for not wanting to know, see Ullmann-Margalit (2000).

But if, all-things-considered, I want to teach rather than play tennis, the fact that I some-things-considered want to play tennis gives me no reason to put on my tennis shoes.

The plausibility of the proposed derivation of QSHA, then, depends on equivocating between an all-things-considered and some-things-considered sense of “wants”. We need to read it as “some-things-considered wants” to maintain the WTK Theory, but as “all-things-considered wants” to maintain ZIP\*. It seems that the prospects for deriving QSHA from a zetetic instrumental principle are not so bright after all.

### 2.3. Why We Still Need QSHA

So much the worse for QSHA, one might respond. After all, we’ve only been looking at cases where there are all-things-considered desires to know, so ZIP\* should account for all the cases. Al and Deidre all-things-considered wanted to find their keys, Bea all-things-considered wanted to know why Santa brought the rich kids better presents, and S2 all-things-considered wanted to know why the wider spring would stretch further, so ZIP\* can explain why they felt pressure to come up with answers they hadn’t rejected. Why not, then, just accept ZIP\* instead of QSHA, if they really do diverge?

Modifying the examples, though, shows that we need to stick with something like QSHA. Suppose at the point they have rejected every answer they can think of, our protagonists receive credible threats that if they discover the answer to their questions, something terrible will happen. They should stop investigating, and no longer all-things-considered want to know the answer. But they might still wonder and appropriately feel puzzlement and discomfort that no answer they can think of could be right. Or suppose in the reviewing case I came to violate QSHA—every answer I can think of I reject (it couldn’t have been anyone in Group A, because X; it couldn’t be anyone in Group B, because Y; but I can’t think of any other group that the author might be a part of). Here, my other reasons might *outweigh* the pressure from the QSHA-violation to try to come up with a non-rejected answer, but that

pressure is still there. It still makes sense for me to feel puzzlement and discomfort about my epistemic position. But I don't feel anything like this in a parallel to the tennis case. The fact that I haven't put on my tennis shoes doesn't bother me in the slightest. When one goes from all-things-considered wanting to merely some-thing-considered wanting, the reasons to pursue the necessary means are not outweighed, but rather disappear. But the reason to find non-rejected answers seems to remain in QSHA-violations where one merely some-things-considered wants to know.

One limitation of ZIP\*, then, is that it doesn't get the results we want for cases of wondering that involve merely some-things-considered wanting to know. There's another kind of case where QSHA goes beyond what ZIP\* demands: cases where one has good reason to think that knowing the answer is going to be unattainable.

When one knows that what one wants is unattainable, one has no reason to pursue what would have been necessary means, had the outcome been attainable.<sup>18</sup> Suppose the only way one can win the Grand Prize is by paying a token to play, then performing some act of skill that I know I am unable to perform. In such a case, I can want the prize, yet feel no pressure at all to buy a token. We see a similar phenomenon for pursuit of means to zetetic ends. Even if the Grand Prize is a piece of knowledge that would settle some question I want to know the answer to, I still have no reason to buy the token. Even though I want to know how long humanity will continue to exist, I don't feel pressure from this desire to do preliminary research that would be required to settle the question, since I recognize that I will never get what I want.

In cases of QSHA violation, by contrast, there seems to be something going wrong, even when I have no hope of knowing the answer. I know that I will not know the answer to the question of why there is

---

<sup>18</sup>Cf. Raz (2005, pp. 6–7). On some ways of spelling out 'necessary means', everything is trivially a necessary means to the unattainable. One should probably try to avoid this by building it into the definition of 'necessary means' that the end in question is possible (as, for example, Kolodny (2018, p. 738) does). So if an end is unattainable, there are no necessary means to obtaining it, hence we can at most talk about what would have been necessary means, had the outcome been attainable.

anything at all. I recognize knowing such a thing is beyond me.<sup>19</sup> Even if I could come up with some answer that I don't reject, I would not take that to show that this is the only would-be answer, just that it's the only one I could think of. So I see that coming up with a would-be answer isn't going to settle the question. Nevertheless, not being able to think of an answer I don't reject bothers me. When I actively wonder the question, I don't just feel the hopelessness of not ever knowing the answer, but a kind of agitation at not being able to think of any good would-be answers. If we're thinking of the reasons QSHA provides in (zetetic) instrumental terms, this is an unexpected result.

A related problem is that in these cases, it seems that knowing a would-be answer is sometimes good enough, and one can end the inquiry there. If the reason comes from wanting to know the answer, this would not make sense: one should either inquire until one has an answer, or until one realizes one will probably not get an answer without more effort than its worth. But in fact reaching a would-be answer is often satisfying enough, though stopping beforehand would not be. Recall S2, who was having trouble coming up with answers to the question of why the wider spring would stretch further. Strikingly, the position at which S2 arrives does not seem to be an outright belief that his hypothesis about how the spring works is the right one. As the interview draws to a close, he says "I now feel pretty good about my understanding about the way a spring works although I realize at the same time I could be quite wrong" (Clement 2008, p. 81). It seems that what satisfied his search was an answer that is *possibly* right. Apparently, what S2 was trying to do was avoid being in a state of wondering some question without being able to conceive any answers that he had not rejected.<sup>20</sup>

---

<sup>19</sup>If you think this is too easy a question (or a confused question), feel free to substitute in your favorite alternative conundrum.

<sup>20</sup>We could try to explain this in instrumental terms by appealing to an additional basic want: wanting to know what *might* be the true answer to the question. But this just pushes the issue back. Why should one want this? It would be as if in cases like the Grand Prize scenario, one developed a non-instrumental desire for the token because it would have been a necessary means for something else one wants. Without some further explanation, this would be a strange pattern of desire formation. It would be better to understand our desires to discover would-be answers in a way that makes more sense of them than this.

What's going on in the cases we want QSHA to explain seems not to have its source in a straightforward instrumental principle like ZIP\*. QSHA can neither be derived nor replaced.

### 3. Does QSHA Demand Too Much?

QSHA explains what's going wrong in various cases, including cases that other norms we've considered cannot explain. But an overgenerating principle that generates apparently good explanations of some phenomena should be rejected nevertheless. We should check whether QSHA ever demands more of agents than is really rationally required.

Saying that one has failed to meet a rational requirement or violated a rational norm can sound like harsh criticism, equivalent to saying one is irrational. This looks a problem for my case for QSHA. Consider the cases I've adduced in its favor: Al is a little forgetful, but irrational? Bea has a false belief inculcated by otherwise trustworthy adults, but is dealing with it admirably. And Deidre and S2 also seem to be inquiring just as they should. If violation of a rational norm constitutes irrationality, then QSHA does not seem to be a rational norm.

I deny, though, the implication from violation of rational norm to blameworthy irrationality. As Parfit puts it, 'irrational', like 'wicked', "is at one end of a range of criticisms" (Parfit 1984, p. 119). Failing to abide by rational norms does mean one falls short of perfect rationality, but this is unremarkable for creatures like ourselves and is far from the more egregious errors and incoherences on which we ordinarily remark with 'irrational'.

When we fall short of a rational requirement in some specific way that we recognize, it is fitting to feel some pressure to improve, but actually making the improvements is often a process that takes time, and is sometimes a process that takes more effort than one can or should muster, given one's limitations. Indeed, staying in such a state—having intransitive preferences, having inconsistent beliefs, or, if I'm right, violating QSHA—while trying to get out of it is often praiseworthy. The force of the kind of rational norm I take QSHA to be, then, is primarily that it is fitting to feel pressure to get oneself out of a violation

of it, not that one deserves censure until one has succeeded in doing so. The cases of Al, Bea, Deidre, and S2 do not, then, seem to me to show that QSHA is too demanding to be a rational norm.

This spin on rational requirements may be too weak for some philosophers' taste, but this is not the place to enter into a debate about such general issues. Moreover, I should say that I am unsure whether such a debate would resolve itself in favor of the picture I'm working with here, as opposed to one which does not allow for excusable violations of rational requirements, or one that does without rational requirements altogether. But even if the right way to think about normativity ends up differing from this picture I'm using, I suspect arguments along the lines of those I have been developing will still show a need for something like QSHA, if not as a requirement of rationality, then as a principle of some other kind.

Thus far I've only been addressing the very cases QSHA is meant to account for. But we also need to consider whether it overgenerates requirements outside of those cases. I doubt that there is a general argument showing there is no such overgeneration, and so think potential problem cases will have to be raised and addressed one at a time. In the rest of the section I raise and give some options for addressing what I take to be the most serious potentially problematic implication of QSHA.

### 3.1. Wondering Beyond One's Abilities

That one can violate QSHA through an inability to come up with adequate hypotheses is crucial to my argument that we should posit QSHA as a norm of rationality. However, it also leads to a potential problem: that QSHA will sometimes require someone to stop wondering a good question, retract a good rejection, or else do something that they are incapable of doing.

S2 was able to fill a conceptual lacuna on the fly, and so was able to escape his QSHA violation in a satisfactory way. But this might have gone differently. Sometimes, presumably, we wonder questions whose answers aren't just inconceivable to us at the moment, but

beyond what we would ever be able to conceive. If one wonders such a question, one may have only rejected false answers, yet still violate QSHA, even though the question does have a true answer.

This does not imply that there are cases in which one cannot satisfy QSHA. QSHA is a wide-scope norm, and can always be satisfied by simply not wondering the question.<sup>21</sup> It does imply, though, that there are cases in which satisfying QSHA requires one to either do what one cannot do or else stop wondering a good question. This seems like a bad result.

### 3.2. Three Options for Responding

How is a defender of QSHA to respond? I see three plausible options.

#### *Option 1: Quietism*

Bite the bullet: admit that one shouldn't wonder what one cannot and will not be able to grasp the true answer to.

Philosophers of a certain temperament may be attracted to a stringent form of this kind of quietism, and will welcome the verdict that we should just not wonder such questions. Pascal, for example, thought that "the principal illness of man is restless curiosity about things he cannot know. And it is not as bad for him to be in error as to be curious to no purpose" (Pascal [1662] 2004, S618). I myself do not find such a view plausible, at least not without some further elaboration. Fortunately, there are a couple of points that can be made that soften the quietist response, making it less anti-inquiry and, to my mind, more plausible.

The first point is that even if we are incapable of acquiring some concept needed for grasping the true answer to a question, we are

---

<sup>21</sup>It could imply that one cannot satisfy QSHA without giving up a true belief if there are some questions that we cannot help but wonder, as Kant (e.g.) seemed to think. He starts the preface of the first edition of the *Critique of Pure Reason* as follows: "Human reason has the peculiar fate in one species of its cognitions that it is burdened with questions which it cannot dismiss, since they are given to it as problems by the nature of reason itself, but which it also cannot answer, since they transcend every capacity of human reason" (A vii–viii; see also B21). If this is the case, then QSHA might really put us in a bind; we might be incapable of satisfying it, at least without violating the Negative Truth Norm.

rarely, if ever, in a position to realize this. When we're faced with some question which may be beyond our capabilities, the sensible response will often be to say, with a character from Smullyan (1983, p. 165), "[o]f course, it *may* be beyond our power [to understand], but why should we give up so easily?" Even if QSHA is violated in such a case, there's a sense in which we should continue to wonder and its violation will be excusable.

Here's the second point. Question contents come in different fineness of grain (see §1.1), and some questions are refinements of others.<sup>22</sup> So one thing the more reluctant quietist can say is that when the answer to a question is beyond what one will be able to grasp, one should settle for wondering a question only at a coarser grain, of which the original question is a refinement. Suppose, for example, you reject all the specific physicalist and non-physicalist answers to "Why do we have conscious experiences?" Perhaps you reject all the non-physicalist answers on general grounds, but only reject the specific physicalist ones for specific reasons. At a certain grain, one which demands a fairly specific theory for the question to be resolved, you reject all the answers you can conceive. And supposing that the true theory is one you will not be able to grasp, then the quietist will say you should stop wondering the question at this refinement. But they can add that you may continue to wonder a related coarser question that could be expressed by the same words, a question that has a would-be answer that you haven't rejected: that some kind of physicalist theory or other explains why we are conscious. You can also wonder a somewhat less coarse question than that, one with an answer that you can't grasp with your current concepts, but which you would be able to grasp with some conceptual development that is within your reach. So even the quietist can encourage a certain kind of sustained inquiry, even if they say there are some good questions that you should not wonder.

### *Option 2: Dilemmas*

Bite the bullet in a different way: admit that one would be violating

---

<sup>22</sup>See Ciardelli, Groenendijk, and Roelofsen (2019, Ch. 2) for one way of thinking about this.

a norm if one continues to wonder the question, but add that one would *also* be violating some other norm if one stops wondering the question. No matter what one does in such a situation, one will violate some norm of rationality. What is best to do in a case like this may sometimes involve continued violation of QSHA, in which case we get the plausible result that one should continue to wonder in such cases, and we wouldn't have to reject QSHA to do so.

This option requires that we accept the controversial possibility of rational dilemmas<sup>23</sup> and also accept that one can be rationally required to wonder certain questions. These are not trivial commitments, but nor are they indefensible.

### *Option 3: Ability Relativity*

Grant the objection, revise the norm by relativizing to the agent's abilities.

It is a familiar idea that normative requirements of various kinds should not demand from us more than what we are able to do. As noted above, the problem of questions with answers beyond what we would be able to grasp is not quite that QSHA demands we do what we are unable to do, since we could satisfy it by giving up on the question. Nevertheless, an addition of ability-relativity in right place does avoid the problem. What this option proposes is that conceptual expansion should only be required by QSHA when the agent is able to do it. Instead of adding to QSHA only "If you are able," we would need to add "If you are able to grasp the true answer to *Q*". The new, ability-relative form of QSHA says that if you can conceive the true answer to *Q*, you must not wonder *Q* while rejecting all the answers to *Q* you can conceive. If you cannot grasp the answer to a question, this new formulation of QSHA simply doesn't apply to you.

The new formulation has an apparent problem. Read in a certain way, the modified version of QSHA is nearly trivial. If you can conceive the true answer, isn't rejecting the true answer the only way to reject

---

<sup>23</sup>On the possibility of moral, epistemic, and rational dilemmas see Marcus (1980), Conee (1982), Priest (2002), Caie (2011), Lasonen-Aarnio (2014), Schoenfield (2015), Worsnip (2018), Pryor (2018), and Hughes (2019).

all the answers to *Q* you can conceive? But then this doesn't seem to go beyond a simple Negative Truth Norm in any interesting way, and it won't account for cases like Deidre's and S2's.

The solution is to read the two "can conceive"s in different ways. The second "can conceive"—the one which also appeared in the original formulation of QSHA—designates an ability: what one is currently able to conceive. The other, which appears in the new ability-relativization clause, designates an ability to have an ability: what abilities to conceive one is able to have. Spelling out the metaphysics of the needed distinction between first-order and higher-order abilities may be tricky, but it seems that there is such a distinction, so I don't think it's a problem that the new version of QSHA needs to avail itself of it. I am not able to speak Finnish, but I think I am able to become able to speak it, whereas I am neither able to run faster than the speed of light nor able to become able to do this.<sup>24</sup> When we read the revised version of QSHA with this distinction in mind, the triviality problem disappears. Deidre and S2 aren't currently able to think the true answers to the questions they wonder, but they are able to have these abilities.

Still, we might think that the ability-relative version of QSHA lets thinkers off the hook too easily. Suppose S2 were incapable of coming up with the concept needed to make sense of the spring puzzle. Then he wouldn't be violating the new form of QSHA. But shouldn't he still feel some discomfort and puzzlement in such a situation?

One response to this is to repeat a point made to help the quietist proposal: that even if we are unable to have the ability to grasp the true answer to a question, we often won't have strong evidence that this is so. We might say, then, that the puzzlement and discomfort is fitting not because the new version of QSHA is violated, but because it seems to one that there's a good chance that is being violated.

Additionally, we might maintain a more full-fledged version of QSHA as a rational ideal, something which is supererogatory rather than required to satisfy.<sup>25</sup> We could say that if you wonder a question

---

<sup>24</sup>See Vetter (2005, pp. 135–139). Aristotle also seems to have something like this distinction in mind in his discussion of first and second potentiality in *De Anima* II.5 (417a 21ff.).

<sup>25</sup>On the idea of rational supererogation, see Benn and Bales (2019).

Q, it is ideal for you to be able to conceive an answer that you do not reject. The puzzlement one feels may point one towards a goal that may be required or may be supererogatory to satisfy, depending on one's abilities. Given one's uncertainty about one's abilities, then, it will make sense to strive to generate new hypotheses.

I've now presented three options for responding to the concern about over-demandingness. Each involves contentious commitments, so no one of them will satisfy everyone, but my hope is that most philosophers will find at least one option that seems promising to them. If so, then together they make it plausible that some form of QSHA is defensible, even if we can't be sure at this point what exact form it will take.

I am inclined, then, to maintain the posit of QSHA in some form or other. But why should there be such a norm? Where does QSHA come from?

#### 4. A Conceivability Approach

It seems that neither doxastic norms nor zetetic instrumental principles will lead to QSHA. In particular, these norms aren't able to explain what is going wrong with violations of QSHA due to inadequate hypotheses. So if QSHA is a genuine norm, what is its source? Why should there be something wrong with these kinds of QSHA-violations?

If we cannot find a plausible account of why there should be a norm like QSHA, we might rethink whether we should really accept that there is such a norm, as opposed to explaining away at least some of our judgements about such cases as confused or otherwise aberrant. In the remainder of the paper, I'd like to propose a partial source for QSHA: norms concerning what contents one *can conceive*; norms, in other words, governing what concepts one has available.

To see what I have in mind, it will be useful to start with an implausibly strong conceivability norm, to see how such an account might work in principle. We will then see whether we can find a more limited and plausible version that can do the work we need it to do.

Suppose that we were governed by the following norm:

**STRONG CONCEIVABILITY (SC)**

Be able to think all contents.

Such a norm would make it easy to see what would be going wrong with QSHA-violations due to inadequate hypotheses. For the time S2 can't think the relevant content about torsion stress, there is some content he is unable to think, so he violates SC. For the time that Deidre can't think [Deidre's keys are in the door], she violates SC. If doxastic norms can explain what's going wrong with the bad rejection and bad question kinds of cases, then it looks like SC would be sufficient to fill in the rest.<sup>26</sup> SC can't be the missing ingredient we're looking for, though, since it's too strong to be plausible, and too strong to be sensitive to what's going wrong in QSHA-violations but not in other scenarios.

It doesn't bother me that there are all sorts of contents which I cannot grasp. Nor should it. Take a recent issue of a speciality journal in some academic field I don't know much about. It will be chock full of expressions of contents which I am not able to think. But I should feel no worse about this than about the realization that there are many facts that I will never know. It would be inappropriate for me to feel here the kind of discomfort S2 or Deidre have in their QSHA-violating states. SC may be some sort of cognitive ideal, in the same way that knowing every fact may be some sort of cognitive ideal, but SC seems not to be a requirement of rationality for thinkers like us, and so doesn't explain what's going on in the QSHA-violations we're interested in. If we're going to appeal to a conceivability norm to partially explain QSHA, we'll need to look for something more constrained.

We might worry, though, that to pursue this would be to go down a garden path, due to a problem with conceivability norms in general.

---

<sup>26</sup>An alternative, conceivability-based way we might capture the bad questions cases:

**STRONG NON-DEFECTIVE CONCEIVABILITY (SNDC)**

Be able to think all and only *non-defective* contents.

Using this would involve spelling out what it takes for a content to be defective. But one plausible candidate for defectiveness is having false presuppositions. Then the explanation for Bea's violation of QSHA could be traced the question content which she wonders is, on this view, a defective content, which is one she should not be able to entertain.

Notice that SC is a norm that requires agents to have certain abilities. But is ability possession a kind of state that can be directly subject to rational evaluation? It seems odd to say that agents just should have certain abilities or others. Should I be able to dance the can-can? One can imagine circumstances that require me to *act* in various ways to make it so that I can dance the can-can, but it is hard to believe that there's some kind of rational requirement simply for agents to have such an ability. And this is not just a problem for specific abilities that not all agents need. Even the very general abilities we might expect all rational agents to have (be able to adjust one's representations on the basis of evidence, be able to act on the basis of one's preferences) don't seem like they should be thought of as rational requirements. If something lacks such abilities, it just isn't an agent. Perhaps one should act in ways to maintain these abilities, but if so, that would seem to come from other sources, not a norm that says one just must have these abilities.<sup>27</sup>

I think this is an important worry, and I will not be able to fully address it here. Nevertheless, I think finding a more constrained conceivability norm can help with this objection, too. Categorical ability norms—ones that are meant to apply to all rational agents qua rational agents—indeed seem odd.<sup>28</sup> Less odd, though, are hypothetical ability norms. If you're a dancer specializing in historical cabaret, you should be able to do the can-can. If you're a philosophy professor, you should be able to give an example of a valid argument. If you've made a promise, you should be able to keep it. It is an interesting question how to understand and explain such requirements—and there may well be different explanations for different requirements—but for now

---

<sup>27</sup>So would say anyone who agrees with the criticism of constitutivist accounts of normativity from Enoch (2006) and Enoch (2011). But even constitutivists who resist Enoch's criticism generally don't say that one simply *ought* to be an agent (or have the abilities required for being an agent). Instead they say it just is an inescapable fact that we are agents and/or one can only make sense of normative questions from an agent's standpoint (Korsgaard (2009, p. 2), Velleman (2009, p. 137), Ferrero (2009), Silverstein (2015)), which they think gives unavoidable normative force to the constitutive norms of agency. But this need not involve any normative requirement to have the abilities required for being an agent.

<sup>28</sup>This is not to say there can be no such thing, just that one would have to do some work in defending their possibility.

I just want to observe that they seem commonplace enough.

What I propose we look for, then, is a hypothetical conceivability norm that can avoid the earlier overgeneration worry. What should be the hypothetical? Given the content of QSHA, and the cases we are trying to capture, a plausible option to consider is simply that of wondering a question  $Q$ . But this alone is not enough of a constraint. Tacking it on to SC, we would get:

If you wonder a question  $Q$ , be able to think all contents.

Clearly, this will not do, since it overgenerates in basically the same way as SC—most people wonder something, after all.

We want to link the contents that one is required to conceive to the question that one wonders. But what should that link be? The natural suggestion is that the requirement is to be able to conceive some answer(s) to the question one wonders. I will call these Question-Answer Conceivability norms. Here are a few salient possible such norms:

$QAC_{\forall}$

If you wonder  $Q$ , be able to think all would-be answers to  $Q$ .

$QAC_{\top}$

If you wonder  $Q$ , be able to think a true answer to  $Q$ .

$QAC_{\exists}$

If you wonder  $Q$ , be able to think some would-be answer to  $Q$ .

We can do away with  $QAC_{\forall}$  and  $QAC_{\exists}$ . The latter is too weak. Even after she has eliminated all the answers she can conceive, Deidre can conceive various would-be answers to the question she's wondering. Namely, those which she has eliminated. So if we are trying to explain what's going wrong when Deidre is violating QSHA,  $QAC_{\exists}$  will not be sufficient.  $QAC_{\forall}$ , though, is too strong. At the beginning of her search, Deidre might realize that there are places she's never heard of that appear in some would-be answers to the question she wonders. But since she has no reason to think they are somewhere she hasn't heard of, this doesn't bother her. Exploring the would-be answers she's aware of

is a fine strategy to start with; no need to start trying to come up with unconsidered possibilities at this point. So if we were to propose that if she wonders where her keys are, Deidre must be able to conceive all would-be answers to the question of where her keys are, we would be placing a requirement on Deidre that seems not to be there.<sup>29</sup>

$QAC_{\tau}$ , however, seems more promising. On the one hand, it avoids the overgeneration cases faced by SC. The fact that I can't grasp the contents expressed in some specialty journal of a field I am unfamiliar with isn't a problem as far as  $QAC_{\tau}$  is concerned. Nor is it too weak to account for what's going wrong in the case of Deidre and S2. Both of them, at the time they are violating QSHA through inadequate hypotheses, are indeed unable to think the true answers to the questions they wonder, and so violate  $QAC_{\tau}$ .

It may seem that it's still too strong a requirement for our purposes, however, since Deidre and S2 also already violate  $QAC_{\tau}$  *before* they've begun to violate QSHA. At the beginning of their investigations, Deidre and S2 are already violating  $QAC_{\tau}$ , but at that point QSHA is unviolated, and they should not yet feel that anything is going wrong. It would be strange for them to feel pressure to come up with new thoughts at that point. This is the same problem we saw for alternatives to QSHA in §1.3.

The solution here, I think, is not to give up  $QAC_{\tau}$ , but instead say how exactly it gives rise to QSHA. It's plausible that there is a way of deriving evidence-relative norms from other norms.<sup>30</sup> If there's a norm

<sup>29</sup>Note that in certain cases, other norms may require us to think of all—or at least all relevant or reasonable—would-be answers. If we are making an eliminative or contrastive inference, such as an inference to the best explanation or to the 'only game in town', we must have reason to think we aren't overlooking better answers. See Harman (1965, p. 89), Earman (1992, Ch. 7) and Stanford (2006). One general way to do this is to coarsen the question so as to include a catchall hypothesis, but then one is left with the task of determining how likely the catchall is.

<sup>30</sup>See note 8 above. How to relate more objective normative statuses to more subjective ones (and to one's evidence about such statuses) is a much disputed issue in epistemology, ethics, and metaethics. For a variety of recent perspectives, see, e.g., Parfit (2011, Part 1), Broome (2013), Gibbons (2013), Zimmerman (2014), Wedgwood (2017), Smith (2018), Lord (2018), Sepielli (2018), Mason (2019), Weatherson (2019), Srinivasan (forthcoming), and Johnson King (ms). Many philosophers will take issue with particular way of deriving evidence-relative norms that I suggest here, but still could accept the idea of basing QSHA on  $QAC_{\tau}$ , recast in terms more to their liking.

that requires one not to cause needless suffering, then it's plausible that there is a derivative norm that requires one not to do what one's evidence entails would cause needless suffering. Here is a way of generalizing this that will suffice for our purposes: when there is a norm  $N$  that requires one not  $\phi$ , there is a norm  $N'$  that requires one not to  $\psi$ , if one's evidence entails that  $\psi$ -ing would be a case of  $\phi$ -ing.<sup>31</sup> Once they reject all the answers they can think of, violating QSHA, Deidre and S2 have enough evidence to believe they are violating  $QAC_{\top}$ . This is what we'd expect if QSHA were an evidence-relative derivative of  $QAC_{\top}$ .

Interestingly, this doesn't just work for QSHA-violations due to inadequate hypotheses, but also those due to bad questions. In wondering her bad question, Bea violates  $QAC_{\top}$ , since there's no true answer to the question for her to be able to think. More generally, bad questions do not have true answers, so nobody can conceive their true answers. And once one has rejected all the answers one can think of, as Bea does when she comes to violate QSHA, one has enough evidence to believe one is violating  $QAC_{\top}$ .<sup>32</sup>

Cases of bad rejection, though, do not need to involve violations of  $QAC_{\top}$ , nor must they give one enough evidence to believe one has violated  $QAC_{\top}$ . We can suppose that Al has often lost his keys, but whenever he finds himself stuck, he always finds his keys in one of the places he had already ruled out.  $QAC_{\top}$  alone, then, does not suffice to explain QSHA. We've already seen in §1.1, though, how we can fill in the gap. We can appeal to the Negative Truth Norm for belief, again with a way of deriving evidence-relative norms. QSHA, in my view, is the result of  $QAC_{\top}$  together with a Negative Truth Norm, and a way of deriving evidence-relative norms from other norms.<sup>33</sup>

<sup>31</sup>Ultimately, we'll likely want some more subtle way to derive evidence-relative norms, due to miners cases. For discussion see, among many others, Kolodny and MacFarlane (2010), Parfit (2011, p. 159), Carr (2015), and Muñoz and Spencer (forthcoming).

<sup>32</sup>This is not quite right, for a reason I address in footnote 33.

<sup>33</sup>Sometimes you might violate QSHA without having enough evidence to think you've violated  $QAC_{\top}$  or enough evidence to think you've violated the Negative Truth Norm, but only that you've violated at least one of them. To account for this, we either need a way of conjoining norms, or a way of deriving evidence-relative norms from sets of norms, rather than single norms. On the former approach, we'd

Truth norms have been much discussed, as have derivations of evidence-relative norms from more basic norms. What's new about QSHA, then, can be traced to a question-answer conceivability norm,  $\text{QAC}_\top$ . Why there should be such a norm is a topic that deserves further investigation.  $\text{QAC}_\top$  is not, for the reasons given in §2, simply a special case of a more general instrumental norm. So where does it come from? I'd like to conclude with a speculative suggestion.

There is a well-known dilemma, confronted in many kinds of situations, between exploring one's environment and exploiting the information one already has.<sup>34</sup> Agents typically find themselves with some information, but not all the information there is, and have the option of trying to acquire more information or instead deciding on the basis of the information they already have. Suppose, for example, that I want to decide which restaurant to go to. I already have some information about what my options and their relative merits are, but of course do not have full information. I am thus faced with a further decision: attempt to acquire more information or just pick the option that looks best given my current information. Exploring has a potential benefit: I might find a new place that's better than all the others, or find out a change of management has brought down the quality at the place I would have otherwise picked, and so on. But this exploration comes with a cost—taking effort, time, and increasing the risk of arriving too late to get a table. Agents would do well to be sensitive to these kinds of costs and benefits, so that they explore enough, but not too much.

Here is my suggestion. The explore-exploit dilemma doesn't just arise for information, but also for concepts. One finds oneself with certain concepts available: able to think certain thoughts, but not all the thoughts it's possible to think. Just as one can realize one doesn't have

---

first get a norm that says not to violate  $\text{QAC}_\top$  or the Negative Truth Norm, then we could derive QSHA as an evidence-relative norm based on that conjoined norm. On the latter approach, we say that when there are norms  $N, N', \dots$ , there is a norm that requires one not to  $\phi$  if one has enough evidence to think that  $\phi$ -ing would violate any of  $N, N', \dots$ . Being in state of QSHA-violation, one has evidence enough to think that one is violating either the Negative Truth Norm or  $\text{QAC}_\top$ .

<sup>34</sup>Going back to Thompson (1933), this has been explored in statistics (Berry and Fristedt 1985), control theory (Powell 2011, Ch. 12), machine learning (R. S. Sutton and Barto 2018, Ch. 2), biology (Stephens and Krebs 1986, Ch. 4), cognitive neuroscience (Cohen, McClure, and Yu 2007), and economics (Stigler 1961), among other fields.

all the relevant information and that one might have some relevant mistaken beliefs, one can realize that one doesn't have all the relevant concepts and that one might have some relevant defective concepts. But we have limited time and cognitive capacities, so there are costs to improving our conceptual repertoires. So we are faced with a dilemma: when shall we exploit the conceptual resources we have for use in thought and action, and when shall we 'explore' conceptual space and develop our conceptual resources? I think a question-answer conceivability norm like  $QAC_{\top}$ —and, derivatively, QSHA—may be driven by this dilemma. Not being able to conceive the true answer to a question one wonders signals a relevant problem with one's abilities to conceive, and so gives one some reason to do some conceptual exploration.

## References

- Aloni, Maria (2001). “Quantification under Conceptual Covers”. PhD thesis. Institute for Logic, Language and Computation, at the University of Amsterdam.
- Benn, Claire and Adam Bales (2019). “The Rationally Supererogatory”. In: *Mind*. doi: [doi:10.1093/mind/fzz055](https://doi.org/10.1093/mind/fzz055).
- Berry, Donald A. and Bert Fristedt (1985). *Bandit Problems: Sequential Allocation of Experiments*. Monographs on Statistics and Applied Probability. Springer.
- Bledin, Justin and Kyle Rawlins (2020). “Resistance and Resolution: Attentional Dynamics in Discourse”. In: *Journal of Semantics* 37, pp. 43–82.
- Boër, Steven E. and William G. Lycan (1975). “Knowing Who”. In: *Philosophical Studies* 28.5, pp. 299–344.
- Bromberger, Sylvain (1965). “An Approach to Explanation”. In: *Analytical philosophy*. Ed. by R. J. Butler. Vol. 2. Blackwell Oxford, pp. 72–105. Reprinted in Bromberger (1992).
- (1992). *On What We Know We Don’t Know: Explanation, Theory, Linguistics, and How Questions Shape Them*. University of Chicago Press.
- Broome, John (2013). *Rationality Through Reasoning*. Wiley Blackwell.
- Caie, Michael (2011). “Belief and Indeterminacy”. In: *Philosophical Review* 121.1, pp. 1–54.
- Carey, Susan (2009). *The Origin of Concepts*. Oxford University Press.
- Carr, Jennifer (2015). “Subjective Oughts”. In: *Ergo* 2.27, pp. 678–710.
- Ciardelli, Ivano, Jeroen Groenendijk, and Floris Roelofsen (2011). “Attention! Might in Inquisitive Semantics”. In: *Proceedings of Semantics and Linguistic Theory* 19, pp. 91–108.
- (2019). *Inquisitive Semantics*. Oxford University Press.
- Clement, John J. (2008). *Creative Model Construction in Scientists and Students: The Role of Imagery, Analogy, and Mental Simulation*. Springer.
- Cohen, Jonathan D., Samuel M. McClure, and Angela J. Yu (2007). “Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration”. In: *Philosophical Transactions of the Royal Society B* 362, pp. 933–942.

- Conee, Earl (1982). "Against Moral Dilemmas". In: *Philosophical Review* 91, pp. 87–97.
- Davis, Wayne A. (1984). "The Two Senses of Desire". In: *Philosophical Studies* 45.2, pp. 181–195.
- Earman, John (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. The MIT Press.
- Elga, Adam and Agustín Rayo (forthcoming). "Fragmentation and information access". In: *The Fragmented Mind*. Ed. by Dirk Kinderman, Andrea Onofri, and Cristina Borgoni.
- Engel, Pascal (2013). "Doxastic Correctness". In: *Proceedings of the Aristotelian Society* 87, pp. 199–216.
- Enoch, David (2006). "Agency, Shmagency: Why Normativity Won't Come from What is Constitutive of Action". In: *Philosophical Review* 115.2, pp. 169–198.
- (2011). "Shmagency Revisited". In: *New Waves in Metaethics*. Ed. by Michael Brady. Palgrave Macmillan, pp. 208–233.
- Ferrero, Luca (2009). "Constitutivism and the Inescapability of Agency". In: *Oxford Studies in Metaethics*. Ed. by Russ Shafer-Landau. Vol. 4. Oxford University Press, pp. 303–333.
- Fogal, Daniel (forthcoming). "Rational Requirements and the Primacy of Pressure". In: *Mind*.
- Friedman, Jane (2013). "Question-Directed Attitudes". In: *Philosophical Perspectives* 27, pp. 145–174.
- (forthcoming). "The Epistemic and the Zetetic". In: *The Philosophical Review*.
- Gibbons, John (2013). *The Norm of Belief*. Oxford University Press.
- Ginzburg, Jonathan (1995). "Resolving Questions, I". In: *Linguistics and Philosophy* 18, pp. 459–527.
- Groenendijk, Jeroen and Martin Stokhof (1984). "Studies on the Semantics of Questions and the Pragmatics of Answers". PhD thesis. University of Amsterdam.
- Hamblin, Charles (1958). "Questions". In: *Australasian Journal of Philosophy* 36 (3), pp. 159–168.
- (1973). "Questions in Montague English". In: *Foundations of Language* 10 (1), pp. 41–53.

- Harman, Gilbert (1965). "The Inference to the Best Explanation". In: *The Philosophical Review* 74.1, pp. 88–95.
- Hintikka, Jaakko (1978). "Answers to Questions". In: *Questions*. Ed. by Henry Hiz. D. Reidel Publishing Company, pp. 279–300.
- Hughes, Nick (2019). "Dilemmic Epistemology". In: *Synthese* 169, pp. 4059–4090.
- Johnson King, Zoë (ms). "Radical Internalism". Version 4.5, Spring 2020.
- Kant, Immanuel ([1787] 1998). *Critique of Pure Reason*. Ed. and trans. by Paul Guyer and Allen W. Wood. Cambridge University Press.
- Kelp, Cristoph (2020). "Theory of Inquiry". In: *Philosophy and Phenomenological Research*. doi: [10.1111/phpr.12719](https://doi.org/10.1111/phpr.12719).
- Kolodny, Niko (2007). "How Does Coherence Matter?" In: *Proceedings of the Aristotelian Society* CVII.3, pp. 229–263.
- (2018). "Instrumental Rationality". In: *The Oxford Handbook of Reasons and Normativity*. Ed. by Daniel Star. Oxford University Press, pp. 731–763.
- Kolodny, Niko and John MacFarlane (2010). "Ifs and Oughts". In: *The Journal of Philosophy* 107.3, pp. 115–143.
- Korsgaard, Christine M. (2009). *Self-Constitution: Agency, Identity, and Integrity*. Oxford University Press.
- Kuhn, Thomas S. ([1962] 1996). *The Structure of Scientific Revolutions*. 3rd ed. Chicago University Press.
- Lasonen-Aarnio, Maria (2014). "Higher-Order Evidence and the Limits of Defeat". In: *Philosophy and Phenomenological Research* 88.2, pp. 214–345.
- Lord, Errol (2018). *The Importance of Being Rational*. Oxford University Press.
- Marcus, Ruth Barcan (1980). "Moral Dilemmas and Consistency". In: *The Journal of Philosophy* 77, pp. 121–136.
- Mason, Elinor (2019). *Ways to be Blameworthy: Rightness, Wrongness, and Responsibility*. Oxford University Press.
- McHugh, Conor (2012). "The Truth Norm of Belief". In: *Pacific Philosophical Quarterly* 93, pp. 8–30.
- Moss, Sarah (2018). *Probabilistic Knowledge*. Oxford University Press.

- Muñoz, Daniel and Jack Spencer (forthcoming). “Knowledge of Objective ‘Oughts’: Monotonicity and the New Miners Puzzle”. In: *Philosophy and Phenomenological Research*.
- Nersessian, Nancy J. (2008). *Creating Scientific Concepts*. The MIT Press.
- Parfit, Derek (1984). *Reasons and Persons*. Oxford University Press.
- (2011). *On What Matters*. Vol. One. Oxford University Press.
- Pascal, Blaise ([1662] 2004). *Pensées*. Ed. and trans. by Roger Ariew. Hackett Publishing Company.
- Peirce, Charles S. (1908). “A Neglected Argument for the Reality of God”. In: *Hibbert Journal* 7, pp. 90–112. Reprinted in *The Essential Peirce: Selected Philosophical Writings, Volume 2 (1893–1913)*. Ed. by the Peirce Edition Project. Indiana University Press.
- Pérez Carballo, Alejandro (2016). “Structuring Logical Space”. In: *Philosophy and Phenomenological Research* 42.2, pp. 460–491.
- Powell, Warren B. (2011). *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. 2nd ed. Wiley.
- Priest, Graham (2002). “Rational Dilemmas”. In: *Analysis* 62.1, pp. 11–16.
- Pryor, James (2018). “The Merits of Incoherence”. In: 59.1, pp. 112–141.
- Raz, Joseph (2005). “The Myth of Instrumental Rationality”. In: *Journal of Ethics & Social Philosophy* 1.1, pp. 2–28.
- Schipper, Burkhard C. (2015). “Awareness”. In: *Handbook of Epistemic Logic*. Ed. by Hans van Ditmarsch et al. College Publications, pp. 77–146.
- Schoenfield, Miriam (2015). “A Dilemma for Calibrationism”. In: *Philosophy and Phenomenological Research* 91.2, pp. 425–455.
- Sepielli, Andrew (2018). “Subjective and Objective Reasons”. In: *The Oxford Handbook of Reasons and Normativity*. Ed. by Daniel Star, pp. 784–799.
- Shah, Nishi (2003). “How Truth Governs Belief”. In: *The Philosophical Review* 112.4, pp. 446–482.
- Silverstein, Matthew (2015). “The shmagency question”. In: *Philosophical Studies* 172, pp. 1127–1142.
- Smith, Holly (2018). *Making Morality Work*. Oxford University Press.
- Smithies, Declan (2012). “The Normative Role of Knowledge”. In: *Noûs* 46.2, pp. 265–288.

- Smullyan, Raymond (1983). *5000 B.C. and Other Philosophical Fantasies*. New York: St. Martin's Press.
- Srinivasan, Amia (forthcoming). "Radical Externalism". In: *The Philosophical Review*.
- Stahl, Irene (1995). "The Appropriateness of Predicate Invention as Bias Shift Operation in ILP". In: *Machine Learning* 20, pp. 97–117.
- Stanford, P. Kyle (2006). *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford University Press.
- Stephens, David W. and John R. Krebs (1986). *Foraging Theory*. Princeton University Press.
- Stigler, George J. (1961). "The Economics of Information". In: *Journal of Political Economy* 69.3, pp. 213–225.
- Sutton, Jonathan (2005). "Stick To What You Know". In: *Noûs* 39.5, pp. 359–396.
- Sutton, Richard S. and Andrew G. Barto (2018). *Reinforcement Learning: An Introduction*. 2nd ed. The MIT Press.
- Thompson, William R. (1933). "On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples". In: *Biometrika* 25.3/4, pp. 285–294.
- Ullmann-Margalit, Edna (2000). "On Not Wanting to Know". In: *Reasoning Practically*. Ed. by Edna Ullmann-Margalit. Oxford University Press, pp. 72–84.
- Van Rooy, Robert (2003). "Questioning to Resolve Decision Problems". In: *Linguistics and Philosophy* 26, pp. 727–763.
- Velleman, J. David (2009). *How We Get Along*. Cambridge University Press.
- Vetter, Barbara (2005). *Potentiality: From Dispositions to Modality*. Oxford University Press.
- Weatherson, Brian (2019). *Normative Externalism*. Oxford University Press.
- Wedgwood, Ralph (2002). "The Aim of Belief". In: *Philosophical Perspectives* 16, pp. 267–297.
- (2017). *The Value of Rationality*. Oxford University Press.
- Whitcomb, Dennis (2010). "Curiosity was Framed". In: *Philosophy and Phenomenological Research* 81.3, pp. 664–687.

- Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.
- (forthcoming). “Justifications, Excuses, and Sceptical Scenarios”. In: *The New Evil Demon: New Essays of Knowledge, Rationality and Justification*. Ed. by Julien Dutant and Daniel Dorsch. Oxford University Press.
- Worsnip, Alex (2018). “The Conflict of Evidence and Coherence”. In: *Philosophy and Phenomenological Research* 96, pp. 3–44.
- Yalcin, Seth (2011). “Nonfactualism about epistemic modality”. In: *Epistemic Modality*. Ed. by Andy Egan and Brian Weatherson. Oxford University Press, pp. 295–332.
- (2018). “Belief as Question-Sensitive”. In: *Philosophy and Phenomenological Research* 97.1, pp. 23–47.
- Zimmerman, Michael J. (2014). *Ignorance and Moral Obligation*. Oxford University Press.