



# Knowledge and inquiry—the missing key for a knowledge-based decision theory

Moritz Schulz<sup>1</sup>

Received: 8 May 2023 / Accepted: 8 September 2023 / Published online: 6 October 2023

© The Author(s) 2023

## Abstract

Fassio and Gao (2021) object to a knowledge-based decision theory on the ground that it cannot deal with unsuccessful inquiry. One way for inquiry to fail is not to know what one should know. If one's inquiry fails in this way, is a subsequent choice in any way wrong when based on one's limited actual knowledge? This paper discusses two strategies for dealing with this problem. On a first strategy, there is nothing wrong with such a choice (but something went wrong prior to one's choice). On a second strategy, there is something wrong with one's choice and a knowledge-based decision theory should require that one's decisions be based on the knowledge one should have, not merely on the knowledge one actually has.

**Keywords** Knowledge · Inquiry · Decision theory · Belief · Belief norm

## 1 Introduction

On a knowledge-based decision theory, your decisions should only be based on what you know. You should decide to perform an action if that action is favored by what you know about yourself and the world. Standard decision theory is usually formulated in terms of certain probabilities that go into a calculation of expected value. One way of tying decision theory to knowledge is by requiring these probabilities to exclude everything you know to be false and to leave open everything you do not know to be true. (Technically, one can implement this constraint by working with probabilities that are conditionalized on what one knows.)

A knowledge-based decision theory can be seen as giving a precise decision-theoretic sense to the idea that knowledge is central for action. One version of such a theory is defended by Schulz (2017) (see also Schulz, 2021a, b). In a recent paper, however, Fassio and Gao (2021:7057) compare Schulz's theory to "an unlucky slalom

---

✉ Moritz Schulz  
moritz.schulz@tu-dresden.de

<sup>1</sup> Institut für Philosophie, TU Dresden, 01062 Dresden, Germany

skier who succeeds in missing all the gates.” In this paper, I wish to focus on one specific challenge that strikes me to be particularly interesting: the challenge from *unsuccessful inquiry*, as one might call it.

Fassio and Gao describe the following case that takes its inspiration from a similar case by Gibbons (2013:179):

### MISSING KEYS

Grace is at home searching frantically for her keys. She remembers that she left them somewhere in the dining room, but doesn’t remember where. She looks in the direction of the dining table. The keys are in clear sight, staring her in the face on the otherwise empty table. Yet Grace surprisingly fails to revise any of her doxastic attitudes and continues searching for the keys in other places (Fassio and Gao, 2021:7036).

Fassio and Gao suggest that MISSING KEYS is a counterexample to a knowledge-based decision theory, for two claims can easily appear to be true:

- (1) Grace should not keep searching.<sup>1</sup>
- (2) A knowledge-based decision theory predicts that Grace should keep searching.

Before going into details, let us quickly see how both of these claims could gain support from the example. It helps to picture Grace as an otherwise normal epistemic subject without any medical conditions that may impair her vision or cognition. Her failure to believe is a one-off event that may be due to, say, an extremely unlikely chance event in her brain (that might feel as if one had been absent-minded for an instant).

Re (1): The thought would be that Grace should not keep searching because the answer to her question is right in front of her. What she should do is to form the belief that the keys are on the table, and with that belief in hand she should stop searching.

Re (2): As Grace does not form the belief that the keys are on the table, she does not come to know this fact. As a result, she knows as little about the keys as she knew before she started her search. Hence, in the light of what she knows, she has not made any progress and should keep searching. To make this a little more precise, let us assume that before entering the dining room, searching for her keys had maximal expected value when calculated based on what she knew. Assuming that her state of knowledge has not changed, her probabilities conditional on what she knows have not changed, and so looking for her keys still has maximal expected value.

In this paper, I hope to resolve MISSING KEYS on behalf of a knowledge-based decision theory. But before entering the discussion, there are a couple of preliminary remarks I would like to make.

First, in order for MISSING KEYS to be a challenge for a knowledge-based decision theory, one has to add certain assumptions to the case. In one realistic scenario, Grace will have the impression after looking at the table that the keys are not on the table, believe that the keys are not there, and search elsewhere (this is actually how Gibbons 2013:179 construed the case). But if this were so, Grace would base her decision on a false belief that does not constitute knowledge. And this would mean that Grace does not base her choice exclusively on knowledge as a knowledge-based decision

<sup>1</sup> For the purposes of this paper, I use “should” and “ought” interchangeably.

theory would require. As it stands, a knowledge-based decision theory would not face a counterexample.

For the case to have bite, one must assume that there is no relevant change in Grace's epistemic attitudes after looking. Epistemically speaking, it is as if she had never looked (in all respects relevant to her choice).

Let us observe that if nothing changes regarding Grace's knowledge, a knowledge-based decision theory will most likely recommend that she simply looks again. To see why, we may assume that Grace's original knowledge favored to start her search by looking at the table. But if her knowledge does not change after her first look, it will keep favoring to look at the table. One possible exception would result if, before her first look, it was permissible but not mandatory to start the search with the table. If this were the original situation, then after her first unsuccessful look, it will still be permissible but not mandatory to look again. Realizing that "keep searching" will most likely mean "look again" takes away some of the heat from a knowledge-based decision theory.

One possible complication results when one considers a variant of the case where it merely *seems* to Grace as if the table were empty and, as a consequence, Grace comes to know that it seems to her this way (yet she does not form the belief that the table is empty).<sup>2</sup> Based on this knowledge, Grace may decide to look elsewhere. So, instead of acting on a false belief that the table is empty, Grace would respond to something she knows to be true, namely how things seem to her. However, in this variant of the case, it is less clear that a knowledge-based decision theory will recommend to search elsewhere. After all, Grace does not form the belief that the table is empty, even though it seems to her this way. Hence, she does not exclude, and from the perspective of what she knows cannot exclude, that the keys are on the table. In this type of case—knowledge of seemings without belief—it would probably be most efficient to look again before searching elsewhere. But I grant that if knowledge of seemings makes it sufficiently likely that the keys are not on the table, and if the costs of looking elsewhere are sufficiently close to the costs of looking again, a knowledge-based decision theory could recommend to search elsewhere (but with all these details made explicit, this seems much less problematic).

On a second note, it may be worth pointing out that MISSING KEYS is not a problem specifically tailored to a knowledge-based decision theory. As a matter of fact, it is a problem for any *actualist* decision theory that describes what an agent should do in terms of attitudes the agent actually has.<sup>3</sup> It does not matter whether the relevant attitude is knowledge. One would still face the problem of MISSING KEYS if instead one thought that decisions should be made based on (justified) outright beliefs, or based on what one is certain of, or based on one's credences. As long as these attitudes and their respective epistemic status do not change when Grace looks at the dining table, there is the kind of problem Fassio and Gao wish to identify. With this in mind, I will continue to cast my discussion in terms of knowledge, but opponents of

<sup>2</sup> Thanks to an anonymous referee for pressing me on this issue.

<sup>3</sup> The distinction between actualist and non-actualist decision theories, or theories of practical rationality more generally, can also be made in the realm of epistemic rationality. Feldman and Conee (1985: 21f.) discuss whether a belief's justification depends on the evidence the agent actually has or rather on the evidence the agent should have (they answer this question in the negative).

a knowledge-based decision theory that nonetheless favor an actualist decision theory can for the most part replace knowledge with their favorite attitude they would prefer to see in its stead.

Given these preliminaries, the plan for the paper is now the following. I start by making a few remarks about inquiry. I will then sketch a positive norm for belief. The remaining sections then discuss two ways of reacting to MISSING KEYS: reject the case as a genuine counterexample or accommodate it by modifying a knowledge-based decision theory.

## 2 Inquiring into a question

In MISSING KEYS, Grace is engaged in a process of inquiry: she inquires where her keys are. So, to address our problem, it will be helpful to outline a few general features of inquiry. For our discussion, three aspects of inquiry will be important: a salient *question*, a phase of *evidence search*, and a phase of *evidence uptake*.

1. *The question.* One may follow Chris Kelp (2021:2f.) and distinguish between two kinds of inquiry: inquiry into a specific question and inquiry into a phenomenon such as the Brexit.<sup>4</sup> Grace's inquiry in MISSING KEYS is of the first kind, with the relevant question being "Where are my keys?" The goal for Grace would be to find an answer to this question.<sup>5</sup> Kelp (2021, 2014) argues that there is a normative constraint that is partly constitutive of inquiry: inquiry only ends properly when one has gained knowledge of the answer to one's question. In typical cases, one would, if Kelp is right, expect Grace to stop searching when she forms a belief as to where the keys are that is internally indistinguishable from a state of knowledge.

Quite plausibly, there is a link between a knowledge-based decision theory and the idea that knowledge is the goal of inquiry. It is not a far-fetched assumption that the larger role of inquiry is to provide information that has the right pedigree to be relied upon in decision-making. If this is right, knowledge may be the goal of inquiry because knowledge is what good decisions should be based on (see Schulz forthcoming).

2. *Search for evidence.* When inquiry starts, there are a lot of actions one can take to search for evidence. One can walk somewhere. Look around. Focus on a specific aspect of what is in one's visual field. With an eye towards MISSING KEYS, one can, for example, walk into the dining area, look in the direction of the table, and focus on whether there are keys on the table. All these activities are within one's voluntary control. One can decide to walk into the dining area, decide to look in the direction of the table, and decide to focus on whether there are any keys on the table.

These evidence searching activities are directed towards an epistemic goal, namely finding an answer to the salient question of one's inquiry. But at least *prima facie* they are nevertheless within the realm of practical reasons, broadly construed to include

<sup>4</sup> It may well be possible to reduce inquiry into a phenomenon to inquiry into a question or set of questions (consider "What is the nature and cause of this phenomenon?").

<sup>5</sup> For more on inquisitive attitudes in relation to inquiry, see Friedman (2013).

everything with a world-to-mind direction of fit.<sup>6</sup> Grace's decision to go to the dining area is informed by her remembering to have left them there and it is motivated by her wish to find the keys. Given that she remembers to have left the keys in the dining area, and given her wish to find them, the best she can do is first to go there and then to start looking where they most likely are.

All these decisions right up to the point when Grace focuses on the surface of the dining table pose no obstacle for a knowledge-based decision theory. Up to this point, inquiry is a process that is constituted by actions that are under one's voluntary control. This aspect of inquiry is therefore best seen as being governed by standards of practical rationality. A knowledge-based decision theory is one way of articulating such a standard. A decision to first search the dining area would be what Grace should do if the corresponding action maximizes expected value in the light of what Grace knows about the keys and what she rightfully wishes to achieve.

The phase of evidence search ends, on the present construal, when one has successfully positioned oneself in front of a piece of evidence, right before one starts forming an epistemic attitude in response to the evidence.

3. *Evidence uptake.* From a certain point onwards, inquiry no longer involves an action directed at where one is most likely to find an answer to one's question. At a certain point, inquiry involves the uptake of whatever one sees or hears. So, when Grace focusses her eyes on the surface of the dining table and attends to whether her keys are there, she should unleash her belief-forming methods to get an update on where she stands with her question.

An update on where one stands with one's question can in general result in a belief, in a disbelief, or in suspension of judgment. Thus, Grace may (a) come to believe that the keys are on the table, close her inquiry and grab them, or (b) come to believe that the keys are not on the table and start to look elsewhere, or (c) not believe either. The latter option can be instantiated in various ways: one may refrain from forming a belief because one forms no attitude at all (as in MISSING KEYS), or one may not yet form a belief when one thinks that the keys are probably not on the table but one cannot exclude that they are, say, behind a vase on the table, or one may deliberately suspend judgment because one worries that something could have been off with one's attempt of information uptake.

The phase of evidence uptake no longer belongs to the realm of practical reasons.<sup>7</sup> The rough and ready explanation for this is that evidence uptake is no longer an action (though the act of initiating such an update may well be). Unless one thinks that beliefs can permissively be influenced by practical reasons (as pragmatists do), one will think that the uptake of evidence has a purely world-to-mind direction of fit (one tries to fit one's beliefs to the world). So, one may decide where to look, how long to look, how much to concentrate, on which aspects to focus, etc., but at a certain point, these choices run out and proper inquiry involves taking in whatever one finds.

---

<sup>6</sup> On the face of it, this claim seems incompatible with norms of inquiry being epistemic, as argued for by Jane Friedman (2020: 512). But here I use "practical" in a broad sense that includes everything pertaining to what one should do. And this is nothing Friedman wishes to dispute (cf. Friedman 2020:fn. 15).

<sup>7</sup> I here wish to sidestep the question of voluntarism about belief. See Feldman (2008) for discussion in relation to norms of belief.

With this in mind, it is clear that a decision theory can be expected to govern the search for evidence. But it does not govern the uptake of evidence, for this process does not involve a decision. As it is natural to think that the search for evidence is governed by norms of practical rationality, it is equally natural to assume that the uptake of evidence is governed by norms of theoretical (or epistemic) rationality.

### 3 You should have seen that!

If a decision theory sets the normative standard for the practical aspects of inquiry, it is, of course, a natural question what governs the theoretical aspects of inquiry. Are there norms that say what we should believe in certain situations?

This is the question of whether there are *positive norms* for belief. Negative norms for belief tell us which beliefs to avoid. Candidate negative norms include: avoid false beliefs, avoid unjustified beliefs, avoid incoherent beliefs, and, most relevant in the present context, avoid beliefs that do not constitute knowledge. A positive norm for belief, on the other hand, would say which beliefs to have. For a start, one would naturally say that in MISSING KEYS *Grace should believe that the keys are on the table*.<sup>8</sup> On the face of it, this is a positive normative claim about a belief for Grace to have. As a working hypothesis, I will therefore assume that there are positive norms for belief. But I leave open whether positive norms for belief always have a practical-cum-epistemic source (see Nelson 2010) or whether there are also purely epistemic norms for belief that are positive (see Simion 2023).

In his discussion of positive norms for belief in a knowledge-first framework, Ichikawa (2022: 18) considers the following norm (I simplify Ichikawa's label):

(QP) If one considers or should consider the question whether  $P$ , and is in a position to know that  $P$ , one ought to believe that  $P$ .

The basic thought behind this norm is that in favorable epistemic situations one should believe those propositions that are answers to a question one has or should have. According to (QP), a situation is favorable when it puts one in a position to know.<sup>9</sup> In this way, a positive norm for belief would be linked to the question-directed attitudes (in the sense of Friedman 2013) of the agent.

If we apply the norm to Grace in MISSING KEYS, then all conditions for the norm to apply are satisfied: Grace considers the question where her keys are, it also seems right to say that she should do so, and she is in a position to know that the keys are on

<sup>8</sup> Gibbons (2013:179) actually starts from the observation that Grace ought to know that the keys are on the table. But as every way of knowing is a way of believing, a positive norm on knowledge seems to imply a positive norm on belief. This raises the interesting question whether positive norms on belief may be grounded in positive norms on knowledge (we ought to believe certain things because we ought to know certain things).

<sup>9</sup> Jenkins Ichikawa considers (QP) on behalf of proponents of a knowledge-first epistemology. He is ultimately skeptical whether such a norm sits well with a knowledge-first framework. For my purposes, all that matters is whether a norm like (QP) is correct—if this implies that one has to scale back one's knowledge-first ambitions, so be it.

the table. So, we would get from (QP) the verdict that Grace should believe that the keys are on the table.<sup>10</sup>

There is a bit of unclarity about the disjunctive first part of (QP)'s antecedent, as Ichikawa (2022: 18, fn. 44) notices himself. If we take the disjunction literally, then if someone considers a question even though they should not, the disjunction would still be true and they should have the corresponding belief (provided they are in a position to know).<sup>11</sup> Conversely, if one does not consider a question even though one should, then the disjunction would also be true and they should have the corresponding belief (provided again they are in a position to know). If one finds any of these implications doubtful, one may retreat to a weaker norm by strengthening the norm's antecedent from a disjunction to a conjunction:

(QP\*) If one considers and should consider the question whether  $P$ , and is in a position to know that  $P$ , one ought to believe that  $P$ .

The norm may be weaker than it could be.<sup>12</sup> But for present purposes, we can leave it open if and how it can be strengthened.

As it stands, the norm may not yet get it fully right. One source of difficulty is that there is a temporal aspect to this norm. In its present form, one could interpret the norm in such a way that one ought to believe  $P$  at the same time at which one ought to consider the question of whether  $P$ . This might be slightly too demanding, for making good on what one is in a position to know takes some time, if only very little. For instance, one may have to direct one's eyes, focus on a certain aspect of what is in one's visual field, and process the information that presents itself. One way of dealing with this issue would be to merely require the agent to believe the proposition in question at a slightly later time.

However, once one acknowledges a temporal gap between the antecedent conditions and a potential duty to believe, there is also room for unwanted interventions, either with practical or with epistemic impact. For one, it becomes possible that one should stop considering the question because of an unexpected event of greater practical importance (think of modifying MISSING KEYS so that there is suddenly smoke coming from the kitchen). For another, it becomes possible that one is presented with a defeater that robs one of being in a position to know (think of adding to MISSING

<sup>10</sup> What we do not get, however, is that Grace's duty to believe is epistemic in the narrow sense of only being grounded in purely epistemic states. A norm such as (QP) cites at least one non-epistemic factor, namely the question to be considered. For Nelson (2010), this would mean that the duty is not epistemic. But Ichikawa (2022: 18f.) rightly points out, in my mind, that there is also a legitimate broader notion of an "epistemic norm" on which (QP) would come out as epistemic.

<sup>11</sup> Situations of this kind could arise if one considers a question but considering another question would be much more urgent. They could also arise as the result of intrusions of privacy (reading someone else's diary). In such situations, a belief formed in response to a question one should not have considered may be fine in an epistemic sense (because it is well adjusted to the evidence). At the same time, it may be a belief one all-things-considered should not have because one should not even have formed an epistemic attitude with respect to the question under consideration.

<sup>12</sup> For instance, the norm does not apply to situations in which one is permitted to consider a question without having a duty to do so. In such situations, one may still have a duty to believe if one decides to answer the question and one is in a position to know the answer.



KEYS that a scientist jumps into the room and (falsely) asserts that Grace has been given a drug that makes it impossible for her to see any metal objects).<sup>13</sup>

I have one further and very minor concern: it seems possible to me that one could have a duty to consider a question without having a duty to find an answer to this question. If one should consider a question, does this strictly require that one should inquire into answering it? For this reason, but primarily in order to make the link with inquiry explicit, I am going to substitute “inquiring into a question” for “considering a question.”

As an attempt to work in all the issues just identified, I propose the following formulation (“IB” for “from inquiry to belief”):

(IB) If one inquires, and ought to inquire, into the question of whether  $P$ , and provided one is in a position to know that  $P$ , and assuming further there is no relevant change either to one’s practical or to one’s epistemic situation, then one ought to believe that  $P$  at an adequate later time.

This norm strikes me as a good first pass at a positive norm for belief. In its present form, the norm cites both a practical (the questions one should answer) as well as an epistemic (what one is in a position to know) source. For now, I must leave open whether the practical and the epistemic aspect of the norm can be disentangled. Is there a positive norm for belief that refers in its antecedent conditions only to purely epistemic matters (as Simion 2023 proposes)? Or is a duty to believe something always partly grounded in practical reasons regarding the questions we should answer (as Nelson 2010 objects)? While I must leave this business unfinished, (IB) has enough substance to discuss whether MISSING KEYS is a counterexample to a knowledge-based decision theory.

## 4 Reject

In the previous section, we have seen that it is possible to criticize Grace in MISSING KEYS for not having formed the belief that the keys are on the table. There is a positive norm for belief that applies in MISSING KEYS. Moreover, we can even say that Grace makes an *epistemic mistake*, for she does everything right up to the point where merely an uptake of the evidence would be required.<sup>14</sup>

What we have not yet seen is what this means vis-à-vis the question of whether MISSING KEYS constitutes a counterexample to a knowledge-based decision theory. What I hope to show in the present as well in the upcoming section is that our brief sketch of a theory of inquiry generates quite a bit of wiggle room for a knowledge-based decision theory. It makes it possible to reject the potential counterexample,

<sup>13</sup> A similar problem is familiar from the debate about knowledge’s closure (see Hawthorne 2005). If I competently deduce a conclusion from known premises, do I have knowledge of the conclusion? Not necessarily, for while I am reaching the conclusion, I might be presented with a defeater against the conclusion.

<sup>14</sup> I assume that one can fail to comply with a norm such as (IB) either by making a practical mistake or by making an epistemic mistake. The reason why this is possible is that (IB) cites both a practical as well as an epistemic condition. To illustrate, one may fail to believe the relevant proposition because one failed to even look at the table (a practical mistake in the phase of evidence search) or one may fail to believe because one’s belief forming methods misfire (an epistemic mistake in the phase of evidence uptake).



as I will explain in this section (call this approach “REJECT”). But it also makes it possible to accommodate the counterexample in terms of a slight modification of a knowledge-based decision theory, as I will explain in the next section (call this approach “ACCOMMODATE”).

The problem for a knowledge-based decision theory regarding MISSING KEYS was that such a theory predicts that Grace should keep searching. The reason being that Grace, after failing to uptake the evidence presented to her, does not know where the keys are. As a matter of fact, her total epistemic state is in all relevant respects the same as it was at the beginning of her search. As we noted above, the most likely way in which Grace should keep searching according to a knowledge-based decision theory would be by looking again. This prediction would be born out if it was prior to her search most likely that the keys are on the dining table or if the fact that she still stands in front of the dining table would make it most efficient to continue her search by looking again (recall that she does in no way believe that the keys are not on the table).

Here is a consideration in favor of the initially problematic prediction. Yes, Grace did not properly form her beliefs in response to her environment. But after she completed the phase of evidence uptake at which she made a mistake, she does not know where the keys are. And without any such knowledge, there is nothing else she can do than to keep searching. And if pressed why she kept searching, it seems fine for her to respond “Well, I did not know where the keys were. What other choice did I have than to keep searching?”

On this line of defense, the counterexample could be rejected. But one could still accommodate the intuition that Grace does something wrong in MISSING KEYS. However, her mistake would not lie in the subsequent choice to keep searching but rather in the prior formation of her beliefs that preceded her choice. As explained above, she should have come to believe that the keys were on the table. So she did make a mistake. But after this mistake was made, her choice to keep searching was not another mistake. Rather, it was the rational thing to do in the epistemic situation she found herself in that was, due to a past mistake, inferior to an epistemic situation she could have been in.

One can try to bolster the presently envisaged response to MISSING KEYS by considering a variation of MISSING KEYS that lets some time pass between the failure to properly uptake the evidence and the time of Grace’s choice. Let us consider this variation:

### MISSING KEYS #2

Grace is at home searching for the keys to the summer house. She remembers that she left them somewhere in the dining room, but doesn’t remember where. She looks in the direction of the dining table. The keys are in clear sight, staring her in the face on the otherwise empty table. Yet Grace surprisingly fails to revise any of her doxastic attitudes. She then gets an important call. A week later, she is about to drive to the summer house and really needs the keys. She starts searching again.

It seems very clear to me that in MISSING KEYS #2, Grace's choice to start searching again is less objectionable. If this is correct, then one may suspect that the original case seems problematic because the time of the choice is almost identical to the time at which Grace made her epistemic mistake. Given this close temporal proximity, one may be more inclined to blame Grace for her choice even though in fact her choice is fine but she makes a related but different kind of mistake at almost the same time.

With all this in hand, we have the outline of a defense of a knowledge-based decision theory against MISSING KEYS. There clearly is something Grace does wrong: she should have come to believe that the keys are on the dining table. So we can keep the intuition that Grace is to blame for something. But what she is to blame for would not be the subsequent choice she makes in the light of her limited knowledge, but rather a prior mistake in the uptake of evidence.

## 5 Accommodate

My preferred choice would be to stick to a clear-cut defense of a knowledge-based decision theory along the lines of REJECT. Nevertheless, I anticipate that some readers might feel a residue of doubt about REJECT (myself included). So in this section I briefly discuss whether a knowledge-based decision theory could also be modified to accommodate Fassio and Gao's verdict about MISSING KEYS.

Why might one think that Grace not only makes an epistemic mistake, but that her subsequent choice is to blame as well? A possible explanation might go like this. In our decisions, we rely on our beliefs. According to a knowledge-based decision theory, this kind of reliance is normatively proper only if these beliefs constitute knowledge. But now there might also be a sense in which the absence of certain beliefs that would be relevant to one's choice should be normatively proper. On this line of thought, one would ask of the beliefs one has that they constitute knowledge, and one would ask of relevant beliefs one does not have that they do not violate a positive norm for belief. All in all, such a decision-theory would both require the presence and the absence of our beliefs to be normatively proper. In sum, the proposal would be that proper decision-making is not to be based on the knowledge one actually has but on the knowledge one should have at the time of choice.

In line with this proposal, one may envisage a decision theory that would require two things of rational agents. First, rational agents would be required to only have beliefs that constitute knowledge (a negative norm on belief). And second, they would be required to have all the beliefs that they should have according to a positive norm on belief. If we follow (IB) from Sect. 3, our preliminary take on a positive norm, rational agents would be required to believe certain relevant propositions they are in a position to know at (or shortly before) the time of their choice.<sup>15</sup>

Fassio and Gao (2021: 7038f.) consider a reaction along similar lines. They object to this strategy by pointing to the fact that being in a position to know is not closed under logical consequence, in particular, it is not closed under conjunction agglomeration.

<sup>15</sup> Let me emphasize that regimenting the notion of knowledge-one-should-have in this way usually leaves the agent with substantial uncertainty. For example, before starting her search, Grace will be uncertain as to where her keys are because she is not yet in a position to know where they are.

One may be in a position to know  $p$  and one may also be in a position to know  $q$ , but one may be in no position to know  $p \& q$ . The standard example that illustrates this is one where  $p$  is of the form “I do not know  $r$ ” and  $q$  is simply  $r$ . For example, I may be in a position to know that there are five plants in my office. But as I currently do not know this, I may also be in a position to know that I currently do not know that there are five plants in my office. If we now consider the conjunction, “There are five plants in my office but I do not know that there are five plants in my office,” it is clear that I am not in a position to know this (it is a Fitch-paradoxical sentence).

Now, this observation becomes a problem for decision theory when one attempts to define the domain for a probability function to be used for calculating expected values. The standard way to delineate such a domain in a knowledge-based framework would be to consider the conjunction of everything one knows. But one cannot simply do the same with the notion of being in a position to know, for one would end up with a proposition that is in a certain sense incoherent, for it would be something the agent could not possibly know. In our example above, one would require of an agent to act as if they were certain that there are five plants in the office while at the same time also act as if they were certain that they do not know this. This would be so because the probability function relevant to one’s decisions would be conditionalized on “There are five plants in my office” as well as on “I do not know this.”

Alternatively, one could consider the strongest proposition one is in a position to know. Relying on a proposition of this kind would avoid the above problem. But it then turns out that there is more than one such proposition (e.g., one including  $p$  and one including  $q$  in the schematic example above). Which one is to be used in one’s decision-making? Fassio and Gao show skepticism that a good answer to this question can be found.

Before responding, let me make two preliminary remarks. First, there is a problem for almost everybody here. If one agrees that the domain for proper decision-making is partly constituted by merely potential information, everyone must have an answer to the question of which information has to be taken into account and which can legitimately be ignored.

Secondly, on a very general level, it is unlikely that we are faced with an unsurmountable problem. It would be quite surprising if the very idea of *things you should know at the time of your choice* would be incoherent. It may well be that certain proposals of how to analyze this idea are incoherent (one can read Fassio and Gao’s objection as a refutation of the proposal that we should characterize what we should know in terms of everything we are in a position to know). But this does not mean that the concept itself is incoherent. It may be that we simply do not yet know how to best analyze it or it may be that it resists a satisfactorily informative analysis. Be that as it may, it would be surprising if one could not sensibly speak of things one should know at the time of one’s choice. And this is all one needs to set up a rudimentary solution along the lines of ACCOMMODATE.

Now, regarding Fassio and Gao’s more specific objection, one may observe that the picture sketched at the beginning of this section in no way implies that what one should know coincides with everything one is in a position to know. One can agree with Fassio and Gao that a domain for one’s choices would be too large if it were to include everything one is in a position to know.

When one looks at the positive norm for belief discussed in the previous section, one immediately notices that it makes use of a question. The idea was that inquiry is structured by salient questions, and so what one should believe is partly constrained by the questions one can and should answer. With this in mind, what one should know at the time of choice will be constrained by which answers to salient questions one is in a position to know. For example, when Grace in *MISSING KEYS* decides to continue her search, it was a salient question where the keys were and she was in a position to answer that question. So, one may naturally conclude that Grace should know at the time of her choice that the keys are on the table.

We may note that in *MISSING KEYS* it is not a salient question for Grace whether she knows where the keys are. The question she should answer by the time of her choice is where her keys are. So, by the time of her choice, she should know that the keys are on the table. It is not the case that she should also know that she does not know that the keys are on the table.<sup>16</sup> Rather, because she should come to know that the keys are on the table, she should not believe that she does not know this. Thus, if what one should know is relativized to questions whose answers one is in a position to know, it is less likely that there will be an issue with closure.<sup>17</sup>

Although what I have said so far is still a far cry from a comprehensive analysis of what one should know, I think enough has been said to indicate that the question-sensitivity of inquiry could shed some light on what one should know at the time of one's choice. Generalizing the present idea a little further, we may say: one should know those things one is in a position to know that constitute answers to questions salient to one's choice.

## 6 Conclusion

When we fail to know what we should know, is a decision wrong when it is based on our actual knowledge alone? One way to deny this question would be to say that there is nothing wrong with a choice that is licensed by one's actual knowledge. Although one may have made a mistake by not acquiring all the knowledge one should have acquired, subsequent reliance on the limited knowledge one did acquire does not constitute another mistake.

Yet a knowledge-based decision theory can also embrace the thought that proper decision-making is not to be based on the knowledge one has but rather on the knowledge one should have. A positive norm on belief can inform our understanding of what we should know. According to one such norm, what we should know partly depends on the questions we should answer. These are the questions that structure inquiry. But what we should know also depends on what we are in a position to know. Taken

<sup>16</sup> Interestingly, this suggests that sometimes knowledge we actually have is not knowledge we should have. We could imagine a variant of the case in which Grace knows that she does not know where the keys are even at the time of her choice. But this is knowledge she should not have. She should know that the keys are on the table. And so she should not be in any state incompatible with this knowledge. But knowing that the keys are on the table implies that she does not know that she does not know that the keys are on the table (for "Grace knows that she does not know that the keys are on the table" implies, by the factivity of knowledge, "Grace does not know that the keys are on the table").

<sup>17</sup> For an in-depth discussion of the relation between questions and closure (of belief), see Hoek (2022).

together, what we should know, and should believe, are those propositions that (a) we are in a position to know and (b) constitute answers to questions that we should try to answer.

Whichever way one prefers, the prospects for a knowledge-based decision theory are better than they may initially seem. Unlike a slalom skier who misses all the gates, it is a theory that can solve the problem of unsuccessful inquiry, though it may still remain open how much luck it will have further down the slope.

**Acknowledgements** I would like to thank two anonymous referees and the editors for their very helpful comments. An earlier version of this paper was discussed in a research seminar on “Knowledge and Decision” at the TU Dresden. Many thanks to all its participants for extremely valuable feedback!

**Funding** Open Access funding enabled and organized by Projekt DEAL.

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

## References

- Fassio, D., & Gao, J. (2021). Do we really need a knowledge-based decision theory? *Synthese*, 199, 7031–59.
- Feldman, R. (2008). Modest deontologism in epistemology. *Synthese*, 161, 339–355.
- Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48, 15–34.
- Friedman, J. (2013). Question-directed attitudes. *Philosophical Perspectives*, 27, 145–74.
- Friedman, J. (2020). The epistemic and the zetetic. *The Philosophical Review*, 129, 500–36.
- Gibbons, J. (2013). *The norm of belief*. Oxford: Oxford University Press.
- Hawthorne, J. (2005). The case for closure. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 26–43). Oxford: Blackwell.
- Hoek, D. (2022). Questions in action. *Journal of Philosophy*, 119, 113–43.
- Ichikawa, J.J. (2022). You ought to have known: positive epistemic norms in a knowledge-first framework. *Synthese* online first, pp. 1–23.
- Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51, 227–32.
- Kelp, C. (2021). *Inquiry, knowledge, and understanding*. Oxford: Oxford University Press.
- Nelson, M. T. (2010). We have no positive epistemic duties. *Mind*, 119, 83–102.
- Schulz, M. (2017). Decisions and higher-order knowledge. *Noûs*, 51, 463–83.
- Schulz, M. (2021). Partial reliance. *Canadian Journal of Philosophy*, 51, 436–51.
- Schulz, M. (2021b). Practical reasoning and degrees of outright belief. In Heil, R., Koscholke, J., Rich, P., & Schulz, M. (Eds.): *Synthese topical collection “knowledge and decision”*, vol. 199, pp. 8069–8090.
- Schulz, M. (forthcoming). The point of knowledge – Is to make good decisions! In Logins, A., & Vollet, J.-H. (Eds.): *Putting knowledge to work: New directions for knowledge-first epistemology*, Oxford: Oxford University Press.
- Simion, M. (2023). Resistance to evidence and the duty to believe. *Philosophy and Phenomenological Research*, pp. 1–13.

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