



Practical reasoning and degrees of outright belief

Moritz Schulz¹

Received: 1 June 2020 / Accepted: 3 April 2021 / Published online: 29 April 2021

© The Author(s) 2021

Abstract

According to a suggestion by Williamson (Knowledge and its limits, Oxford University Press, 2000, p. 99), outright belief comes in degrees: one has a high/low degree of belief iff one is willing to rely on the content of one's belief in high/low-stakes practical reasoning. This paper develops an *epistemic norm* for degrees of outright belief so construed. Starting from the assumption that outright belief aims at knowledge, it is argued that degrees of belief aim at various levels of *strong knowledge*, that is, knowledge which satisfies particularly high epistemic standards. This account is contrasted with and shown to be superior to an alternative proposal according to which higher degrees of outright belief aim at higher-order knowledge. In an “Appendix”, it is indicated that the logic of degrees of outright belief is closely linked to *ranking theory*.

Keywords Belief · Degrees of belief · Practical reasoning · Knowledge-first epistemology · T. Williamson

1 Belief and practical reasoning

We believe many things. We rely on our beliefs when it is time for action. Yet we rely on some of our beliefs more than others. Belief is a category in which practically significant distinctions can still be made.

To fix ideas, consider the following two variants of practical reasoning:

Bar vs. Airport

Ada is about to go to a local bar where they sometimes want to see an ID card. She believes that she has her ID in her purse. She relies on this belief when deciding to leave without double-checking. She also believes that she is well.

This article is part of the topical collection on Knowledge and Decision, edited by Moritz Schulz, Roman Heil, Jakob Koscholke, and Patricia Rich.

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

¹ Institut für Philosophie, TU Dresden, Dresden, Germany

Although this belief does not necessarily play a very active role in her reasoning, she relies on it as well. If she did not believe herself to be well, she would stay at home. Now imagine Ada in a different context having the same beliefs. This time she is about to go to the airport. She does not rely on her belief that the ID is in her purse. Rather, she checks again. But she does rely on her belief that she is well. If she did not believe herself to be well, she would not go on her flight.¹

Ada has the same two beliefs in both contexts: that her ID is in her purse and that she is well.² But these beliefs are of different strengths: her belief that she is well is stronger than her belief that her ID is in her purse. This difference shows in her decision-making. She is disposed to rely on one of the two beliefs more than the other.

We could describe the situation in slightly more theory-laden ways. Described in this way, Ada has two beliefs which differ in *degree*. The degree of one belief is higher than the degree of the other. In the lower-stakes scenario (bar), she relies on both beliefs. But in the higher-stakes situation (airport), she relies only on the belief of higher degree. That she double-checks her ID can be seen as an attempt to boost the strength of this belief so that it can be relied upon in solving her decision problem.

In the light of such cases, one may conjecture that degrees of belief are closely linked to levels of stakes. The more that is at stake, the higher a degree of belief must be for it to be relied upon in practical reasoning. Timothy Williamson has articulated a position along such lines:

What is the difference between believing p outright and assigning p a high subjective probability? Intuitively, one believes p outright when one is willing to use p as a premise in practical reasoning. Thus one may assign p a high subjective probability without believing p outright, if the corresponding premise in one's practical reasoning is just that p is highly probable on one's evidence, not p itself. Outright belief still comes in degrees, for one may be willing to use p as a premise in practical reasoning only when the stakes are sufficiently low. Nevertheless, one's degree of outright belief in p is not in general to be equated with one's subjective probability for p ; one's subjective probability can vary while one's degree of outright belief remains zero. Since using p as a premise in practical reasoning is relying on p , we can think of one's degree of outright belief in p as the degree to which one relies on p . [Williamson (2000, p. 99)]

We may extract from this two hypotheses about belief:

Outright Belief. One believes p outright iff one is willing to use p as a premise in practical reasoning as long as the stakes are low.

¹ This story draws on the “Airport Case” in Cohen (1999, p. 58f.) and the “Train Cases” discussed in Fantl and McGrath (2002), but it is used here to illustrate aspects of belief rather than aspects of knowledge or justification.

² Some, such as Nagel (2010) and Weatherson (2005), disagree with this verdict. Ada may not count as believing that the ID is in her purse when the stakes are high, either because of psychological factors (Nagel, 2010) or because of the very fact that she is not prepared to rely on the proposition in question (Weatherson, 2005). Ross and Schroeder (2014, p. 271) offer something of a compromise by distinguishing two senses of ‘belief’: Ada would count as *having the belief* but she would not count as *occurently believing*. For the purposes of this paper, I shall assume that there is at least one sense in which Ada has the same beliefs in both scenarios.

Degrees of Outright Belief. One believes p outright to a high/low degree iff one is willing to use p as a premise in practical reasoning when the stakes are high or low/only when the stakes are low.

For starters, some preliminary remarks are in order.

First, I do not take either of the two hypotheses to *define* outright belief or degrees of belief. For all I am assuming, the two hypotheses might merely capture the (*functional*) role belief has for practical reasoning.³

Second, the hypotheses are about *outright* belief. They are not about credences or subjective probabilities, as Williamson points out.⁴ This means that degrees of outright belief should not be equated with more familiar *degrees of partial belief* also known as *credences* or *subjective probabilities*. Degrees of outright belief track differences within the class of outright beliefs. If the degree of outright belief is zero, this would merely mean that the proposition is not believed outright. Hence, it could in principle be assigned any subjective probability (cf. the Williamson quote above).

Third, it is legitimate to question how to understand the notion of *being willing to use a proposition as a premise in practical reasoning*. Drawing on Ross and Schroeder (2014), a willingness to use a proposition as a premise can be taken to consist in a certain kind of *disposition*. This disposition is typically triggered when the proposition becomes *relevant* to a decision problem at hand and it manifests itself by assigning the proposition a certain status during one's reasoning process. For present purposes, I take practical reasoning to be a goal-oriented activity which aims at solving a particular decision problem. This activity typically ends when a decision is made.

Now, premises of practical reasoning serve, on one salient way of construing them, to impose limits on the range of possibilities to be reckoned with in one's decision-making.⁵ If p is used as a premise in practical reasoning, then no possibility incompatible with p is taken into consideration. For example, if I adopt the proposition that my lottery ticket lost as a premise in my reasoning about whether to throw it away, this means not to reckon with the possibility that it might have won (cf. Hawthorne and Stanley, 2008, p. 572). If the ticket doesn't serve any other purposes, throwing it away would seem to be a rational choice. Conversely, if throwing it away does not seem rational, because it is only highly probable on my evidence that it lost, then I should not adopt this assumption as a premise in my practical reasoning.⁶

³ See Weatherson (2005) and Ross and Schroeder (2014), but also Stalnaker (1984).

⁴ And neither are they about *acceptance* in the sense of Bratman (1992). There are, however, points of contact with the phenomena Bratman is interested in. Bratman considers cases where one takes something one does not believe for granted in order to simplify or speed up one's decision-making. He also considers cases where one brackets a belief because of the risks involved. What Bratman sees as a homogenous class of similar cases would, on the picture developed here, be described in heterogenous ways. The latter class of cases require a higher degree of belief, while the former concern decision-making without belief which is content with potentially second-best choices.

⁵ Weatherson (2012) makes a similar claim. He argues that a possibility should be reckoned with in decision-making iff one does not know that it won't be actual.

⁶ A consequence of this line of thought is that when a given proposition is adopted as a premise in one's decision-making, it is assigned probability 1 for the decision problem at hand (cf. Greco, 2013). However, a premise which receives probability 1 in a low-stakes situation may be assigned a probability less than 1 in a high-stakes scenario. A particularly salient way of implementing this behavior would be to construe the

Fourth, and finally, what is the relation between the two hypotheses **Outright Belief** and **Degrees of Outright Belief**? If we take degrees of outright belief to mark distinctions within the class of outright beliefs, then it is natural to assume that *outright belief simpliciter* as described in the first hypothesis corresponds to having some positive degree of outright belief.⁷ If one's degree of belief is low, then one is willing to rely on it as long as the stakes are low. If one's degree of belief is high, then one is willing to rely on it even when the stakes are high, but, of course, also when the stakes are low. An outright belief, construed in terms of a positive degree of belief, would therefore validate the first hypothesis, **Outright Belief**: a proposition is believed outright iff it is a proposition one is willing to use as a premise in practical reasoning as long as the stakes are low.

Moreover, we can see that the first hypothesis, **Outright Belief**, follows from the second hypothesis, **Degrees of Outright Belief**, under the present assumption: if outright belief simpliciter is taken to correspond to a positive degree of outright belief and a high/low degree of outright belief corresponds to a willingness to rely on the belief in high/low stakes practical reasoning, then an outright belief corresponds to a willingness to rely on the belief in practical reasoning as long as the stakes are low.

In this paper, I develop these two hypotheses further. Special attention is given to degrees of outright belief as they have not been discussed much in the debate about knowledge and action.⁸ But if degrees of outright belief help in navigating through decision problems involving different stakes, they play a vital role in our decision-making and might also shed some light on various problems in epistemology caused by stake-effects.

Before getting started, let me contextualize the paper's ambitions with respect to the wider debate about practical reasoning. It may at first be surprising that I am going to say much about degrees of outright belief but little on how they should inform one's decision-making. This is because I think there are two separate, albeit related questions here (cf. Jenkins Ichikawa 2012). There is, first, the normative question regarding the *entry conditions* of practical reasoning. What epistemic status should a proposition have for being relied upon in practical reasoning? This is the question the present paper addresses. But there is also a second question. Once one has said what kind of premises should be used in practical reasoning, one faces the question of how one should decide in the light of these premises. This question asks for a *decision rule* which would provide one with a recommended choice as output when the premises of one's practical reasoning are supplied as input. I am not going to answer this second

Footnote 6 continued

probabilities used in decision-making as stemming from a prior probability function by being conditionalized on the stake-sensitive set of one's premises of practical reasoning (see Schulz, 2017 for a proposal along such lines).

⁷ For the sake of simplicity, we shall assume that degrees of belief are measured in terms of the natural numbers. A positive degree of outright belief would therefore mean a degree greater than or equal to 1 (cf. Spohn, 2012, ch. 5). See also the "Appendix".

⁸ Degrees of outright belief have, of course, been discussed in other contexts. See, for example, Spohn (2012). Particularly relevant to the topic of this paper is Carter et al. (2016).

question. But there are a number of decision theories which could naturally be adapted for these purposes (cf. Lin, 2013; Schulz, 2017).⁹

During my discussion, I will make use of the influential assumption that outright belief *aims at* knowledge. Based on this, I will ask: If belief aims at knowledge, what do degrees of outright belief aim at? Do they have any kind of aim, and if so, what are they aiming at?

In the upcoming three sections, I will make these questions more precise (Sect. 2) and discuss two rival proposals for answering them (Sect. 3 and 4), ultimately opting for the second option. Section 5 concludes. An “Appendix” contains an initial discussion of the logic of degrees of outright belief.

2 The project: a norm for degrees of belief

There is a growing trend in epistemology to give an account of belief based on the concept of knowledge (Bird, 2007; Engel, 2004; Smithies, 2012; Williamson, 2000, 2005b). Speaking somewhat metaphorically, knowledge is taken to be the *aim of belief*. Less metaphorically speaking, it is assumed that there is an epistemic norm to the effect that one ought to believe that p only if one knows that p . In yet other words, knowledge can be seen as the *standard of success* for belief, so that an instance of belief counts as successful only if it constitutes knowledge. Here I will treat these variations as largely verbal and focus on the following formulation:

Knowledge Norm of Belief. One ought: believe p outright iff one knows p .

According to this thesis, there is an epistemic norm to the effect that this norm is violated iff one forms a belief which does not constitute knowledge.¹⁰ As knowledge implies belief—or so I will assume in this paper (*pace* Radford, 1966; Myers-Schulz and Schwitzgebel, 2013)—the right-to-left direction of the embedded biconditional is something one cannot violate. But, of course, one frequently forms beliefs which do not constitute knowledge and in doing so one violates the left-to-right direction.

The knowledge norm faces considerable resistance. For instance, Hawthorne et al. (2016) argue that belief is *weak*: belief is a much weaker attitude than the knowledge norm would have it. One can rationally believe that it will rain on the grounds that a weather forecast predicts a 90% chance of rain even when it is very clear that one does not know this.¹¹

In this paper, I cannot offer a proper defense of the knowledge norm. What I can do, though, is to sketch one way of reconciling the problematic data with a certain way of

⁹ There is also a decision theory by Giang and Shenoy (2000) specifically designed for ranking theory (for discussion, see Spohn, 2012, ch. 10.4, 2017, sec. 4), whose exact implications are, however, still somewhat under-explored.

¹⁰ Let me note, though, that the present paper requires only minimal assumptions about knowledge being the norm of belief. The central tenets of this paper are compatible with knowledge being merely a derivative goal of belief, parasitic on truth being its primary aim (see Wedgwood, 2002, who builds on Williams, 1978, pp. 37–45).

¹¹ McGlynn (2013) and Whiting (2013) object to the knowledge norm in a similar way. For an alternative explanation of such observations, see Stanley (2008), who argues that belief only appears to be weak because it is often used for *hedging* assertions (as in e.g. ‘Mary’s friend Ada is in town, I believe’).

construing the knowledge norm.¹² The suggestion I would like to make is that there is a notion of *outright* (or *full*) belief to which the knowledge norm applies.¹³ On this proposal, ‘outright’ (or ‘full’) is not redundant when used to qualify ‘belief’. Thus, having an outright belief would be incoherent with also believing that one does not know this proposition. But one could grant that a considerable number of our beliefs are ‘weak’ in the sense of Hawthorne et al. (2016) and do not answer to as strong a requirement as knowledge.

If belief is weak in the sense of Hawthorne et al. (2016), then knowledge implies more than just belief. For instance, if I take myself to know that my lottery ticket lost, I take myself to be in a stronger epistemic position (along some salient dimension of comparison) than when I believe this proposition on probabilistic grounds alone. Thus, here is a first role for outright belief: it can serve as the subjective epistemic attitude which is the strongest kind of belief always implied by knowledge. Without the notion of outright belief, it would be hard to describe how ‘knowledge feels from inside’.

Interestingly, the case against a knowledge norm for belief does not extend to the *knowledge norm for assertion* (on the latter, see Williamson, 2000, ch. 11). If belief is weak, then it is occasionally fine to believe that one’s (lottery) ticket lost even though one has merely probabilistic evidence. But it seems always problematic to flat-out assert ‘My ticket lost’ when one does not know this. Now, there is the question of what kind of epistemic state one expresses when one makes an unqualified assertion. It cannot be weak belief, for the standards of assertion are higher than those for weak belief. Outright belief, on the other hand, would serve this role well, as Adler (2002, p. 231) contends.

Lastly, there seems to be a clearly identifiable difference in how propositions are treated in practical deliberation depending on whether we just have a ‘weak belief’ or whether we fully believe the proposition in question (Hawthorne & Stanley, 2008, p. 582f.). As Hawthorne and Stanley (2008, p. 572) note, we are normally not inclined to throw a lottery ticket away when we do not yet know the outcome of the lottery.¹⁴ This changes when we acquire the relevant knowledge. Once known, we would *treat the relevant proposition as true* in our practical reasoning and be prepared to act as if this proposition were true. Thus, on the present construal, outright belief would come with (or even consist in) a defeasible disposition to treat its content as true in one’s decision-making (see also Ross & Schroeder, 2014).¹⁵

To briefly summarize, I think there are good reasons to adopt a notion of outright belief to which the knowledge norm applies. Outright belief is, on the present suggestion, the strongest subjective attitude implied by knowledge and comes with a defeasible disposition for unqualified assertion and for treating its content as true in one’s practical reasoning.

¹² I expand on this line of defense in Schulz (forthcoming).

¹³ This also seems to be Adler’s (2002) view, see especially Adler (2002, p. 36) and Adler (2002, p. 40f.).

¹⁴ As a reviewer pointed out, this may not always be true. Think, for example, of passing through a country in which there is a fine on owning a lottery ticket because gambling is prohibited.

¹⁵ I should note that ‘treating as true’ is intended to include propositions with a probabilistic content: an outright belief that it is *likely* going to rain can just as well be a starting point of practical reasoning.

With the present notion of outright belief in place, I would now like to focus on the picture of degrees of outright belief best combined with a knowledge norm for outright belief. If outright belief aims at knowledge, how should degrees of outright belief be construed? More precisely, the project of the paper can be described succinctly as finding a solution for the following equation:

The Project. One ought: believe p outright to degree n iff $\phi(p, n)$.

What content must ϕ have for this to be correct? Is there even any (systematic) relation of the desired kind? A positive answer to these questions would most likely describe an epistemic attitude towards a proposition p which comes in different levels of strength n . Just to give an example (to be discussed properly in the next section), such an attitude could be *higher-order knowledge*. Degrees of outright belief, on this construal, would aim at higher iterations of knowledge; one should believe p to degree n iff one possesses n iterations of knowledge.

In order to tackle this project, it is instructive to keep in mind that degrees of outright belief have been introduced by linking them to practical reasoning. A high degree of belief means that one is willing to rely on the belief in high-stakes decision-making. This raises the question of how the knowledge norm for belief relates to practical reasoning.

Details vary, but the missing link is usually provided by a norm which speaks to the appropriateness of using something as a premise in practical reasoning.¹⁶ Here I wish to focus on a principle which is, just like the principle for outright belief above (= **Outright Belief**), explicitly restricted to low-stakes practical reasoning:

Appropriate Premises. It is appropriate to rely on p as a premise in low-stakes practical reasoning iff one knows p .¹⁷

If outright belief (i) necessarily comes with a willingness to be relied upon in practical reasoning (= **Outright Belief**) and (ii) is norm-conforming iff it constitutes knowledge (= **Knowledge Norm of Belief**), then it is only a small step to think that one's use of a proposition as a premise in practical reasoning is appropriate iff one knows it. (Appropriateness can in this context be understood as norm-conformity: willingness to rely on p is appropriate iff doing so does not violate any epistemic norms.¹⁸)

On the view explored in this paper, degrees of outright belief are tied to practical reasoning. In seeking a norm for degrees of outright belief, it is therefore natural to focus on norms for practical reasoning. The principle **Appropriate Premises** develops

¹⁶ The principle below is a weakened version of a stronger claim made by Williamson (2005a) and discussed in the upcoming section. Similarly stronger principles, which are not restricted to low stakes, are defended by Fantl and McGrath (2002) and Hawthorne and Stanley (2008).

¹⁷ Similar to their non-normative counterparts in Sect. 1, principles of this kind are usually restricted to premises which are *relevant* to the decision problem at hand. I leave this restriction implicit.

¹⁸ It should probably be granted that there can be pragmatic, non-epistemic reasons to violate **Appropriate Premises**. As practical reasoning absorbs time and energy, lack of time and energy can supply practical reasons for relying on an unknown proposition to speed up one's decision-making. On the present proposal, this point can be assimilated to the fact that one occasionally has all-things-considered reason to break a promise even though one always has a pro tanto duty to keep one's promises. In a similar vein, one can have all-things-considered reason to rely on an unknown proposition (e.g. "It won't rain") even though there is an epistemic pro tanto norm to the effect that one should not do so.

such a norm for low-stakes practical reasoning. The next step would be to investigate norms for high-stakes practical reasoning. This is the topic of the upcoming two sections.

3 Aiming at higher-order knowledge?

If I know p but do not know whether I know p , this casts some doubt on p . The lack of higher-order knowledge is a source of uncertainty. By the same argument, the presence of higher-order knowledge makes for more certainty. If I know p and also know that I know p , I am in a stronger epistemic position regarding p . It is therefore a natural thought that the more iterations of knowledge one possesses for p , the more one should rely on p in practical reasoning.

Williamson (2005a) probes the idea that judgements of appropriateness concerning practical reasoning are driven by judgements about higher-order knowledge:

For suppose that the agent in C knows p and the agent in C^* knows p^* , but the agent in C is in no position to know that she knows p and the agent in C^* is in no position to know that she knows p^* . Since stakes are higher in C^* than in C , we as theorists may view the failure of second-order knowledge in C^* more sternly than its failure in C , and therefore regard p as appropriate in C but p^* as inappropriate in C^* . The agent herself may take the same view, if she is aware of the difference in stakes. Similar explanations are possible on the assumption that the agent in C has n iterations of knowledge of p and the agent in C^* has n iterations of knowledge of p^* , while the agent in C is in no position to know that she has n iterations of knowledge of p and the agent in C^* is in no position to know that she has n iterations of knowledge of p^* , for fixed n . How many iterations are relevant depends on how much is at stake. [Williamson (2005a, p. 231)]

Williamson considers two agents who both possess first-order knowledge but lack corresponding higher-order knowledge. One agent is in a low-stakes situation while the other agent is in a high-stakes situation. Given that more is at stake in the second context, we may regard the proposition in question, p , as inappropriate in the second context. Relying on p in a high-stakes situation while lacking higher-order knowledge could easily seem out of proportion.

Going with this thought, there might be a close connection between appropriateness, stakes, and high-order knowledge:

[One] could try to build variation in the required number of iterations of knowledge into appropriateness itself, [...]: in some cases q would be appropriate iff one knew q , in others iff one knew that one knew q , and so on, depending on the stakes. [Williamson (2005a, p. 232)]

On this conception, what propositions should be relied upon in practical reasoning depends not only on their epistemic status but also on the stakes. The more that

is at stake, the more iterations of knowledge must one possess for a premise to be appropriate (see also Schulz, 2017). We may capture this idea in the following way:

Appropriate Premises (Higher-Order). It is appropriate to rely on p in practical reasoning of stakes n iff one possesses n iterations of knowledge for p .

Thus, there would be a close correspondence between appropriateness and higher-order knowledge. The more that is at stake, the more higher-order knowledge is required for appropriately relying on a premise in practical reasoning.

What could this tell us about degrees of outright belief? Recall that higher degrees of belief go with more willingness to be relied upon in high-stakes decisions. If relying on a belief in higher-stakes reasoning is appropriate just in case one possesses higher-order knowledge, then a certain degree of belief seems to be in order iff one possesses a corresponding number of iterations of knowledge. Higher degrees of belief would stand to higher-order knowledge like outright belief simpliciter stands to (first-order) knowledge. This would allow us to derive a first hypothesis about what degrees of outright belief might be aiming at:

Higher-Order Norm. One ought: believe p outright to degree n iff one possesses exactly n iterations of knowledge for p .

On this picture, outright belief of degree 1 aims at first-order knowledge, while beliefs of higher degree aim at higher-order knowledge. Thus, degrees of belief qualify the strength of outright belief. Outright belief simpliciter can be equated with a degree of belief of 1 or higher.¹⁹ This makes good on the idea that degrees of belief introduce further distinctions into the category of outright belief simpliciter. It also fits well with the idea that outright belief simpliciter aims at knowledge independently of whether or not one possesses further iterations of knowledge. Outright belief simpliciter has met its aim when one knows p without knowing that one knows but also when one knows and knows that one knows.

Interestingly, with this constraint in place, one can offer a simple story about what degrees of outright belief are: they are outright beliefs simpliciter about higher-order knowledge. That is, having a belief of degree 1 (or higher) means to believe p . Having a belief of degree 2 (or higher) means to believe that one knows p . And having a belief of degree 3 (or higher) means to believe that one knows that one knows p . To see this, just note that the idea of belief in p aiming at knowledge of p contains no restrictions as to the content of p . If p is an ordinary proposition, it implies that belief in p , Bp , aims at knowledge of p , Kp . By substituting ' Kp ' for ' p ' in this claim, we see that BKp aims at KKp . Further applications of this then show that $BKKp$ aim at $KKKp$, and so on and so forth.

Nice and neat though this account might be, there are various reasons for concern.

A *first reason* has to do with the danger of over-intellectualization. As Mikkel (2011, p. 539) points out, beliefs about knowledge are only possible if the agent possesses the concept of knowledge and is generally capable of engaging in higher-order reasoning.

¹⁹ See also Carter et al. (2016) who take belief to be the weakest attitude normally sufficient for rational acceptance. A similar view is developed by Wedgwood (2012) who suggests that outright belief could be equated with a disposition to treat a proposition as certain for all normal practical purposes; see also Ross and Schroeder (2014).

This point quickly fuels an objection to the present account of degrees of belief. If degrees of beliefs consist in beliefs about higher-order knowledge, one can only have degrees of higher order if one possesses the concept of knowledge. But a stakes-dependent willingness to rely on different propositions to different degree does not seem to require that one possesses the concept of knowledge. One can imagine an agent without the concept of knowledge who still trusts some of her beliefs more than others. Such an agent would be causally sensitive to how well confirmed a given belief is when relying on it in her decision-making. At least *prima facie*, the fact that this agent does not possess the concept of knowledge does not make her stakes-adjusted practical reasoning non-rational. However, as knowledge is a very basic concept acquired already at early age, the described agent is a fairly remote possibility. So one should probably not let this point carry too much weight.

A *second reason* is mentioned by Williamson himself:

Moreover, it [= a principle like **Appropriate Premises (Higher-Order)**] unperceptibly mixes together considerations at different levels. In knowing that one knows *q*, one knows a truth about one's own epistemic state; in knowing *q* itself, one may know a truth simply about the external world. One might therefore prefer the attractively simple KPR+ [= a version of **Appropriate Premises (First-Order)**; see below]. [Williamson (2005a, p. 232)]

Higher-order knowledge is partly knowledge of one's own mind. This is true even if one does not count knowledge as a state of mind (as Williamson, 2000, ch. 1 does). To the extent that first-order knowledge implies belief, higher-order knowledge is partly knowledge that one believes the proposition in question (at least it must be strong enough to validate this inference). Why should knowledge about what one believes matter for rational decision-making? This has been dubbed the *wrong content problem* by Carter et al. (2016, sect. 2) (see also Gao, 2019, p. 101).

The higher-order approach would posit a peculiar *discontinuity* between low- and high-stakes decision-making. In a low-stakes situation, knowledge of *p* would be sufficient. No knowledge of one's mental state would be required. But in high-stakes situations, not only the epistemic standing of *p* would matter, but also whether one is in a certain mental state.

To follow up on this objection, let us go back to the case of Ada who is about to go to the airport (see p. 1). Can she rely in her reasoning on the proposition that her ID is in her purse? Given the importance of the matter, it is plausible that her knowledge must be somewhat stronger than usual. She should have checked shortly before leaving, or else it must be highly certain that her ID is in her purse. Nevertheless, it seems clear that the salient question for her is *whether her ID is in her purse*. It is this information which has to meet a somewhat higher epistemic standard. If Ada somehow lacked access to her own mind or failed for other reasons to address this higher-order question, but double-checked that the ID is in her purse, she is epistemically well positioned to rely on this information in deciding to leave for the airport.

To be fair, it is often very natural to engage in second-order reasoning when the stakes get high.²⁰ What is more, it could seem that Ada in the example above has little reason to double-check unless she asks herself “Do I really know this? Couldn’t I misremember?”. Given that I assume in this paper that Ada still believes and knows that her ID is in her purse, what reason could she have to double-check? To begin with, note that a first-order approach need not deny that higher-order questions can be relevant. On the contrary, one way of casting doubt on p is to question whether one’s belief in p was properly formed or might, say, be an instance of false memory. Thus, insofar as higher-order questions are relevant to the pertinent first-order question, they are expected to naturally arise even on the assumption that merely the first-order question ultimately matters.

Although second-order reasoning about one’s knowledge comes fairly easily, perhaps the strongest objection to the higher-order approach is that it is very hard to find cases where we naturally engage in third-order reasoning or higher. “Do I know that I know that I know that I know?” seems to be a question one is hardly ever inclined to ask oneself even in high-stakes situations. Thus, the plausibility of the higher-order approach regarding second-order reasoning does not translate to higher levels.

One could react to this observation by finding some way of stopping at the second level.²¹ Thus, one would follow the higher-order approach half-way by holding that higher stakes require that one should check on whether one knows. Although I do not see a principled objection to such a proposal, my impression is that it will not be easy to explain why the difference in first-level and second-level stakes requires a shift from the first-order question (“ p ?”) to the second-order question (“Do I know p ?”), while a shift from second-level stakes to third-level stakes does not require a similar shift from the second-order question to the third-order question (“Do I know that I know p ?”). It strikes me that one would gain an advantage over such a theory if one could stay at the first-level throughout and explain the plausibility of second-order questions in some other way (as I tried above).

What happens if one rejects the higher-order approach (and with it the hypotheses **Appropriate Premises (Higher-Order)** and **Higher-Order Norm**)?

Williamson suggests that one may simply stick to the simpler idea that, independently of how much is at stake, a premise is appropriate iff one knows it (to wit, see the last quote above). This would result in the following simple principle:

Appropriate Premises (First-Order). It is appropriate to rely on p as a premise in practical reasoning (no matter how high the stakes) iff one knows p .

Clearly, this principle avoids the aforementioned worries concerning the higher-order approach. It merely requires first-order knowledge in any kind of decision situation. But what about the intuitive data which motivated the higher-order approach in the first place?

Williamson offers a slightly modified story in return (here ‘Hi’ is an agent in a high-stakes context while ‘Lo’ is an agent in a low-stakes context):

²⁰ Occasional use of second-order reasoning might be a very general phenomenon, which can arise for any norm. “Do you know whether you are complying with this norm?” seems to be a question which naturally arises for any norm in certain critical contexts (see also Gao, 2019, p. 101f and Schulz, 2017, p. 480).

²¹ I thank an anonymous reviewer for this suggestion.

However, Hi has far more reason than Lo has to check on such practical reasoning, to engage in second-order practical reasoning about whether to trust the first-order practical reasoning. [Williamson (2005a, p. 232f.)]

Thus, both for Hi and for Lo a given instance of knowledge, p , may be equally appropriate. But because more hinges on the correctness of Hi's reasoning, she has more reason to reflect on it, according to Williamson. Should it turn out that she does not know that p is appropriate because she does not know whether she knows p , she has reason not to go through with the results of her first-order reasoning.

Opting in this way for a simple first-order principle about practical reasoning frees appropriateness from higher-order commitments. However, it does not free one's overall theory of rationality from higher-order commitments. One is still supposed to engage in higher-order reasoning when the stakes are high. Thus, in one sense (a *first sense*), if Hi knows p but does not know that she knows p , acting on p would be appropriate. But in another sense (a *second sense*), it would still be inappropriate.²² According to one's overall theory of rationality, Hi should have reflected on whether she knows that she knows. Thus, in a way one has saved one's primary notion of appropriateness from higher-order commitments by introducing a second notion of appropriateness which again comes with higher-order commitments.

Of course, for the envisaged secondary notion of appropriateness, similar worries as above apply (for discussion, see Brown, 2008; Gao, 2017, 2019; Gerken, 2011, 2015; Vollet ms). Moreover, as Jie Gao (2019, p. 101) points out, it is unclear whether the envisaged two senses of appropriateness are really there. It seems we are not inclined to say that in some sense, acting on one's knowledge is okay, but in another sense, it may not be if we lack the relevant higher-order knowledge.²³

Interestingly, rejecting the higher-order approach does not force one to accept the simple **Appropriate Premises (First-Order)**. One can reject the thought that higher stakes require higher-order reasoning without thinking that they do not require anything at all, i.e. that mere first-order knowledge is always enough. The alternative idea would be that higher stakes raise the epistemic standards without requiring higher-order reasoning (cf. Gerken, 2011). To this option I shall now turn.

4 A better way: stronger knowledge

Central to the proposal I would like to make is the idea that first-order knowledge can be more or less strong. Going back to our sample case, Ada has after rechecking

²² Note that the present distinction does not map perfectly onto the standard distinction of following a norm and being blameless in violating a norm, where the latter is often characterized as having a *good excuse*. In the present case, Hi actually follows the envisaged norm, and therefore does not violate it, blamelessly or not (Vollet ms makes this point). For a recent explication of the standard distinction, see Williamson (forthcoming).

²³ Could one perhaps fall back on a position according to which one merely needs to be *in a position to know that one knows*? Note first that this option is not open when the agent cannot even form the relevant higher-order thoughts. But even if the agent has the necessary abilities, the evidence which strengthens one's first-order beliefs (like double-checking) need not be evidence enabling higher-order knowledge. For instance, double-checking one's purse typically is not evidence for being in a certain mental state as required for knowledge. See Gao (2019, pp. 99–103) for discussion.

stronger knowledge than before that the ID is in her purse. This does not mean that she did not have knowledge before she checked again. She might well have known before because she *remembered* that her ID is in her purse. This knowledge merely gets strengthened when she checks again.

How does knowledge get stronger? It seems clear that knowledge imposes a certain *epistemic standard*. This standard can be satisfied to different degrees, while only a certain degree is necessary and, possibly in conjunction with other conditions, sufficient for knowledge. To illustrate, consider a familiar analogy. Height can be measured on a scale. One needs to be of a certain height in order to count as tall. Tallness imposes a certain standard on one's height. Yet two tall people can, obviously, still be of different height. Similarly, two instances of knowledge can satisfy the epistemic standard necessary for knowledge, while one instance satisfies this standard to a higher degree.

The present considerations are often considered problematic because they seem to suggest that expressions like 'know' must be *gradable* (Stanley, 2005, p. 35f.). Yet 'know' does not seem gradable in the same way as a predicate like 'tall' is gradable, although there are expressions like 'he knows better than I do' or 'she knew full well'. However, it is fairly clear that a linguistic hypothesis about the gradability of a verb like 'know' does not follow from invoking epistemic standards in one's epistemology, as Romy Jaster (2013) points out in response to Jason Stanley. On some way of understanding epistemic standards, anyone (including Stanley) can agree (i) that epistemic standards can be more or less demanding and (ii) that knowledge requires one to satisfy a certain—not necessarily contextually invariant—epistemic standard. For this to be possible, the term 'know' does not have to be gradable. Rather, it merely requires that knowledge is *qualifiable* in a very broad sense: of two agents who satisfy the epistemic standard required by knowledge, one can still sensibly ask whether one agent satisfies a higher standard than the other.

What kind of epistemic standard does knowledge impose? This is a hard question. Prominent candidates include but are not limited to: *justification, safety, sensitivity, reliability, skillfully exercised cognitive success*. For example, consider justification. If the epistemic standard imposed by knowledge is justification, then knowledge requires a certain amount of justification (of a certain kind). But when comparing two justified beliefs, one belief can still be better justified than the other. Stronger knowledge, on this approach, would be knowledge with more justification than necessary. Or consider safety. If the epistemic standard imposed by knowledge is safety, then knowledge imposes a certain amount of safety. But when comparing two safe beliefs, one belief can still be safer than the other. Stronger knowledge, on this approach, would be knowledge which is safer than necessary.

This is not the place to adjudicate between different theories of knowledge. For this reason, I will explicate the proposal in terms of a general notion of epistemic standards without fleshing out what kind of epistemic standard knowledge in fact requires. A minimal condition on epistemic standards is that they can be taken to come in degrees, whether it is degrees of justification, degrees of safety or degrees of some other kind.

Here, then, is the alternative norm I would like to propose to govern degrees of outright belief:

Strong Knowledge Norm. One ought: believe p outright to degree n iff one knows p by satisfying an n -high epistemic standard for p (but no higher one).

To fix ideas, let a 1-high epistemic standard be the minimal epistemic standard imposed by knowledge. Consequently, a 0-high epistemic standard is a standard insufficient for knowledge. An outright belief of degree 0 is then best interpreted as no outright belief at all. An outright belief of degree 1 would aim at knowledge which satisfies the minimal epistemic standard for knowledge. Higher degrees of outright belief aim at knowledge satisfying higher epistemic standards.²⁴

The bracketed qualification (“no higher one”) takes care of the fact that in forming a belief of degree n , one can both overestimate and underestimate. If one forms a belief of a high degree without possessing strong knowledge, one has overestimated one’s epistemic position. Conversely, if one forms a belief of a low degree while possessing strong knowledge, one has underestimated one’s epistemic position. In both cases, one’s degrees of belief are not in line with the strength of the knowledge one possesses. Although overestimating may strike one as worse than underestimating, there seems to be something wrong in both cases. This is why the present norm penalizes divergences in either direction.²⁵

If degrees of outright belief are constrained by a norm like the **Strong Knowledge Norm**, they do not reduce to higher-order beliefs about knowledge. Rather, the content of higher degrees of belief continues to be p on all levels. If p itself is unconcerned with knowledge, so will be higher degrees of belief concerning p . In this regard, the present proposal is similar to the one put forward by Gerken (2011). According to Gerken, various pragmatic factors including the stakes influence how much *warrant* an agent needs for a proposition to be relied upon in practical reasoning. But as Gerken’s notion of warrant is for various reasons not knowledge-entailing (e.g. it is non-factive), his account differs from the present one which requires knowledge (of various strength) across the board.

The **Strong Knowledge Norm** dovetails nicely with the idea that higher degrees of outright belief come with a willingness to be relied upon in high-stakes decisions. The higher one’s degree of outright belief, the stronger is the knowledge one takes oneself to possess. And the stronger one’s knowledge, the more reasonable it is to rely on it when much is at stake.²⁶

The latter thought can be made more explicit by adjusting the constraint (labeled “**Appropriate Premises (Higher-Order)**” above) which states under what conditions it is appropriate to rely on a premise in practical reasoning. Previously, the constraint required higher-order knowledge. If rejected, it seemed one might have to accept that no matter how high the stakes, first-order knowledge is always sufficient. Yet with the present proposal in place, there is an alternative worth considering:

²⁴ Carter et al. (2016) offer a picture of belief which is *structurally* similar. One major difference in *content* is that they take *truth* rather than knowledge to be the primary aim of belief.

²⁵ Not much hinges on this feature. If preferred (as Carter et al., 2016 do on their account), one could equally well introduce a notion of degree of belief which just requires that the epistemic standard satisfies a given threshold.

²⁶ See Schulz (2017) for a (formal) decision theory which could accommodate such a norm.

Appropriate Premises (Generalized). It is appropriate to rely on p in practical reasoning of stakes n iff one knows p by satisfying an n -high epistemic standard for p .

On this account, higher stakes require something more: they require a stronger epistemic position. Stronger knowledge is required in order to rely on a proposition in high-stakes practical reasoning. But stronger knowledge remains first-order throughout. It is always an epistemic attitude towards p , the basic content. This way, one can simultaneously do justice to the impression that (a) higher stakes demand a stronger epistemic position without holding that (b) higher stakes require one to investigate one's own mental states.

One may wonder whether **Appropriate Premises (Generalized)** does not make it superfluous to posit degrees of outright belief.²⁷ With this constraint in place, it is already clear that one should aim in one's practical reasoning at a correspondence between the level of stakes and the strength of one's epistemic position. The concern would be that degrees of outright belief are not needed to explain good practical reasoning. All one needs to assume is that outright beliefs do not all come, and should not all come, with the same kind of willingness to fuel practical reasoning. Depending on the strength of one's epistemic position, only some beliefs would, and should, come with a willingness to be relied upon in high stakes situations.

My response, in a nutshell, is that such a story still implicitly commits one to degrees of outright belief. Recall that such degrees are meant to introduce further distinctions within the class of outright beliefs. That is, any belief of positive degree is an outright belief, while a degree of zero simply signifies the absence of an outright belief. Thus, degrees of belief do not introduce any new kind of epistemic state, they just make finer distinctions within an already existing type of state, that is within the class of outright beliefs. Now, on the objection just sketched, outright beliefs are still distinguished in terms of the circumstances in which they are relied upon. Moreover, there is a normative constraint on this distinction: reliance in higher stakes requires a stronger epistemic position. On my view, this is all it takes for positing degrees of outright belief. They merely describe and systematize the kind of distinctions one would like to make within the class of outright beliefs. Take a belief and look at the highest kind of stakes in which the agent is still willing to rely on it and you know its degree.

One might expect that the differences between degrees of belief are to some extent introspectively accessible. So, one might expect that beliefs of higher degree are felt to be stronger. But on the present account, these psychological features of outright beliefs are not wholly definitive of their nature. The present characterization fits better with taking degrees of outright belief to be *properties* of outright beliefs. Beliefs of different degree differ in their potential to influence decision-making.²⁸ If one favors a dispositional account of belief (see e.g. Schwitzgebel, 2002), one may cast this in dispositional terms: degrees of outright belief would be properties of a complex disposition which describe the range in which the disposition is triggered, that is when one is disposed to activate the belief in one's reasoning. On the present construal, degrees of outright belief would not be an independent kind of entity. Rather, they would

²⁷ I owe this concern to an anonymous reviewer.

²⁸ A higher degree of belief may also come with more resistance to (belief) revision.

supervene on the modal profile of one's beliefs which pertains to one's reasoning. It is important to note that being supervenient does not mean to be useless. Mass might be a supervenient property of the modal profile of particles, but it is still extremely useful for doing physics. In a similar way, degrees of outright belief are supervenient properties of the modal profile of beliefs, but they are still extremely useful for one's theory of reasoning, or at least that is the contention of this paper.

Naturally, there are a variety of open questions. One concerns the possible *observability* of degrees of outright belief.²⁹ Credences in part earned their keep by being associated with an agent's *betting behavior*. Very roughly, it is taken to be indicative of a certain credence in p when an agent is inclined to accept bets on p which would be favorable in light of this credence. Although the exact relation between credences and betting behavior is a matter of dispute, it seems to provide a pathway to showing that a cardinal scale of probabilities might be required for explaining human action.

What kind of observable consequences can be expected from degrees of outright belief? Recall that the central point of degrees of outright belief is to account for the stake-sensitive variability of reliance. Depending on how much is at stake, one may rely on one belief but not on another when facing a certain decision. Thus, for any belief an agent has, there is a certain trajectory through a hierarchy of stakes: for any stakes of a given level, we can ask whether the agent would still be prepared to rely on this belief when facing a decision involving this kind of stakes. An outright belief of a certain degree then corresponds to the highest kind of stakes where the agent is still prepared to rely on the belief in question. This already indicates that degrees of belief have observable consequences on the basis of which one could make informed predictions about which degrees of outright belief an agent is likely to have.

The present line of thought can be made more precise by relating it closely to the betting interpretation of credences.³⁰ The latter usually focuses exclusively on the *odds* an agent is prepared to accept when betting on a proposition. But note that a bet not only comes with certain odds but also with certain *stakes*. A bet with a potential gain of 1\$ and a potential loss of 100\$ has the same odds as a bet which makes one 100\$ while risking to lose 10.000\$. It is clear that the stakes of these two bets differ, so that someone may be prepared to accept the first bet without being prepared to accept the second bet.

Now, if outright belief is taken to be a kind of practical certainty (as has been proposed in this paper), then one would expect that an agent is prepared to accept any odds on an outright belief as long as the stakes are low. This is because certainty in a proposition makes all odds look favorable. The higher the stakes get, however, less and less beliefs are treated as practically certain because less and less beliefs are of

²⁹ In this regard, it should be mentioned that there is a *representation theorem* for degrees of outright belief in the context of ranking theory by Hild and Spohn (2008) (see also Spohn, 2012, ch. 8). As Hild and Spohn show, if one adopts certain axioms for *iterated contraction*, then the agent can be represented as changing her beliefs according to ranking theory. In addition, there is a representation theorem by Giang and Shenoy (2000) for ranking theory which closely resembles the representation theorem by von Neumann and Morgenstern. For discussion, see Spohn (2012, ch. 10.4) and Spohn (2017, sec. 4). A proper discussion of representation theorems for degrees of outright belief as construed in this paper is, unfortunately, beyond its scope.

³⁰ Of course, the following interpretation requires the same bells and whistles as the original betting interpretation of credences: one has to factor in all sorts of psychological factors, including risk aversion.

a sufficiently high degree. Thus, the proposal would be that there is a *level-adjusted betting interpretation* of degrees of outright beliefs: the agent has a high/low degree of belief if she is prepared to accept any odds on the proposition believed in high/low stakes.³¹

5 Conclusion

Beliefs serve us in making decisions. We rely on them in practical reasoning. But we rely on some beliefs more than others. In the category of belief, useful distinctions can still be made. Some beliefs are such that we rely on them only when the stakes are low. Others are such that we rely on them even when the stakes are high. The first are beliefs of low degree, the second are beliefs of high degree.

When is it appropriate to rely on a belief of a certain degree? According to a prominent thought, it is appropriate to rely on a belief iff this belief constitutes knowledge. However, this claim should be qualified. More than knowledge seems required to appropriately rely on a proposition when the stakes are high.

Perhaps what is additionally required is higher-order knowledge. In higher stakes, one should only rely on a premise when one knows that one knows. But on closer inspection, this suggestion seems doubtful. Do we have to reflect on our own state of mind when making high-stakes decisions?

This paper makes a different, albeit related proposal. High-stakes decisions do not require higher-order knowledge. But they require stronger knowledge, that is, knowledge which satisfies a higher epistemic standard. The resulting picture has it that degrees of outright belief aim at degrees of (first-order) knowledge just like unqualified outright belief aims at unqualified (first-order) knowledge.

Some readers might wonder whether the kind of strong knowledge posited in this paper is the kind of knowledge various contextualists and sensitive invariantists posit for unqualified knowledge ascriptions to be true in high-stakes contexts. As far as I can see, this is indeed the case. Knowledge of various levels of strength provides the anti-skeptical insensitive invariantist with a conceptual tool to mimic what others would describe as contextual variation or pragmatic encroachment of some kind or other.³² Developing this tool was the aim of this paper. Putting it to use will, unfortunately, have to wait for another occasion.

Acknowledgements I would like to thank Alexander Dinges, Roman Heil, Jakob Koscholke, Andy Müller, Patricia Rich, Sergui Spatan, Wolfgang Spohn and Jacques Vollet. Thanks are also due to two anonymous referees for *Synthese*. This work was supported by the Deutsche Forschungsgemeinschaft under Grant SCHU 3080/3-1/2/3 and was conducted within the Emmy-Noether research group ‘Knowledge and Decision’.

³¹ A crucial question for such a proposal is whether stakes are measured on an ordinal or on a cardinal scale. What is clear, however, is that no matter how one answers this question, one should give the same answer for degrees of outright belief. If stakes are discrete/continuous, so should be degrees of outright belief. On some interpretations of outright belief like the interpretation based on higher-order knowledge, measuring stakes on a continuous scale would therefore not be an option. Note that although I present degrees of belief as being discrete, the interpretation I offer does not foreclose construing them as continuous: epistemic standards are compatible both with a discrete and a continuous construal. Thanks to an anonymous referee for pressing me on this.

³² Gerken (2011, 2015) uses his account for a similar purpose.

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/>.

Appendix

A The logic of degrees of belief

It is commonly assumed that subjective probabilities of an ideal rational agent satisfy the axioms of probability theory. These axioms can therefore be seen as specifying the *logic of partial belief*. In the light of this, a natural question is: what is the *logic of outright belief*?

If degrees of outright belief aim at knowledge of a certain strength, one expects the logic of knowledge to inform the logic of outright belief. Following up on this assumption, this section starts with a discussion of *degrees of knowledge*, that is, knowledge of various levels of strength as introduced in Sect. 4 (recall that this requires a much weaker assumption than that 'know' is gradable). It will be shown that degrees of knowledge are, under idealized conditions, subject to the axioms of *ranking theory* (as presented e.g. in Spohn, 2012, ch. 5).

Let us start by making the notion of degrees of knowledge a little more precise. The strength of a given instance of knowledge can be defined in terms of the *maximal standard*:

Degrees of Knowledge. Let p be a proposition known by a subject at a given time. Then p is known to degree n iff n is the maximal number such that p satisfies an n -high epistemic standard. If p is not known, we stipulate that the degree of knowledge is 0.

Thus, degrees of knowledge keep track of the highest epistemic standard a given instance of knowledge satisfies. It immediately follows—just as in the case of outright belief—that a proposition is known without qualification iff it is known to some degree n greater than or equal to 1. We can then say that degrees of outright belief aim at degrees of knowledge: the standard of success for an outright belief of degree n is to possess knowledge of degree n . For the sake of simplicity, it is assumed that epistemic standards are measured in the realm of the natural numbers (including infinity).³³

³³ This assumption can be relaxed. Degrees of outright belief can be made continuous and defined over the real numbers (Spohn, 2012, ch. 5). If epistemic standards are provided in terms of safety, this would mean that safety conditions could be assumed to increase continuously as is, for instance, assumed by Pritchard (2014). However, as it stands, the framework developed in this section is incompatible with the idea suggested by Carter et al. (2016) that the different strengths of belief could be *incomparable*.

Are there any laws for degrees of knowledge so defined? As in the case of credences or subjective probabilities, general structural features can only be expected to be forthcoming under idealized conditions. For instance, knowledge is not automatically closed under logical consequence. But knowledge may well be closed under competent deduction in favorable circumstances (see Hawthorne, 2005, p. 29 and Williamson, 2000, p. 117).

For present purposes, we may assume that we are dealing with subjects whose logic skills are spotless and who have enough computational power to draw all available consequences of what they know. In particular, they have perfect knowledge of any logical theorem and do not believe any contradiction. This would suggest the following two stipulations:

- (i) Degrees of knowledge are measured within the realm of the natural numbers including 0 and ∞ .
- (ii) A contradiction always receives degree 0 while a logical theorem always receives degree ∞ .

The first property is a straightforward consequence of the simplifying assumption that epistemic standards are measured in terms of the natural numbers. The second feature stems from the fact that perfect thinkers never believe a contradiction while they have perfect knowledge of logical theorems.³⁴

Admittedly, these are no more than plausible stipulations having little to do with the nature of knowledge. One may therefore wonder whether there are any other valid principles more intimately linked to the nature of knowledge.

What has only been partially explored so far is the closure of knowledge under competent deduction. In favorable circumstances, competent deduction extends knowledge from premises to conclusion. This implies that the epistemic standard necessary for knowledge is such that if it is met for the premises, it is met for the conclusion as well. In the same way, it is plausible that stronger knowledge is also closed under competent deduction. That is, if knowledge of degree 1 is closed under competent deduction, one would expect knowledge of degree 2 (or higher) to be equally closed. For instance, if knowledge is closed under competent deduction partly because a safety condition on knowledge is so closed, it would be implausible to assume that this depends on the size of the safety zone.³⁵ In sum, if knowledge which is safe to degree 1 is closed, then knowledge which is safe to degree 2 (or higher) is equally closed.

Under idealized conditions, one may therefore assume the following closure constraint to hold:

Closure. For any propositions p_1, \dots, p_i which imply a proposition q and any n (relative to a subject and a time): if p_1, \dots, p_i all have a degree of knowledge of n or higher, then so does q .

³⁴ Of course, ordinary subjects do not always, if ever, have perfect knowledge of logical theorems. See Williamson (2009) for a way of integrating a safety condition without forcing logical omniscience.

³⁵ Depending on the exact details of one's notion of safety, one might even be able to show that this is so. See Schulz (2020) for relevant details. In general, there is some discussion on how best to construe a safety condition so that it can explain the closure of knowledge under competent deduction. See Sosa (1999, p. 143) and Pritchard (2005, p. 167f.). A pessimist stance is defended by Alspecter-Kelly (2011).

According to **Closure**, epistemic standards as required by knowledge are closed under logical inference.

With all this in place, we can identify two further interesting properties of degrees of knowledge:

- (iii) The degree of knowledge of two logically equivalent propositions is the same.
- (iv) The degree of knowledge of a conjunction equals the minimal degree of its conjuncts.

The first of these two properties is straightforwardly verified. By **Closure**, if p and q are logically equivalent, the degree of knowledge of one cannot be higher than the other.³⁶ From this it is only a small step to the second property, (iv). As a conjunction implies its conjuncts and vice versa, (a) each conjunct must, by **Closure**, have a degree of knowledge at least as high as the conjunction while at the same time, again by **Closure**, (b) the conjunction must have a degree at least as high as the minimal degree of the conjuncts.

The point of establishing these properties is to indicate that degrees of knowledge satisfy under idealized conditions the axioms of a *positive ranking function* (Spohn, 2012, p. 75). A positive ranking function is (i) any function β into the real numbers including 0 and ∞ which additionally satisfies (ii) $\beta(\perp) = 0$ and $\beta(\top) = \infty$ (with arbitrary contradiction \perp and arbitrary logical theorem \top), (iii) does not distinguish logically equivalent propositions and (iv) $\beta(p \wedge q) = \min\{\beta(p), \beta(q)\}$. These are exactly the conditions (i)–(iv) above. Succinctly put: degrees of knowledge under perfect conditions are (positive) ranks.

The present discussion provides a possible structure for degrees of knowledge which set the standard of success for degrees of outright belief. Does it follow that rational degrees of outright belief satisfy—again, under ideal conditions—the same structural axioms? Although there is a strong pull towards a positive answer, there is no direct link.³⁷

The missing link could be provided by principles which relate the aim of belief to when a belief is rational. For example, on a very strong conception, a belief would only be rational if it meets its aim. Now, if knowledge is the norm of belief, an outright belief would be rational if it constitutes knowledge.³⁸ Similarly, a degree of outright belief would be rational if it constitutes knowledge of the same degree. It would then follow that the logic of rational degrees of outright belief is the logic of degrees of knowledge, simply because rational degrees of outright belief *are* degrees of knowledge.

The missing link could also be provided by a weaker principle. One may think that a belief is rational if it would have met its aim *under normal conditions*. Rational

³⁶ If propositions are construed as sets of possible worlds, then assumption (iii) is a triviality, for on such a construal, logical equivalence amounts to identity. The assumption only has bite and requires justification

Footnote 36 continued
when propositions are more finely individuated (as many think they must be in order to serve as objects of belief and knowledge).

³⁷ Of course, one may think, as an anonymous referee pointed out, that similar axioms for outright belief enjoy independent plausibility. For instance, one could argue directly for a closure condition on belief.

³⁸ A proponent of a knowledge norm for belief like Williamson hesitates, though, to draw such a strong conclusion about rationality (although he accepts that justified belief requires knowledge); see Williamson (forthcoming).

belief could then be construed as would-be-knowledge under normal conditions.³⁹ On such a conception, the logic of rational degrees of belief could still be derived from the logic of degrees of knowledge.

But not all ways of relating the aim of belief to the question of rationality would provide a similar link. To see this, consider a view which takes a belief to be rational if it is *reasonably likely* that it has met its aim. If knowledge is the aim of belief, it would have to be reasonably likely that a belief constitutes knowledge. On such a conception, rational belief may well turn out not to be closed under competent deduction: it can happen that the premises of an argument constitute likely knowledge, while the conclusion constitutes unlikely knowledge. In such a case, it would be rational to adopt the premises, but it would not be rational to believe the conclusion. For this reason, the logic of degrees of outright belief would differ from the logic of degrees of knowledge.

To sum up, degrees of outright belief aim at degrees of knowledge. Under ideal conditions, degrees of knowledge satisfy the axioms of ranking theory. However, one cannot straightforwardly infer the logic of rational degrees of outright belief from the logic of degrees of knowledge. To bridge the gap, one has to make assumptions about how norms for belief constrain rational belief.

References

- Adler, J. E. (2002). *Belief's own ethics*. MIT Press.
- Alspector-Kelly, M. (2011). Why safety doesn't save closure. *Synthese*, 183, 127–42.
- Bird, A. (2007). Justified judging. *Philosophy and Phenomenological Research*, LXXIV, 81–110.
- Bratman, M. E. (1992). Practical reasoning and acceptance in a context. *Mind*, 101, 1–15.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42, 167–89.
- Carter, A. J., Jarvis, B. W., & Rubin, K. (2016). Belief without credence. *Synthese*, 193, 2323–2351.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, Epistemology, 13, 57–89.
- Engel, P. (2004). Truth and the aim of belief. In D. Gillies (Ed.), *Laws and models in science* (pp. 77–97). King's College Publications.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *The Philosophical Review*, CXI, 67–94.
- Gao, J. (2017). Rational action without knowledge (and vice versa). *Synthese*, 194, 1901–17.
- Gao, J. (2019). Against the iterated knowledge account of high-stakes cases. *Episteme*, 16, 92–107.
- Gerken, M. (2011). Warrant and action. *Synthese*, 178, 539–547.
- Gerken, M. (2015). The role of knowledge ascriptions in epistemic assessment. *European Journal of Philosophy*, 23, 141–61.
- Giang, P. H., & Shenoy, P. P. (2000). Linear utility theory for Spohn's theory of epistemic beliefs. In C. Boutilier & M. Goldszmidt (Eds.), *Uncertainty in artificial intelligence* (Vol. 16, pp. 220–29). Morgan Kaufmann.
- Greco, D. (2013). Probability and prodigality. In T. S. Gendler (Ed.), *Oxford studies in epistemology* (Vol. 4, pp. 82–107). Oxford University Press.
- Hawthorne, J. (2005). The case for closure. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 26–43). Blackwell.
- Hawthorne, J., Rothschild, D., & Spectre, L. (2016). Belief is weak. *Philosophical Studies*, 173, 1393–1404.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105, 571–90.

³⁹ The present idea is inspired by work on justification. According to Bird (2007) and Jenkins Ichikawa (2014), a doxastically justified belief is a kind of “would-be-knowledge” in close or normal circumstances.

- Hild, M., & Spohn, W. (2008). The measurement of ranks and the laws of iterated contraction. *Artificial Intelligence*, 172, 1195.
- Jaster, R. (2013). Contextualism and gradability—A reply to Stanley. In T. Spitzley, M. Hoeltje & W. Spohn (Eds.), *GAP.8 proceedings*, pp. 318–323.
- Jenkins Ichikawa, J. (2012). Knowledge norms and acting well. *Thought*, 1, 49–55.
- Jenkins Ichikawa, J. (2014). Justification is potential knowledge. *Canadian Journal of Philosophy*, 44, 184–206.
- Lin, H. (2013). Foundations of everyday practical reasoning. *Journal of Philosophical Logic*, 42, 831–62.
- McGlynn, A. (2013). Believing things unknown. *Noûs*, 47, 385–407.
- Myers-Schulz, B., & Schwitzgebel, E. (2013). Knowing that P without believing that P. *Noûs*, 47, 371–84.
- Nagel, J. (2010). Epistemic anxiety and adaptive invariantism. *Philosophical Perspectives*, 24, 407–35.
- Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.
- Pritchard, D. (2014). The modal account of luck. *Metaphilosophy*, 45, 594–619.
- Radford, C. (1966). Knowledge-by examples. *Analysis*, 27, 1–11.
- Ross, J., & Schroeder, M. (2014). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*, LXXXVIII, 259–88.
- Schulz, M. (2017). Decisions and higher-order knowledge. *Noûs*, 51, 463–83.
- Schulz, M. (2020). Finding closure for safety. *Episteme*, 1–15, <https://doi.org/10.1017/epi.2019.48>.
- Schulz, M. (forthcoming). Strong knowledge, weak belief? Forthcoming in *Synthese*, topical collection on “Knowledge First Epistemology”, edited by Christoph Kelp and Mona Simion.
- Schwitzgebel, E. (2002). A phenomenal, dispositional account of belief. *Noûs*, 36, 249–75.
- Smithies, D. (2012). The normative role of knowledge. *Noûs*, 46, 265–88.
- Sosa, E. (1999). How to defeat opposition to Moore. *Noûs*, 33, 141–53.
- Spohn, W. (2012). *The laws of belief*. Oxford University Press.
- Spohn, W. (2017). Knightian uncertainty meets ranking theory. *Homo Oeconomicus*, 34, 293–311.
- Stalnaker, R. C. (1984). *Inquiry*. MIT Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford University Press.
- Stanley, J. (2008). Knowledge and certainty. *Philosophical Issues*, 18, 35–57.
- Vollet, J. (ms). Anticliminosity, excuses and the sufficiency of knowledge for action. *Ms in progress*.
- Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19, 417–43.
- Weatherson, B. (2012). Knowledge, Bets and Interests. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions* (pp. 75–105). Oxford University Press.
- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives*, 16, 267–97.
- Wedgwood, R. (2012). Outright belief. *Dialectica*, 66, 309–29.
- Whiting, D. (2013). Nothing but the truth: On the norms and aims of belief. In T. Chan (Ed.), *The aim of belief* (pp. 184–203). Oxford University Press.
- Williams, B. (1978). *Descartes: The project of pure enquiry*. Penguin Books.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
- Williamson, T. (2005a). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly*, 55, 213–35.
- Williamson, T. (2005b). Knowledge, context and the agent’s point of view. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy* (pp. 91–114). Oxford University Press.
- Williamson, T. (2009). Probability and danger. *The Amherst Lecture in Philosophy*, 4, pp. 1–35. Available online at <http://www.amherstlecture.org/williamson2009/>.
- Williamson, T. (forthcoming). Justification, excuses, and sceptical scenarios. In J. Dutant, & F. Dorsch (Eds.), *The new evil demon*. Oxford University Press.

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