

Knowledge Dethroned

It is generally agreed that good practical reasoning involves some kind of belief-like or doxastic state, perhaps in conjunction with some kind of desire-like or conative state. But what kind of belief-like state figures in good practical reasoning? That is, when practical reasoning functions well, on the basis of what kind or kinds of doxastic state do we reason? Various answers to this question have been proposed. Some say we should base our practical reasoning on our (outright) *beliefs*.¹ Some say we should base it on our *credences* or partial beliefs.² Some say we should base it on our *justified beliefs* or on our *justified credences*.³ Recently, a number of prominent philosophers have argued that we should base our practical reasoning on, and only on, our *knowledge*. Thus, Timothy Williamson (forthcoming) states that “when practical reasoning functions well, one acts on what one knows.” Similarly, Jason Stanley and John Hawthorne argue that, where a proposition, *p*, is relevant to the choiceworthiness of the available options, “it is appropriate to treat the proposition that *p* as a reason for action iff you know that *p*” (Hawthorne and Stanley 2008: 578.) We may call the idea that we should reason only on the basis of what we know the *knowledge-first* view of practical reasoning.

There are many ways of arguing against this knowledge-first view, but our focus will be on what we will call the *problem of partial belief*. It seems that credences, or partial beliefs, can figure in good practical reasoning even when these credences don’t constitute knowledge. Stephen Schiffer was the first to raise this objection, which he illustrates as follows:

[The] thesis that one should act only on what one knows isn’t plausible. The following sort of example is very common: you are completely justified in carrying an umbrella even though you don’t know that it will rain but merely believe to degree 0.4 that it will rain. Not only is this a *prima facie* counterexample to [the knowledge-first view], insofar as it appears to be an example in which one is justified in acting as one does on the basis not of knowledge but of a justified partial belief, but there is a familiar and widely-accepted Bayesian account of why one’s action is justified, even though one isn’t acting on knowledge: one is performing that action which has the greatest expected desirability. (Schiffer 2007: 189)

¹ This idea is at the heart of belief-desire-psychology. Davidson (1963) holds that a primary reason is a pair consisting of a belief and a pro attitude.

² See Savage (1954).

³ The former is endorsed in Audi (2004) and Gerken (2011), the latter in Jeffrey (1965).

Defenders of the knowledge-first view are well aware of this problem, and have offered a number of responses to it. In this paper, we will consider these possible responses in turn, and argue that none can succeed.

1. The Knowledge of Probabilities Response

Hawthorne and Stanley (2008) suggest that cases where we appear to be acting appropriately on the basis of our credence or partial belief in p are in fact cases where we are acting on the basis of our knowledge, namely, our knowledge *of the probability of p* . Stanley illustrates this idea as follows:

...there are lotteries in which it is rational for me to buy a ticket, even though I do not know that I will win; when pressed to defend my purchase, I will respond that there is a chance I will win. But this is just to say that there are certain types of actions that I perform on the basis of beliefs about chances. In order for these actions to be acceptable, such beliefs must still constitute knowledge. (Stanley 2005: 10).

Similarly, in Schiffer's case in which someone with a .4 credence that it will rain carries an umbrella, Hawthorne and Stanley would maintain that the person acts appropriately only if her reason for carrying the umbrella is her knowledge of the probability of rain. Two kinds of question naturally arise for this view: first, what kind of probability is in question, and, second, what kind of knowledge of this probability is required? Concerning the first question, they say that the relevant kind of probability is epistemic probability, where the latter is understood as probability relative to one's knowledge or evidence (Hawthorne and Stanley 2008: 582). And concerning the second question, they tell us that we don't need to know the exact numerical probability—to justify taking the umbrella, we don't need to know, for example, that the epistemic probability of rain is precisely .4. Rather, what we need to know is “that it is reasonably likely that it will rain, that is, that it is likely enough to warrant the disutility of bringing an umbrella on my walk.” Thus, in general, on their view, in a p -dependent choice situation in which it is uncertain whether p is true, one appropriately ϕ s if and only if one ϕ s on the basis of the knowledge that the epistemic probability of p is sufficiently high (or sufficiently low, as the case may be) to warrant ϕ -ing.

One worry about this view is that it may seem overly intellectualistic. Surely we make choices under uncertainty all the time, and yet we rarely so much as entertain propositions about epistemic probabilities, let alone form beliefs about them. Hence, if the only way we can act appropriately under uncertainty is to act on the basis of our knowledge of epistemic probabilities, then it would seem that we hardly ever act appropriately when we make choices under uncertainty. But there may be ways around this objection. Jonathan Weisberg (2013: 10), for example, has suggested that our beliefs about epistemic probabilities may be *implicit in our credences*. “In general,” he says, “when one has credence x in P , one is disposed to judge that P is x probable. This close cognitive connection between credence and occurrent probabilistic

belief makes it plausible that credences constitute dispositional beliefs about probabilities” (Weisberg 2013: 11). Thus, on Weisberg’s view, when Jane has a .4 credence that it will rain, she is thereby disposed to form the belief that the epistemic probability of rain is .4, and so this partial belief constitutes a dispositional belief about the epistemic probability of rain. A similar, though more radical, view is proposed by Sarah Moss, who argues that credences can *directly constitute* knowledge about probabilities, without the need for any disposition to form explicit beliefs about probabilities (Moss 2013: 12). Moss and Weisberg agree, however, that Jane’s .4 credence that it will rain can constitute knowledge of the probability of rain. Hence, they hold that Jane can act on the basis of her .4 credence that it will rain, without violating the knowledge-first principle that one should reason only on the basis of what one knows.

But there remains a problem. Suppose we grant that credences can constitute knowledge of probabilities. Still, not all credences will constitute such knowledge. Even Moss and Weisberg require that, to count as knowledge, our credences must satisfy certain conditions, such as safety (see Moss (2013: 18) referring to the safety condition knowledge operative in Williamson (2000: 128). However, there are cases in which it seems that we act appropriately on the basis of such credences, but where the safety conditions is not met. As an illustration, consider the following case:

Borderline Probability Case: Elvis cares only about money, and he cares about it in a linear manner. He has no umbrella, and he’s deciding whether to buy one. An umbrella would cost him \$20.00. However, if he doesn’t buy the umbrella and it rains, then he’ll need to replace his blue suede shoes, which will cost him \$50.01.

Elvis’s decision problem can be represented by the following table:

	Rain	No Rain
Buy umbrella	-\$20.00	-\$20.00
Don’t buy umbrella	-\$50.01	\$0.00

Here the indifference point occurs where the probability of rain is .39992. Thus, if the probability of rain exceeds this value, then the expected utility of buying the umbrella exceeds the expected utility of not buying, whereas, if the probability of rain is less than .39992, then the expected utility of not buying is greater. Now suppose that the epistemic probability of rain is .4, and that Elvis rationally has .4 credence that it will rain. In this case, it would seem perfectly appropriate for Elvis to buy the umbrella on the basis of his .4 credence in rain. However, he would not thereby be acting on the basis of his knowledge that the probability of rain is sufficiently high to justify the purchase. For, assuming Elvis is an ordinary human being, the credences he assigns won’t perfectly track epistemic probabilities. Hence, there will be a nearby world where he has a .4 credence in rain but where the epistemic probability of rain is only .3999 (perhaps because the sky is one nano-shade less gloomy and so the evidence for rain is ever-so-

slightly weaker). Thus, even if we maintain Elvis's .4 credence in rain constitutes a belief that the probability of rain is .4, or that the probability of rain is sufficiently high to justify buying the umbrella, we must concede that there is a nearby world in which Elvis has this belief incorrectly. And so we must concede that this belief is unsafe, and hence that it does not constitute knowledge.

2. The Knowledge of Underlying Evidence Response

At this point, the knowledge-firster might reply as follows:

Even in the case just described, Elvis needn't act on the basis of any doxastic state other than knowledge. Suppose, for example, that what underlies Elvis's credence of .4 in the proposition that it will rain is his knowledge that it has rained on two of the past five days. In this case, Elvis could choose to purchase the umbrella on the basis of this knowledge of recent weather, rather than on the basis of his .4 credence in rain. Hence, he could avoid basing his decision on any doxastic state other than knowledge. More generally, whenever we act on the basis of our credences, we could equally act on the basis of our knowledge that underlies, and provides the evidence for, these credences. (This response is proposed in Hawthorne and Stanley, 2008: 584.)

This maneuver rests on the view, defended by Timothy Williamson, that evidence = knowledge, and that the epistemic probability of a proposition is determined by what one knows (see chapters 9 and 10 of Williamson 2000). Elsewhere, one of us has argued against this Williamsonian view (see Ross (2012) and Ross (manuscript)). There, it is argued that two agents having different experiences could know the exact same set of propositions, and yet the epistemic probability of a given proposition could differ for them in virtue of their differing experiences. If this is right, then rational credences don't supervene on known propositions, and so the latter cannot take the place of the former. We will not rehearse these arguments here. Instead, we will point out a different problem for Stanley and Hawthorne's proposal. The problem is that, even if the epistemic probabilities that are relevant to a given decision are determined entirely by some body of known propositions, it doesn't follow that we can make the decision directly on the basis of these propositions. For the propositions on the basis of which we make a given decision need to connect, in the right way, with our desires.

Let r be the proposition that it has rained on two of the last four days. And suppose, for the sake of argument, that Elvis chooses to buy the umbrella on the basis of his knowledge of r together with some conative state. This conative state must be, or entail, a preference for an outcome in which [r and Elvis buys an umbrella] over [r and Elvis does not buy an umbrella]. Note, however, that if Elvis has any such preference, it is a derivative one. No one cares intrinsically about whether [they buy an umbrella and it has rained on two of the past five days]. Rather, we care intrinsically about other things, such as whether we get wet, whether we ruin our shoes, etc. In the case of Elvis, we have supposed, for simplicity, that all he cares about is

money. Hence, if he prefers [r and Elvis buys an umbrella] to [r and Elvis does not buy an umbrella], this must be because he expects to end up richer if the former is true than if the latter is true. And whether this is so will depend on the probability of rain conditional on r . If the probability of rain conditional on r is greater than .3997, then it will be rational for Elvis to prefer [r and Elvis buys the umbrella] to [r and Elvis does not buy an umbrella], since, in this case, he should expect to end up richer if the former obtains than if the latter does. But if the probability of rain conditional on r is less than .3997, then he should prefer the opposite.

Thus, insofar as Elvis bases his decision to buy the umbrella on his knowledge of r together with a preference for [r and Elvis buys the umbrella] over [r and Elvis does not buy an umbrella], the latter preference will be a derivative preference, and will be based on this intrinsic desire together with his doxastic states. It might, for example, be based on his belief that the chance of rain conditional on r is .4, or on his belief that the chance of rain conditional on r is sufficiently high to warrant buying an umbrella. However, we should not expect this belief to constitute knowledge. For, just as the probabilities we assign to propositions don't perfectly track their epistemic probabilities, so the probabilities we assign to propositions conditional on bodies of evidence don't perfectly track their epistemic probabilities conditional on bodies of evidence.⁴ Hence, we should expect our knowledge of epistemic probabilities to involve a margin of error. Consequently, if epistemic probability of rain conditional on r is .4, and if this conditional epistemic probability must be at least .3997 in order for it to be rational for Elvis to purchase the umbrella, then Elvis won't know that the epistemic probability is precisely .4, nor will he know that the epistemic probability is sufficiently high to warrant the purchase. Hence, his instrumental preference for [r and Elvis buys the umbrella] over [r and Elvis does not buy an umbrella] will be based on a belief about probabilities that fails to amount to knowledge. And so his decision will ultimately be based on a doxastic state other than knowledge.

One might object to this argument as follows:

Where E is the total evidence available to Elvis, if his credence in rain conditional on E is .4, then he will thereby count as believing that the probability of rain conditional on E is greater than .3997. Moreover, truths about epistemic probabilities conditional on total bodies of evidence are necessary truths, or at least they are modally robust, not varying in any nearby possible world. Consequently, Elvis's belief that the probability of rain conditional on E is greater than .3997 is safe, since there is no nearby world where Elvis has this belief and it is false. And so this belief constitutes knowledge.

⁴ Here we are assuming, as our opponents all assume, that there is a distinction between epistemic or evidential probabilities, on the one hand, and subjective probabilities or credences, on the other. The subjective probability of a proposition is the degree of belief an agent *has* in that proposition. The epistemic probability of a proposition, for a given agent at a given time, is the degree of belief the agent *ought* to have in that proposition at that time.

But we should reject this argument. While safety is plausibly a *necessary* condition for knowledge (as we have been assuming), it is not plausible that safety, understood in the simple sense assumed above, is *sufficient* for knowledge. For the view that safety, so understood, is sufficient, would make knowledge of necessary truths too easy. It is presumably a necessary truth that the atomic number of Zirconium is 40. Hence, if you believe that the atomic number of zirconium is 40, *that particular belief* could not easily have been false. Hence, on the simple understanding of safety, this belief is safe. But it would be absurd to infer from this that the belief must constitute knowledge. If you formed that belief on the basis of a mechanism that could easily have under- or over-estimated the atomic number of zirconium by as much as two units, then we clearly should not say that you know that the atomic number of zirconium is 40. At most, we should say that you know that the atomic number of zirconium is between 38 and 42. Similarly, it may be a necessary truth that the epistemic probability of rain conditional on E is .4. Hence, if Elvis's credence in rain conditional on E is .4, that particular conditional credence could not constitute an overestimate of the probability of rain conditional on E. However, if Elvis's conditional credences in such propositions do not perfectly track their conditional epistemic probabilities, in the sense that his conditional credences could easily have exceeded or fallen short of the corresponding conditional epistemic probabilities by as much as, say, .02 units, then we should not say that Elvis knows that the conditional probability of rain is precisely .4. At most, we should say that he knows that it is between .38 and .42. Hence, we should deny that Elvis knows that the conditional epistemic probability of rain is sufficiently high to warrant buying the umbrella.

3. The Appeal to Excuses

It seems, as we have seen, that in cases like the Borderline Probability Case, it can be perfectly appropriate to act on doxastic states other than knowledge. Hawthorne and Stanley grant that it appears this way, and they offer an error theory for this appearance, as follows. “[T]here will of course be cases where one excusably acts on beliefs about chances that fall short of knowledge. As always, our diagnosis here will be that one violates the fundamental norm of practical reasoning but in an excusable way” (Hawthorne and Stanley 2008: 582). This appeal to excuses, they maintain, does nothing to undermine the importance of knowledge, because excuses “too are sensitive to the facts about knowledge” (Hawthorne and Stanley 2008: 573). In particular, if a fundamental norm has the form “ ϕ only if C” and if one ϕ 's despite the fact that C does not obtain, then, on the view they propose, one will have an excuse just in case *one does not know that C does not obtain*. Hence, on their view, if we act on the basis of partial beliefs that fail to amount to knowledge, then we will have an excuse just in case we don't know that these partial beliefs fail to amount to knowledge.

Could this appeal to excuses solve the problem raised by the Borderline Probability Case? The answer is no, and for two reasons. First, this appeal applies only where the agent is unaware of her ignorance. However, in the Borderline Probability Case, Elvis seems to act appropriately even if he knows that he lacks the relevant knowledge of probabilities. Suppose his credence in

Rain is .4, and he has calculated that the indifference point is .39992. Given this, he may know that he is in a borderline case. That is, he may know that the epistemic probability of rain is close enough to the indifference point that he could easily be wrong about whether this probability is high enough to justify buying an umbrella. In this case, he will know that he lacks the kind of knowledge required on the knowledge-first view. And yet, if he rationally has a .4 credence in rain and buys the umbrella on the basis of this credence, his choice would seem to be perfectly appropriate.

The second problem with appealing to excuses in the Borderline Probability Case is that it conflicts with the principle that *ought implies can*. For, according to this response, Elvis acts *inappropriately* in failing to act only on what he knows, and yet he has an excuse for this norm violation. And so the proponent of this reply is committed to saying that Elvis *ought* to act only on what he knows. However, Elvis is *not in a position* to act only on the basis of what he knows. For his knowledge is consistent with the possibility that the epistemic probability of rain is above the indifference point (and hence that he should buy the umbrella), but it is likewise consistent with the possibility that the epistemic probability of rain is below the indifference point (and hence that he should not buy the umbrella). Hence, what he knows is not a sufficient basis for making either choice. And since he cannot rationally decide what to do in this case purely on the basis of what he knows, it is implausible to maintain that he ought to do so.

4. Revising Decision Theory

So far we have been assuming, with our opponents, that it is rational to choose an option only if this option has maximal expected value or desirability. Thus, we have been assuming the following:

Maximization: For any agent *S* and option *O*, it is rationally permissible for *S* to choose *O* if and only if there is no alternative option, *O'*, such that the expected value or desirability of *O'* exceeds that of *O*.

One might, however, hold that this account of practical rationality is overly demanding. Consider, once more, the Borderline Probability Case. Here we have been assuming that, if Elvis has a rational credence of .4 in the proposition that it will rain, then he is rationally required to buy the umbrella, since doing so has an expected value of -\$20, whereas doing otherwise has an expected value of -\$20.004. However, since Elvis is not in a position to *know* that buying the umbrella has higher expected value, one might hold that Elvis can't be rationally required to purchase the umbrella. More generally, one might adopt the following, weakened version of decision theory:

Knowledge-Constrained Maximization: For any agent *S* and option *O*, it is rationally permissible for *S* to choose *O* if and only if there is no alternative option, *O'*, such that *S* is in a position to know that the expected value or desirability of *O'* exceeds that of *O*.

This view implies that, in the Borderline Probability Case, it's rationally permissible for Elvis to buy the umbrella, and it is likewise permissible for him not to buy the umbrella since, for all he knows, either of these alternatives might maximize expected value. On this view, Elvis needn't employ his .4 credence in the proposition that it will rain in deciding whether to buy an umbrella. He could instead reason on the basis of the strongest proposition he knows concerning the epistemic probability of rain—say, the proposition that the epistemic probability of rain is between .35 and .45. So long as this knowledge didn't rule out a given option, he could permissibly choose this option on the basis of this knowledge. Hence, he could make his choice entirely on the basis of what he knows.

But there is a problem with this view. For suppose the strongest proposition Elvis knows about the epistemic probability of rain is that it lies between .35 and .45. In this case, his knowledge wouldn't rule out paying \$20 for the umbrella, but his knowledge likewise wouldn't rule out paying, say, \$22 for the umbrella. For, if the probability of rain were .45, then paying \$22 for the umbrella would maximize expected value. Hence, for all he knows, buying the umbrella for \$22 might maximize expected value. Moreover, having purchased the umbrella, his knowledge that the epistemic probability of rain is between .35 and .45 would not rule out selling the umbrella back to the original seller for \$18 since, if the probability of rain were .35, then selling the umbrella for \$18 would maximize expected value. Hence, Knowledge-Constrained Maximization would permit Elvis to purchase the umbrella for \$22, and it would permit him to sell it back again for \$18, resulting in a loss of \$4. Thus, Knowledge-Constrained Maximization permits incoherent sequences of choices.

Thus, the knowledge-firster faces a dilemma. In the Borderline Probability Case, she can maintain either that Elvis is rationally required to buy the umbrella, or she can maintain that either alternative is permissible. If she holds that he is rationally required to purchase the umbrella, then she will have a hard time maintaining that he must base his decision entirely on what he knows, since what he knows doesn't rule out declining the umbrella. If, on the other hand, she adopts a more permissive view that allows him to either buy the umbrella or decline it, then she will have a hard time avoiding the conclusion that he would have both these options even if the price of the umbrella were slightly higher or slightly lower. And in this case it will be difficult to avoid the conclusion that he can permissibly make incoherent sequences of choices.

One might try to solve this problem by supplementing Knowledge-Constrained Maximization with a further principle that requires agents to constrain their present choices by their past choices in such a way as to avoid incoherent sequences of choices. Such a view will face problems in cases involving long sequences of actions in which an agent can't remember all her past choices. But it also faces a deeper problem. If all Elvis cares about is money, then, in deciding whether to sell his umbrella, his choice should be entirely forward-looking—he should be asking himself whether he should expect to end up richer or poorer as a result of selling the umbrella (keeping in mind that, if he sells the umbrella and it rains, he will need to replace his blue suede shoes at a cost of \$50.01). He shouldn't care how much he spent on the umbrella in

the first place. Even if he spent a large sum of money on the umbrella initially, if the chance of rain is sufficiently low and so it is sufficiently unlikely that the umbrella will be required to protect his shoes, then he should now be willing to sell the umbrella for a lower price. If Elvis were to take into account how much he paid for the umbrella initially in deciding whether to sell it now, he would be guilty of something similar to the sunk cost fallacy. And, when paired with Knowledge-Constrained Maximization, any view that requires Elvis to constrain his present choices by his past choices in such a way as to avoid incoherent sequences of actions would require him to do precisely this. And so any such view is implausible.

5 Revising the Knowledge-First View: The Indispensibility of Knowledge

The lover of knowledge might respond as follows:

I concede that the knowledge-first view, as you have defined it, cannot be accepted. That is, I concede that knowledge is not the *only* doxastic state that can figure in good practical reasoning, since credences that do not constitute knowledge can likewise figure in such reasoning. Knowledge remains crucial, however, since mere credences can figure in good practical reasoning only when accompanied by knowledge states. Thus, in the Borderline Probability Case, Elvis can properly reason using his .4 credence that it will rain only if there is something he knows. He needn't know that the epistemic probability of rain is .4, but at the very least he must know that it *might* rain. When he chooses to buy the umbrella, both his .4 credence in rain and his knowledge that it might rain play an essential role: his knowledge that it might rain constitutes his reason for buying the umbrella, and his .4 credence in rain determines the weight that he assigns to this reason in his deliberation. More generally, in order for our positive credence that p to figure in good practical reasoning, it does so by determining the weight that we assign to the known fact that *might* p .

This kind of proposal has been suggested by Jonathan Weisberg (2013: 13). Unfortunately, while this proposal may solve the problem in the Borderline Probability Case, it is insufficiently general. For it fails to solve the parallel problem that arises in cases like the following:

Borderline Epistemic Possibility Case: Rita is shown two lines, such that the first line is x times as long as the second line. She is considering the following proposition:

q : The first line is longer than the second line.

Let us say that q is epistemically possible for Rita just in case Rita does not know that q is false. If the value of x is sufficiently small (e.g., $x = 0.5$), then the first line will be significantly shorter than the second line. Thus, Rita will know that the first line isn't longer, and so q will be epistemically impossible for Rita. By contrast, for sufficiently large values of x (e.g., $x = 1$), q will be epistemically possible for Rita. Let m be the *maximum* value of x such that, if $x = m$, then q is *epistemically impossible* for Rita (or, if

there is no such maximal value, then let m be the minimum value of x such that, if $x = m$, then q is epistemically possible for Rita). Suppose that, in fact, x exceeds m by some very small margin Δ —say, $m = .95$ and $x = .95001$. Thus, if x were a tiny bit smaller, and hence the first line were a tiny bit shorter, then Rita could rule out q , but as it is she cannot quite rule out q . Rather, let us suppose, she rationally has a tiny positive credence in q , say 0.00001.

Since Rita has positive credence in p , it seems that, if she were offered a bet on q at sufficiently favorable odds (say, the bet costs a penny and she gets a million years in paradise if she wins), it would be perfectly appropriate for her to accept the bet. However, in doing so, she would not be acting on her knowledge that q might be true, or that q is epistemically possible. For, assuming she is a human being with ordinary cognitive limitations, she won't be able to discriminate between the actual world, where $x = .95001$ and q is epistemically *possible*, and an extremely nearby world where $x = .94999$ and q is epistemically *impossible*. Hence, any belief she may have that q is epistemically possible will be unsafe, and so it won't constitute knowledge. Thus, while Rita has positive credence in q , she doesn't know that her evidence doesn't rule it out, and so she doesn't know that q might be true. Consequently, when Rita rationally chooses to accept the bet on q , she must do so purely on the basis of her positive credence in q , unaccompanied by any knowledge that q might be true. And so we should reject the view that credences must be accompanied by the knowledge of corresponding modal propositions in order to figure in good practical reasoning.

6. Revising the Knowledge-First View: The Distinctive Role of Knowledge

One last response available to our opponent is the following:

I concede that we sometimes reason purely on the basis of our credences without knowing the probability, or even the epistemic possibility, of the propositions in question. Even so, there remains an important connection between knowledge and action. For knowledge plays a role in practical reasoning that is played by no other doxastic state. When we reason on the basis of our credences, we are treating the propositions in which we have credence as merely *probable*. By contrast, when we reason on the basis of what we know, we treat the known propositions as *true*. In general, it is the propositions we know, and only the propositions we know, that are *reasons* in the strict sense, as these are the only propositions it is appropriate for us to treat as true in our reasoning.

This kind of move is suggested in Jackson (2012: 355) and in Weatherson (2012). And it is also one way of reading Hawthorne and Stanley's claim that "it is appropriate to treat the proposition that p as a reason for action iff you know that p " (Hawthorne & Stanley 2008: 578). We will argue, however, that this position is untenable. For we will argue that either it is sometimes *appropriate* to treat as true propositions that *are not* known, or else that it is sometimes *inappropriate* to treat as true propositions that *are* known. Either way, it is not the case that the

propositions we may treat as true in our reasoning coincide with the propositions we know to be true, and so knowledge does play the distinctive role in practical reasoning that has been proposed.

To show this, we will focus on a case in which an agent is faced with two simultaneous decision problems, and where one of these is a “high-stakes” decision problem while the other is a “low-stakes” decision problem in relation to the same proposition.⁵ That is, we will focus on a case where an agent must simultaneously make two decisions, and in one of these decision problems, but not the other, mistakenly treating a given proposition as true would have disastrous consequences. As an illustration, consider the following case:

Elvis is scrambling to get ready for a concert. He has already put on his jumpsuit, and he must now fetch whatever boots he will wear. He has multiple pairs of white boots both in his dressing room and in his closet, but he has only one pair of gold boots which he placed in the dressing room just moments ago. He is deciding whether to go to the dressing room or to the closet to fetch the boots he will wear. While he is deciding this, a genie offers him a bet on the proposition that his gold boots are in the dressing room. If he accepts the bet, and the boots are indeed in the dressing room, then he will win a penny. But if he accepts the bet and the gold boots are not in the dressing room, then he will be condemned to eternal torment.

In deciding whether to go to the dressing room or to the closet to fetch the boots he will wear, it seems he could reason as follows:

My jumpsuit would go slightly better with gold boots than with white boots. So I should wear the gold boots. And the gold boots are in the dressing room. So I'll go to the dressing room.

Thus, in deciding where to go to fetch the boots he will wear, it seems that Elvis could appropriately treat it as true that the gold boots are in the dressing room, since he remembers putting them there just moments ago. For this is a low-stakes choice, and so he needn't take into account the minute possibility that the boots may be elsewhere. By contrast, in deciding whether to accept the genie's bet, it seems he should reason as follows:

I seem to remember putting my gold boots in the dressing room just moments ago. However, there's a tiny chance that I'm misremembering, or that someone moved my boots in the meantime. Hence, there's a tiny change that my gold boots aren't in the dressing room. Consequently, if I accept the genie's bet, there's a tiny chance I'll be condemned to eternal torment. It's not worth taking this risk for the sake of a penny. So I won't take the bet.

⁵ Further problems for this view are raised by Gettier cases. See Brown (2008: 171).

Thus, in making this decision, Elvis should not treat as true the proposition that his gold boots are in the dressing room, since the cost of being wrong are too high. He should instead treat this proposition as merely extremely probable.

Let t be the time at which Elvis is simultaneously making these two decisions. The proponent of the view we are considering faces a dilemma. At time t , either Elvis knows the gold boots are in the dressing room, or he doesn't know this. If he knows this, then, in deciding whether to take the bet, it is *not* appropriate for him to treat as true some proposition that he *knows*. If, on the other hand, he doesn't know that the gold boots are in the bedroom, then, in deciding where to go to fetch his boots, it *is* appropriate for him to treat as true some proposition that he *does not* know. Either way, we must reject the view that it is appropriate to treat a proposition p as true in one's practical reasoning if and only if one knows that p .

7. Conclusion

Let us return to Williamson's slogan with which we began, namely "when practical reasoning functions well, one acts on what one knows." Or, as he puts it elsewhere in the same paper, "when the cognitive-practical system functions well, one acts on what one knows." We have argued against this view by considering cases where it has implausible implications. But even apart from these considerations, there is reason to be skeptical of this Williamsonian view. To see why this is so, consider an analogous claim: "when the heart functions well, it beats 72 times per minute." Clearly, this claim is absurd. The role of the heart is to pump blood in a manner that enhances our fitness. To fulfill this role, the heart must pump at different rates in different circumstances: when we are at rest, it must pump at a lower rate, and when we are active it must pump at a higher rate. In a given circumstance C , our heart functions well just in case it pumps blood in a manner which, in circumstance C , is fitness-enhancing.

The same applies, *mutatis mutandis*, to the cognitive-practical system. The role of this system is to carry out cognitive processes, including practical reasoning, in a fitness-enhancing manner. In order to fulfill this role, the cognitive-practical system must function in different ways under different circumstances. There are many relevant dimensions along which our circumstances can vary. One such dimension concerns the availability of *time*: when we have all the time in the world, we do best to reason in a careful and methodical manner, weighing every relevant consideration. But when we must act quickly, we do best to appeal to simplifying heuristics. Another such relevant dimension concerns the availability of *knowledge*. Sometimes we are in what we may call *epistemically favorable circumstances*, that is, circumstances in which we know enough that we can effectively reason purely on the basis of what we know. In such circumstances we may indeed do best if we confine ourselves to what we know in deciding how to act. However, as we have seen, we can sometimes find ourselves in *epistemically unfavorable circumstances*, such as the Borderline Probability Case. In such circumstances, our knowledge provides an insufficient basis for making optimal decisions. Hence, the most fitness-enhancing manner of practical reasoning will involve relying on doxastic states other than knowledge. Consequently, our cognitive-practical system must rely on such states in order to function well.

References

- Audi, R. "Reasons, Practical Reason, and Practical Reasoning", *Ratio* 17 (2004): 119-149.
- Davidson, D. "Actions, Reasons and Causes", *Journal of Philosophy* 60 (1963): 685-700.
- Gerken, M. "Warrant and Action", *Synthese* 178 (2011): 529-547.
- Hawthorne, J. and J. Stanley "Knowledge and Action", *Journal of Philosophy* 105 (2008): 571-590.
- Jackson, A. "Two Ways to Put Knowledge First", *Australasian Journal of Philosophy* 90 (2012): 353-369.
- Jeffrey, R. *The Logic of Decision*, New York: McGraw-Hill, 1965.
- Moss, S. "Epistemology formalized", *The Philosophical Review* 122 (2013): 1-43.
- Ross, J. "Rationality, Normativity, and Commitment", *Oxford Studies in Metaethics* 7 (2012): 138-
- Ross, J. "Saving the Appearances: Distinguishing Evidence from Knowledge", manuscript.
- Schiffer, S. "Interest-Relative Invariantism", *Philosophy and Phenomenological Research* 75 (2007): 188-195.
- Stanley, J. *Knowledge and Practical Interests*, Oxford: Oxford UP, 2005.
- Savage, L. *The Foundations of Statistics*, New York: John Wiley and Sons, 1954.
- Weatherson, B. "Knowledge, Bets and Interests" in Jessica Brown & Mikkel Gerken (eds.), *Knowledge Ascriptions*. Oxford University Press. pp. 75-103 (2012)
- Weisberg, J. "Knowledge in Action", *Philosopher's Imprint* 13 (2013): 1-23.
- Williamson, T. *Knowledge and Its Limits*, Oxford: Oxford UP, 2000.
- Williamson, Timothy (forthcoming). "Acting on Knowledge", in Carter et al. (eds), *Knowledge-First* (Oxford: Oxford University Press).