

KNOWLEDGE, PRACTICAL REASONING AND ACTION

Peter BAUMANN

ABSTRACT: Is knowledge necessary or sufficient or both necessary and sufficient for acceptable practical reasoning and rational action? Several authors (e.g., Williamson, Hawthorne, and Stanley) have recently argued that the answer to these questions is positive. In this paper I present several objections against this view (both in its basic form as well in more developed forms). I also offer a sketch of an alternative view: What matters for the acceptability of practical reasoning in at least many cases (and in all the cases discussed by the defenders of a strong link between knowledge and practical reasoning) is not so much knowledge but expected utility.

KEYWORDS: knowledge, practical reasoning, expected utility

Several authors have recently argued that there is a close connection between knowledge and practical reasoning. Williamson for instance says: "One knows q iff q is an appropriate premise for one's practical reasoning."¹ We can call this the "knowledge norm of practical reasoning"; for the sake of brevity we can refer to the claim as "the connection thesis." Stanley states that "one should act only on what one knows."² Hawthorne agrees that knowledge of a proposition is necessary for using it as a premise in acceptable practical reasoning³; he adds that it is both necessary and sufficient.⁴ These kinds of claims are usually introduced by their defenders as intuitively plausible principles, supported by considerations and discussions of examples and cases,⁵ like, for instance, cases involving lotteries

¹ Timothy Williamson, "Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge," *Philosophical Quarterly* 55 (2005): 231; see also Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000), 47.

² Jason Stanley, *Knowledge and Practical Interests* (Oxford: Clarendon Press, 2005), 9.

³ See John Hawthorne, *Knowledge and Lotteries* (Oxford: Clarendon Press, 2004), 29, 85, 174-175.

⁴ See Hawthorne, *Knowledge and Lotteries*, 30; see also Jonh Hawthorne, Jason Stanley, "Knowledge and Action," *The Journal of Philosophy* 105 (2008): 571-590 and Jeremy Fantl, Matthew McGrath, *Knowledge in an Uncertain World* (Oxford: Oxford University Press, 2009). The thesis can easily be extended to multi-premise reasoning; the premises would, according to the connection-thesis, count as acceptable for one's practical reasoning as long as they are all known by the subject.

⁵ See, e.g., Hawthorne, *Knowledge and Lotteries*, 85, passim, and Stanley, *Knowledge and Practical Interests*, 9-10, passim.

(especially in the case of Hawthorne). The recent debate on this topic as a whole (see below) has also been very much driven by such considerations of plausibility. More systematic theoretical motivations for the connection thesis have been rare and if they play a role at all they are rather working in the background: Hawthorne's and Stanley's subject-sensitive invariantism⁶ goes very well with the connection thesis; for Williamson⁷ the connection thesis is part of a *knowledge*-centered systematic epistemological theory. Apart from that, general dissatisfaction with expected utility accounts of practical reasoning (see below) might play some role in the background, too.⁸

There are several problems with such claims about knowledge and practical reasoning – whether knowledge is deemed necessary or sufficient or both necessary and sufficient for acceptable practical reasoning and action.⁹ I will develop my objections step by step; most of them concern the necessity claim. I should stress from the beginning that I will follow the current debate on this topic and take “practical reasoning” to refer to instrumental reasoning here (if not indicated otherwise).

1. *The Main Example.* The example most often used in support of the connection thesis has to do with lotteries.¹⁰ Hawthorne gives the example of

⁶ See Hawthorne, *Knowledge and Lotteries* and Stanley, *Knowledge and Practical Interests*.

⁷ See Williamson, *Knowledge and Its Limits*.

⁸ Some might want to argue that being in a position to know *p* is necessary and sufficient for using *p* as a premise in acceptable practical reasoning. Alternatively, one might want to claim that being justified in believing *p* (or holding a justified true belief in *p*?) is the relevant condition. One would have to see the specific arguments for such claims; these arguments will be sufficiently different from the ones presented for the connection thesis. Hence, we should leave such related claims aside here.

⁹ To be sure, acceptable practical reasoning has to meet further conditions (see John Hawthorne, Jason Stanley, “Knowledge and Action,” *The Journal of Philosophy* 105 (2008): 578). This might make the search for sufficient conditions more difficult than the search for necessary conditions. For the sake of simplicity and because nothing hinges on it here, I will disregard concerns with these additional conditions. – It is very plausible to say (see Hawthorne, Stanley, “Knowledge and Action,” 572) that if a person ought to have known a certain proposition and taken it into account in her practical reasoning but is in fact ignorant of its truth and thus does not take it into account, then her practical reasoning is defect; however, this does not show that it is the knowledge of the proposition she is missing here rather than, say, the justified belief in it; apart from that, there are, of course, also many cases where the subject is ignorant of the truth of a proposition without there being any “obligation” to know its truth. For the sake of simplicity, I will leave cases of justified ignorance aside here.

¹⁰ Hawthorne, Stanley, “Knowledge and Action,” 571-574 present several further cases in favour of their view. I won't go into them because they don't add anything new here.

someone who's been offered a penny for his lottery ticket. The person reasons as follows:

"The ticket is a loser.

So if I keep the ticket I will get nothing.

But if I sell the ticket I will get a penny.

So I'd better sell the ticket."¹¹

Hawthorne claims that it is "clear enough" that this is a piece of bad reasoning and that "ordinary folk" agree with that. They would also agree that it is bad reasoning because the first premise is not known.¹²

First of all, it is not obvious that one can under no circumstances know a "lottery proposition" like the first premise of the quoted piece of practical reasoning.¹³ The problem Hawthorne's book¹⁴ is dedicated to¹⁵ shows why this is not so clear: If I know that I will never be rich (which seems plausible) and if I also know that this entails that I won't win the lottery, then how could I not know that I won't win the lottery? I don't want to go into this difficult problem here but only point out that it is no trivial claim at all that one cannot know a lottery proposition. It is also not obvious – as we will see in more detail below (section 5) – which pieces of reasoning are good or bad and why.¹⁶

¹¹ Hawthorne, *Knowledge and Lotteries*, 29; see also 85.

¹² See Hawthorne, *Knowledge and Lotteries*, 29-30; see also Hawthorne, Stanley, "Knowledge and Action," 571-572.

¹³ See Stephen Hetherington, *Good Knowledge, Bad Knowledge. On Two Dogmas of Epistemology* (Oxford: Clarendon Press, 2001), 102-107; Peter Baumann, "Lotteries and Contexts," *Erkenntnis* 61 (2004): 415-428.

¹⁴ *Knowledge and Lotteries*.

¹⁵ See originally Gilbert Harman, *Thought* (Princeton: Princeton University Press, 1973), 161.

¹⁶ See also Rhys McKinnon, "Lotteries, Knowledge, and Practical Reasoning," *Logos & Episteme* 2 (2011): 225-231 who points out that Hawthorne and others need to make sure that the reasoning in a lottery case like the one above is bad not just because it violates principles of expected utility (see below). – What makes people say that the kind of reasoning discussed by Hawthorne, Stanley and others is bad? Here are two weaker but not unpopular arguments which I want to discuss briefly just to get them out of the way. First, someone might imagine that the person first bought a ticket in the expectation that it might win and then decided to sell it in the expectation that it will lose. In this case, the person – if she has not simply changed her mind – has incoherent expectations. Incoherence would then be the problem but not the lack of knowledge of a premise of the practical reasoning – the connection thesis does not seem relevant here. Apart from that, this kind of incoherence would not show that the practical reasoning leading to the selling of the ticket itself is bad; the incoherence would rather be

Let me begin with some rather straightforward objections against the connection thesis (sections 2-4) before discussing it in the light of an alternative view (sections 5-7).

2. *Gettierization*. Let us start with a point which is important enough to stress right at the beginning. If knowledge of a proposition is necessary for its use in acceptable practical reasoning, then gettierized subjects cannot engage in acceptable practical reasoning. But this does not seem plausible at all.¹⁷ Compare Bo and Ben. Bo sees herself confronted with a raging bull. She comes to know there is a bull, engages in some swift practical reasoning and decides to run. Ben, however, finds himself, unknowingly, in raging bull facade county (where the locals frighten strangers with their exquisite raging bull facades). As it happens, he really is confronted with a raging bull. Given the circumstances, he doesn't come to know that.¹⁸ Despite lack of knowledge, he engages in some swift practical reasoning and

between two different pieces of practical reasoning (the one leading to the buying of the ticket and the other leading to the selling of it).

Second, in discussion philosophers supporting the connection thesis often ask "So, if you decide to sell it – why did you buy it in the first place?" If this is not the incoherence charge again, then it can be taken either as the charge that there is no good reason to buy a lottery ticket in the first place ("and are you not admitting this yourself by deciding to sell it?"); or it can be taken as the charge that there is no good reason to sell it, once bought ("and are you not admitting this yourself by having bought it?"). On the first charge: There are many reasons to buy lottery tickets. If one can get emotionally involved while watching a movie knowing that it is all fictional, then why not get excited in a similar way about the slim prospects of winning a lottery? People have all kinds of reasons for buying lottery tickets and these reasons can be bad but they need not be. Furthermore, people's reasons and the quality of their reasons need not have anything to do with whether they know that their ticket will win or lose. On the second charge: Why not sell a lottery ticket? Some people do that even as their profession (whether or not they buy them in the first place) and there seems nothing wrong with it as such. Or is it specifically because one does not know that it will lose? This question leads us back to the main issue and requires more than just the raising of a rhetorical question ("So, if you decide to sell it – why did you buy it in the first place?").

¹⁷ See also E.J. Coffman, *Old Dog Does New Tricks*, Ms., 2007; Jessica Brown, "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning," *Noûs* 42 (2008): 167-189, sec.5; Jonathan L. Kvanvig, *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press, 2003), 22; Mikkel Gerken, "Warrant and Action," *Synthese* 178 (2011): 535-6.

¹⁸ See Alvin I. Goldman, "Discrimination and Perceptual Knowledge," in his *Liaisons. Philosophy Meets the Cognitive and Social Sciences* (Cambridge & London: MIT Press, 1992), 86 for this kind of example as well as Clayton Littlejohn, "Must We Act Only on What We Know?" *The Journal of Philosophy* 106 (2009): 464-5, 469, and Ram Neta, "Treating Something as a Reason for Action," *Noûs* 43 (2009): 687-688 who use it in a similar way than we do here; the argument here can be easily applied to all kinds of Gettier-cases.

decides to run. It seems implausible to assume that Bo's but not Ben's practical reasoning is fine because Bo but not Ben knows that there is a raging bull in the vicinity.

Gettierized belief in the narrow, original sense of the term¹⁹ is justified true belief which does not amount to knowledge. Given the great variety of Gettier-like examples and the controversial nature of the notion of justification it might be better to use a broader notion of gettierized belief. In that sense, a belief is gettierized just in case it is true but does not constitute knowledge and this by no epistemic mistake of the subject. Epistemically blameless true belief which does not amount to knowledge is gettierized belief (in the broad sense). This explanation suggests that there need not be anything wrong with the practical reasoning of such a gettierized subject: If the subject fails to know the relevant proposition by no epistemic mistake of their own, then why should we blame them for their practical reasoning which is based on the relevant proposition? Sure, one would want to point out to the subject that she did not know what she seemed to know or that what seemed to be true is not in fact true; however, there is no reason to take this as a criticism of the subject's reasoning (how could she have done better?). Knowledge thus does not seem to be necessary for acceptable practical reasoning. One might want to object and argue that in such cases the subject is really violating a rule, though blamelessly. Lack of justified blame and the presence of good excuses, one could point out, do not entail that the rule of reasoning has not been violated. Since similar points will come up in the next two sections, I will reply to this kind of objection at the end of section 4.²⁰

3. *Truth*. Is even the truth of a proposition necessary for the acceptability of using it in practical reasoning? Consider the following two cases. Jill is close to

¹⁹ See Edmund L. Gettier, "Is Justified True Belief Knowledge?" *Analysis* 23 (1963): 121-123.

²⁰ One might also hold that the Gettier-objection only works against the connection thesis if one assumes in addition that the concept of knowledge can be reductively defined in terms of individually necessary and jointly sufficient conditions; according to this idea, it does not work against someone like Williamson who holds that the concept of knowledge cannot be reductively defined (see Williamson, *Knowledge and Its Limits*). I disagree. While it is true that Gettier cases are a problem only for reductive definitions of knowledge, it is also true that they pose a problem for any defender of the connection thesis. Defenders of reductive as well as of non-reductive accounts of knowledge can easily agree that there are gettierized subjects (that is, subjects who meet certain conditions while lacking knowledge); thus, the question whether such subjects are entitled to use the believed proposition as a premise for practical reasoning even if they don't know it makes a lot of sense for both parties. Williamson, for instance, agrees, too, that there are Gettier cases; he thus needs to answer the Gettier-objection above like everyone else.

Peter Baumann

dying of thirst when she finds some water. She comes to know it is water and engages in some quick practical deliberation:

This is water.

If I drink this, I will survive.

If I don't drink this, I will die.

So, I better drink this.

She drinks the water and survives. Jack, however, has recently travelled to twin earth but doesn't know about the differences between earth and twin earth.²¹ He finds himself in a situation identical to Jill's as far as the subject's perspective is concerned. However, what he takes to be water is really something different, namely t-water. Fortunately, t-water is as good for survival as water. Jack engages in some quick practical deliberation:

This is water.

If I drink this, I will survive.

If I don't drink this, I will die.

So, I better drink this.

He drinks the t-water and survives. Should we really say that there is something wrong with Jack's practical reasoning but not with Jill's because Jill knows that there is water whereas Jack's belief that there is water is not even true? This does not seem plausible.²² True belief does not seem necessary for acceptable practical reasoning based on that belief; it follows that knowledge is also not necessary. Perhaps Hume was right after all and belief (and desire) is all we need for an account of good practical reasoning? Again, one might object that all this only shows that the subject is blameless and excused but not that no rule of practical reasoning has been violated. I will get back to this kind of point at the end of the next section.

4. *KK*. What if S knows that *p* but does not know that he knows that *p*? Suppose S is unsure and does not know whether he knows that *p* (see Radford 1966). Assuming that this is compatible with S's knowledge that *p*, S could have a

²¹ See Hilary Putnam, "The Meaning of 'Meaning,'" in his *Mind, Language and Reality. Philosophical Papers, vol. 2* (Cambridge: Cambridge University Press, 1975), 215-271; if one does not like the semantic externalism involved in the example one can easily modify it by replacing reference to t-water by reference to some liquid indistinguishable from water.

²² See also Trent Dougherty, *Knowledge and Context-Sensitive Norms: A Defense of Simple Moderate Invariantism*, Ms., 2007.

good reason not to simply act on the proposition that p . Sure, if “not acting” is not a serious alternative, that is, if not acting on any proposition relevant to the practical issue at hand were to risk bringing about very bad consequences, then S ought to act on some proposition, and if the proposition that p seems more likely to be true than any alternative proposition, then S ought to act on the proposition that p . However, this might not be the case: “Not acting” might be unproblematic or better than acting on a proposition that “might,” according to S ’s worries, turn out to be false; apart from that, there could also be competing propositions which seem to S more likely to be true. In such cases, the subject does not have a good reason to act on the proposition that p and even has a good reason not to act on the proposition that p . Paul might know that the answer to the 1 Million Pound question is “Teheran.” However, he is cautious and really not sure whether he knows that: Didn’t he make bad mistakes about geography before? Suppose Paul can ask a friend about this. In such circumstances he should not simply go ahead and give “Teheran” as his final answer. Knowledge of a proposition is thus not sufficient for its use in good practical reasoning (see fn.1 above). Would knowledge that one knows be? Perhaps – but that thesis is much stronger and much less interesting.²³

What about the reverse case where someone (by no mistake of their own) does not know that p but has very good reasons to think they know that p ? One can think of cases of gettierization here or of cases like the above one about water. Isn’t it acceptable then to act on a proposition which is not known (whether that proposition is true or false)? Even if the practical stakes of getting it right are very high – which would be relevant to the case of an unknown false proposition), the subject ought not go by anything but their best reasons; if these reasons suggest that p , then S ought to act on p . If that is correct, then knowledge is not necessary for good practical reasoning.

Hawthorne and Stanley discuss such cases²⁴ and propose that they are not counter-examples to the connection thesis (or what they call the “knowledge-rule”) but rather cases where the subject has a good excuse and thus deserves no blame for violating the rule. This kind of reply can also be used against the cases mentioned in the last two sections.

²³ Williamson, “Contextualism, Subject-Sensitive Invariantism,” sec.V briefly discusses the appropriateness of practical reasoning from a known premise when one does not know that one knows the relevant proposition; however, what he proposes here (appropriateness depending on and varying with subject’s stakes) does not solve the problem above for the straightforward connection thesis. I will therefore not further go into this here. See also a brief remark in Jonathan L. Kvanvig, “Against Pragmatic Encroachment,” *Logos & Episteme* 2 (2011): 80-1.

²⁴ Hawthorne, Stanley, “Knowledge and Action,” 573, 586; see also Neta, “Treating Something,” 688-9.

However, there is a general methodological problem with this kind of reply: Can't one always say that? Or at least way too often? Hawthorne and Stanley don't make it clear how one can distinguish between a counter-example to the knowledge-rule and the case of an excused violation of the rule. Hawthorne and Stanley use this kind of reply – the interpretation of what might at first sight look like a counter-example as a mere excusable violation of their rule – quite often.²⁵ Without further arguments to the contrary, one is entitled to at least let it cut both ways: While Hawthorne and Stanley might say that something is an excusable violation of their rule, the sceptic about the connection thesis seems to have at least as much reason to see it as a counter-example. I don't want to put too much weight on this point and rather leave it at that; however, this should already raise some doubts about the connection thesis.²⁶

But wouldn't we, after discovering that our belief in some premise was gettierized or that we were mistaken about what it is that we refer to or that we were wrong about our epistemic states (see this and the last two sections), criticize our reasoning as inadequate? And wouldn't that show that knowledge is required for appropriate practical reasoning after all? I don't think so. Even if one acknowledged this kind of "drawback" one wouldn't have to concede and should in fact not concede that there was something wrong with one's original reasoning: there wasn't anything wrong with that. The gettierized person might find herself lucky when noticing that she had been gettierized but she would not have any reason to see her practical reasoning itself as deficient. Sure, she might concede that the basis of her reasoning was not as good and solid as she thought; but even then she would be right to insist that there was nothing wrong with her reasoning as such, given that there was nothing she could have done to improve the basis of her reasoning. Similar things hold for subjects who are mistaken about referents or about their epistemic state (see above). Even if the premises on which her reasoning were based turned out to be false, she ought not to accept blame for her reasoning itself.

5. *An Alternative: Expected Utility.* Consider again the original piece of practical reasoning:

(Case 1)

The ticket is a loser.

²⁵ Hawthorne, Stanley, "Knowledge and Action," 585-586 also remark that our intuitions about what constitutes a counter-example and what not are unclear in cases of failure of luminosity where one is in a particular (mental) state without knowing that one is.

²⁶ For more on excuses see Gerken, "Warrant and Action," 537-544.

So if I keep the ticket I will get nothing.

But if I sell the ticket I will get a penny.

So I'd better sell the ticket.²⁷

Why is it bad? Is it because the first premise is not known? Or is it for some other reason?

Compare (Case 1) with another case. John finds a lottery ticket with the number 666 on the street. He knows a crazy collector of items with the number 666 on them who will offer him \$5,000 for the ticket. He reasons in the following way:

(Case 2)

The ticket is a loser.

So if I keep the ticket I will get nothing.

But if I sell the ticket I will get \$ 5,000.

So I'd better sell the ticket.

What if, in addition, John did not find the ticket but was paid a good sum by the former owner of the ticket to take it because he thought that keeping a ticket with that number will bring terrible bad luck? Hawthorne remarks in a footnote²⁸ – without further argument – that it doesn't make a difference whether the ticket was free or not; I find it hard to see how this could not make a difference.

Finally, Hawthorne himself briefly mentions the case of a 10,000 ticket lottery with a \$5,000 prize where one ticket costs a cent. It would be irrational "to decline [buying a ticket] on the basis of your 'knowledge' that the ticket will lose."²⁹ Here is the reasoning for this case:

(Case 3)

That ticket is a loser.

So if I buy the ticket I will get nothing.

But if I don't buy the ticket I will keep one 1 cent.

So I'd better not buy the ticket.

(Case 1) and (Case3) strike us as cases of bad practical reasoning while (Case 2) is not so bad if not quite good. Why is that? All the cases share the same logical form.

²⁷ See Hawthorne, *Knowledge and Lotteries*, 29.

²⁸ See Hawthorne, *Knowledge and Lotteries*, 29, fn. 74.

²⁹ Hawthorne, *Knowledge and Lotteries*, 85.

We get closer to an answer if we make the relevant implicit assumptions about odds and stakes explicit. In (Case 1) and (Case 2) we can assume that the subject owns a ticket with a minute chance of winning (say 1 out of 1 Million) while the chances are much better in case (3) (1 out of 10 000). Let us assume the prize in (Case 1) and (Case 2) is 10 000 Dollars.

It is tempting to use Expected Utility Theory to analyse such cases and interpret them as cases of expected utility reasoning.. The basic idea is that in cases like the above ones appropriate practical reasoning identifies from the feasible set of available acts the act with the highest expected utility (whether it is explicitly guided by this idea or not). Given a finite and exhaustive set of mutually exclusive set of circumstances $c_1 \dots c_n$,³⁰ given further a utility function U of the subject which maps the outcome of each act in a given circumstance to a (measurable) utility $u_1 \dots u_n$, and given, finally, a probability function P of the subject which assigns each circumstance or outcome a certain probability $p_1 \dots p_n$, the expected utility EU of an act A can be characterized as the sum $p_1 \times u_1 + \dots p_n \times u_n$. Or:

$$EU = \sum_i p_i \times u_i$$

Given EU , we cannot represent practical reasoning any more in the form used above, namely simply as deductive inferences from given premises to a conclusion. Otherwise we could not explain why the pieces of reasoning differ in quality, given that they share the same logical form. Rather, we have to take into account that the reasoner has different credences in the different propositions. This, together with the different utilities, explains why the reasoning is good in one case and bad in the two other cases above.

Let us apply this idea to our 3 cases. For the sake of simplicity we may assume that the scale of the subject's utilities for money can be mapped by a positive linear transformation onto the scale of monetary values (nothing hinges on this simplification). Then the expected utility of keeping the ticket in (Case 1) is 1/100 expected Dollars; the expected utility of selling the ticket is also 1/100 expected Dollars. Both acts have the same expected utility and there is thus no reason (*ceteris paribus*) to prefer one to the other, given the theory. For this reason, the inference to the conclusion of (Case 1) is not a good one.

In (Case 2) the expected utility of keeping the ticket is, 1/100 expected Dollars while the expected utility of selling the ticket is 5000 expected Dollars. So, selling looks like the rational thing to do.

³⁰ If the set is infinite, things are more complicated; fortunately, the assumption of finite sets is unproblematic here and nothing hinges on it.

In (Case 3), finally, the expected utility of not buying a ticket is 0 Cents while the expected utility of buying a ticket is 49 expected Cents. Hence, it would be foolish not to buy a ticket.

According to the connection thesis, all 3 pieces of practical reasoning are bad because some premise is unknown. According to Expected Utility-Theory, the first and last case (1, 3) are cases of bad reasoning while the middle case (2) is a case of good reasoning. In cases like these, Expected Utility-Theory seems to give the correct answers while the connection thesis doesn't.³¹

Here is another case Hawthorne brings up:

"I will be going to Blackpool next year.

So I won't die beforehand.

So I ought to wait until next year before buying life insurance."³²

Given that the life insurance offers a reasonable deal, this is an example of bad reasoning. But, again, the reason it is bad is that it goes against the expected utilities and not the lack of knowledge of the premises of the inference (I leave it to the reader to go through the numbers for particular examples). Similarly, there are perfectly acceptable pieces of reasoning in favour of buying life insurance, even if the subject does not know the relevant premises of her reasoning. While the connection thesis excludes certain propositions from acceptable practical reasoning, namely the unknown ones, the principle of expected utility is much more liberal and accepts them all; the constraints of the latter view rather concern the question how one ought to use a proposition (namely according to Expected Utility-Theory) rather than whether to use it.³³

6. *Expected Utility versus Knowledge?* But is it really true that practical reasoning guided by the idea of maximizing expected utility does not require

³¹ I say "seems" because I have not offered a formal argument that Expected Utility-Theory always gives the right results in such cases. However, plausibility is enough for my purposes here, given that the main aim of this paper is to argue against a view and not so much for an alternative.

³² Hawthorne, *Knowledge and Lotteries*, 175; see also Hawthorne, Stanley, "Knowledge and Action," 571.

³³ On expected-utility accounts as alternatives see also the brief remarks in Hawthorne, Stanley, "Knowledge and Action," 580-585, Richard Feldman, "Knowledge and Lotteries," *Philosophy and Phenomenological Research* 75 (2007): 222-225, Alan Goldman, "Knowledge, Explanation, and Lotteries," *Noûs* 42 (2008): 471-472, Igor Douven, "Knowledge and Practical Reasoning," *Dialectica* 62 (2008): 101-118, and Richard Fumerton, "Fencing Out Pragmatic Encroachment," *Philosophical Perspectives* 24 (2010): 247; for acceptable reasoning using unknown lottery propositions see also Littlejohn, "Must We Act," 471-472.

Peter Baumann

knowledge of its premises? Expected Utility-reasoning can also be characterized informally by the following schema (restricted here, for the sake of simplicity, to the case of 2 feasible acts):

- (1) I can do A_1 or A_2 .
- (2) Given the possible circumstances and outcomes, A_1 has higher expected utility than A_2 .
- (C) Hence, I should do A_1 .³⁴

Or, more briefly (and very roughly):

Doing A_1 is the best means for the attainment of my ends.

Hence, I should do A_1 .

Does acceptable practical reasoning along such lines require that one knows one's feasible set of acts, the set of possible circumstances of action and the different outcomes of given acts in given circumstances (instead of having some belief about this which falls short of knowledge)? Since there are many cases where we engage in acceptable practical reasoning and since we often do not know these things, it seems very plausible to assume that the acceptability of practical reasoning does not require knowledge of these things. In other words, acceptable practical reasoning is compatible with lack of knowledge of at least some of its premises. I won't go more into lack of knowledge of one's options, circumstances and outcomes here because the point just made seems very plausible; it is uncontroversial (or not even an issue) in the debate on the connection thesis.

But what about probabilities and utilities? Doesn't acceptable practical reasoning require knowledge of at least them? We have seen that knowledge of the premises in cases like (Case1), (Case 2) and (Case 3) are not required. This is already an interesting result concerning the standard form in which practical inferences are usually presented by defenders of the connection thesis. But can't we restrict and reformulate the connection thesis and claim that acceptable practical reasoning though not requiring knowledge of all of its premises still does require knowledge of some premises, namely of those stating probabilities and utilities?

Let us take the case of utilities first. It is quite plausible, I think, to assume that agents often have mistaken views about or are ignorant of some of their utilities. However, even in such cases acceptable practical reasoning seems possible. Suppose I wonder whether I should spend the evening out with friends or rather

³⁴ Sometimes it is said that the conclusion of a practical inference is an action rather than a proposition; we can leave this complication aside here.

alone at home.³⁵ I reason in the light of what I take to be best for me. Suppose that I assume on the basis of good evidence (about myself) but falsely that it would be better for me to stay home alone tonight. Still, my reasoning that I should stay home alone strikes me as good while any reasoning resulting (on the same basis) in the conclusion that I should go out would appear unmotivated and foolish even if this option would really be better for me. To know one's own utilities is certainly an important advantage but it seems forced to say that it is also a necessary condition for acceptable practical reasoning.

I don't want to pursue the issue of knowledge of one's utilities any further here but rather, finally, go into knowledge of probabilities. This is a topic defenders of the connection thesis have commented upon – while they haven't said much if anything about knowledge of utilities.

Stanley points out that even if one is dealing with probabilities (or expected utilities) in one's practical reasoning, one still needs knowledge, namely knowledge of those probabilities³⁶). Take an agent who's deliberating about the question whether they should buy a ticket in a 10 ticket lottery with a \$10,000 prize for just one cent. If this reasoning is acceptable, then the agent knows the relevant probabilities, so the idea.

But does she really have to get the probabilities right in order for her reasoning to count as a good practical reasoning? Suppose both Ann and Barbara have received exactly the same information about some lotteries, – Ann about a lottery she is considering and Barbara about a different lottery she is considering. Suppose further that they both have no reason to be suspicious about the information. The only difference is that in Ann's case the information is correct whereas in Barbara's case it is incorrect. Should we really be so "externalist" about practical reasoning as to say that Ann's but not Barbara's practical reasoning is good because only Ann but not Barbara knows the probabilities? This defence of the connection thesis against the expected utility objection comes at a high prize: One would have to accept a very controversial theory of practical reasoning, namely externalism (the idea that the quality of practical reasoning is, at least partly, determined by factors which might not be accessible to the subject; defenders of the connection thesis have not

³⁵ See, e.g., Friedrich Waismann, "Ethics and the Will," in Friedrich Waismann, Josef Schächter, Moritz Schlick, *Ethics and the Will. Essays*, eds. Brian McGuinness and Joachim Schulte (Dordrecht: Kluwer, 1994), 53-137, sec.17 for the issue and similar examples.

³⁶ See Stanley, *Knowledge and Practical Interests*, 10, Jason Stanley, "Replies to Gilbert Harman, Ram Neta, and Stephen Schiffer," *Philosophy and Phenomenological Research* 75 (2007): 203-206, and Hawthorne, Stanley, "Knowledge and Action," 581-585.

offered any argument to this conclusion).³⁷ Apart from that, the theory would demand a lot from deliberators, namely knowledge about (objective) probabilities. But are probabilities “out there” for us to know? The defence discussed here thus comes with a lot of substantial commitments on other topics: One has to accept not only a strong externalism about practical reasoning but also a particular, controversial view of probability. This in itself speaks against this defence manoeuvre.

Perhaps those who want to make this kind of move should then rather think of probabilities as subjective or as epistemic (this is Hawthorne’s and Stanley’s move³⁸). However, if one does that, it becomes unclear why the agent has to know their probabilities and why it is not sufficient just to rely on them (e.g., as her betting dispositions). Her behaviour would then just express his probabilities (together with her utilities). Furthermore, in this case the move towards probabilities seems besides the point. Knowledge of the probabilities wouldn’t be knowledge of external facts, about the world rather than the agent; however, knowledge of the world (and not of one’s own mind) would be what is needed in practical reasoning if knowledge of any interesting kind is needed at all.

Hawthorne and Stanley offers a reply to this objection.³⁹ Suppose my epistemic probability that the restaurant is on the left is .6 because 3 out of 5 persons I asked told me so. Then I can deliberate and act on the known proposition that 3 out of 5 persons told me it’s on the left, and I don’t need to act on the known epistemic probabilities. It is very doubtful whether this reply helps Hawthorne and Stanley. Why is my knowledge that 3 out of 5 persons told me the restaurant is on the left relevant to the question whether I should go right or left – if not because it implies something about or simply reduces to a belief about the relevant probabilities? So, we’re back with knowledge of probabilities.

To summarize: If the defenders of the connection thesis refer to objective probabilities, then they run into serious problems (see above); if they only talk about knowledge of subjective or epistemic probabilities, then the connection thesis loses most of its bite: The knowledge necessary or sufficient for good practical reasoning would be about subjective or epistemic states and not about facts in the external world – which is what the knowledge relevant to the connection thesis should be about if the thesis is supposed to be interesting.⁴⁰

³⁷ To be sure, it is “better” to know the probabilities than not to know but this does not entail that there is something wrong with the reasoning of the unknowing subject.

³⁸ In Hawthorne, Stanley, “Knowledge and Action,” 584-585.

³⁹ Hawthorne, Stanley, “Knowledge and Action,” 584-585.

⁴⁰ See Stephen Schiffer, “Interest-Relative Invariantism,” *Philosophy and Phenomenological Research* 75 (2007): 188-195, sec.1; Schiffer also points out that the subject need not have any concepts of probability.

The upshot of all this is that appropriate practical reasoning does not require knowledge of the premises of inferences like the ones in (Case 1), (Case 2) or (Case 3). It also turned out that the usual representation of practical reasoning in these kinds of cases is a bit elliptical and needs to be reformulated, namely, as we say, in terms of expected utilities. An alternative view of practical reasoning as based on the idea of expected utility accounts better for our judgments about good and bad practical reasoning. This view also allows for acceptable practical reasoning without knowledge of the premises. (We should always remember that “practical reasoning” refers to instrumental reasoning here).

Sure, no practical reasoner is ever completely wrong or ignorant about all the facts relevant to her reasoning. And perhaps some knowledge of some propositions is required for practical reasoning. For instance, a practical reasoner who does not know some basic facts about how actions intervene in the worlds cannot count as a good practical reasoner. Conceding this is, however, far from agreeing with the much stronger connection thesis or the claim that good practical reasoning requires knowledge by the subject of all the propositions used as premises in their reasoning.

I have focused here very much on expected utility and there are several objections one might raise against this view. There are other, non-maximizing, conceptions of practical reasoning. “Satisficing” views, for instance, hold that enough is enough and that one need not or should not maximize goods.⁴¹ However, similar points can be made, *mutatis mutandis*, on the basis of these kinds of views. They also don’t entail or support the connection thesis. I won’t go into any detail here.

There are also choice situations where the subject has no idea concerning the relevant probabilities. In such cases, Expected Utility-Theory is inapplicable. It is not clear⁴² whether there are any principles or rules of good practical reasoning for such cases and, if so, which ones (*Maximin?*). This kind of case, however, does not seem to help the defender of the connection thesis much, given that subjects know even less here than in the cases considered above. What about the other extreme, namely cases where the subject is certain about the outcomes of given acts in given circumstances? These cases also don’t give the connection thesis any advantage over Expected-Utility-theory; the only relevant difference here is that the probabilities for the circumstances go up to 1 in one case and down to 0 in all other cases. It is interesting to notice that the defenders of the connection thesis typically present as

⁴¹ See Herbert Simon, *Reason in Human Affairs* (Oxford: Blackwell, 1983), Michael A. Slote, *Beyond Optimizing: A Study of Rational Choice* (Cambridge: Harvard University Press, 1989), Barry Schwartz, *The Paradox of Choice: Why More Is Less* (New York: ECCO, 2004).

⁴² See the overview in Michael Resnik, *Choices. An Introduction to Decision Theory* (Minneapolis & London: University of Minnesota Press, 1987), ch. 2.

cases of bad reasoning cases where the subject has some idea about the probabilities. The argument above, if it stands, seems sufficient against the connection thesis.

There are other principles or rules of good practical reasoning which often converge with Expected Utility-principles, like, e.g., the principle of dominance.⁴³ An act A dominates another act B just in case under no circumstance is the outcome of A worse than the outcome of B and under some circumstance it is better. The principle of dominance says that one should choose the acts which dominate all other feasible acts in a given situation of choice. Obviously, this principle is compatible with the lack of knowledge of all premises of the practical inference. There is no need to go into the details here. I will rather continue by illustrating the advantage of Expected Utility-theory over the connection thesis a bit more.

7. *Additional Considerations.* I would like to end by adding some further considerations against the connection thesis. The main weight lies on what has been said above but what follows should also be taken into account. I will discuss extreme stakes (a), further explanatory advantages of expected utility accounts (b) and end with some brief remarks on implications for epistemological scepticism and moral cognitivism (c).

(a) Considerations of extreme lotteries (or, more generally: of situations where very much is at stake) add further reasons to doubt the connection thesis. It is worth going into it briefly. Again, this rather supports the idea that it is expected utility and not knowledge that matters for practical reasoning.⁴⁴ Consider a 100 billion ticket lottery with a \$100,000 prize; you got your ticket for free but have the chance of selling it for \$90,000. It would be foolish not to sell the ticket even if you don't know that it won't win. This suggests that knowledge of a proposition is not necessary for acceptable practical reasoning based on that proposition.⁴⁵ Or take something you know for certain, like "if it rains, then it rains." It would still be foolish to bet your life on it, at least in normal circumstances (if you think that one

⁴³ As is well known, plausible principles of practical reasoning can conflict with each other. Newcomb's paradox is one very well-known case where the principle of dominance conflicts (or seems to conflict) with a principle of maximizing expected utility (see, e.g., Robert Nozick, "Newcomb's Problem and Two Principles of Choice," in his *Socratic Puzzles* (Cambridge: Harvard University Press, 1997), 45-73.

⁴⁴ See Stewart Cohen, "Knowledge, Assertion, and Practical Reasoning," *Philosophical Issues* 14 (2004), 487; Hawthorne, *Knowledge and Lotteries*, 148; Brian Weatherson, "Can We Do without Pragmatic Encroachment?" *Philosophical Perspectives* 19 (2005): 438-440; Dougherty, *Knowledge and Context-Sensitive Norms*; Schiffer "Interest-Relative Invariantism," sec.1; Brown, "Subject-Sensitive Invariantism," sec.7; Janet Levin, "Assertion, Practical Reasoning, and Pragmatic Theories of Knowledge," *Philosophy and Phenomenological Research* 76 (2008): 377-380.

⁴⁵ See also Levin, "Assertion, Practical Reasoning," 377-380.

should assume probability 1 for logical tautologies and should bet everything in such cases of maximal certainty, then you should modify the example and consider a case where the probabilities are extremely high but still below 1). This suggests that knowledge of a proposition is also not sufficient for the legitimate use of it in practical reasoning.⁴⁶ Expected Utility accounts, in contrast, have no problem at all accounting for such cases.

Hawthorne shortly mentions the last point and makes the following remark, sketching a response: "... we should consider whether knowledge of any proposition can be destroyed by environments in which a suitable bet is offered. One option, of course, is to think that the sketched connection between knowledge and practical reasoning is only roughly correct."⁴⁷ Let us leave the latter option ("One option ...") aside here: One would have to spell out in detail in what ways this is only roughly correct and how it could be modified; since this has not been done yet, this idea is hard to discuss. Consider rather the first idea: If the stakes are high enough, then the person does not know the relevant proposition (e.g., that if it rains, it rains). How good is that reply? Well, it will be attractive to those who like Hawthorne or Stanley (but not Williamson) hold that knowledge depends on what is at stake for the subject. It won't cut much ice for those who don't accept that theory.⁴⁸ So, the dialectical weight of this rejoinder is limited.⁴⁹ Apart from that, it is hard anyway to imagine circumstances in which "If it rains, then it rains" would become unknown.⁵⁰

⁴⁶ See for a different argument here: Thomas M. Crisp, "Hawthorne on Knowledge and Practical Reasoning," *Analysis* 65 (2005): 138-140, Jeremy Fantl, Matthew McGrath, "On Pragmatic Encroachment in Epistemology," *Philosophy and Phenomenological Research* 75 (2007): 558-589, and the remarks in Keith DeRose, *The Case for Contextualism. Knowledge, Skepticism, and Context, vol.1* (Oxford: Clarendon Press, 2009), 252-254, 262-268; also see fn. 4 above.

⁴⁷ Hawthorne, *Knowledge and Lotteries*, 177, fn.37.

⁴⁸ See also Brown, "Subject-Sensitive Invariantism," sec.7.

⁴⁹ See also the discussion in Hawthorne, Stanley, "Knowledge and Action," 587-589 where the authors do not commit themselves to a particular strategy against the objection above.

⁵⁰ One further criticism of the idea that an agent might know that p but still not be entitled to act on what they know says that this would allow even a knowing agent to further check the evidence. However, to say something like "I know she'll be at the party but let me check!" sounds weird (see Hawthorne, *Knowledge and Lotteries*, 148-149). Does it really? (see Brown, "Subject-Sensitive Invariantism," sec.7, Jennifer Lackey, "Acting on Knowledge," *Philosophical Perspectives* 24 (2010): 361-382 but also Neta, "Treating Something," 697.) What about the surgeon who claims to know that he is supposed to take out the appendix but decides to check the file again, "just in case"? Even if it does sound weird to say something like "I know she'll be at the party but let me check!": It does not seem to be relevant here. In the case of "If it rains, then it rains" we might not even be able to think of further evidence one could check. And

(b) Compare two pieces of knowledge. Suppose you know some very complex proposition about elementary particles and you also know that you exist now. You would certainly bet much more on the latter than on the former. This in itself does not show that knowledge isn't necessary for practical reasoning but it suggests at least that something else, the subject's probabilities in connection with expected utility, is doing the explanatory work. An expected utilities view which takes probabilities into account can, in addition, explain something the connection thesis cannot explain: why one should bet more on the second than on the first proposition.

Compare a strong belief or conviction which doesn't amount to knowledge (let us assume "I won't win the lottery" is an example) with some piece of knowledge the subject is not nearly as certain of ("The Copenhagen view on quantum mechanics is the right one"), assuming here that knowledge does not entail a probability of 1.⁵¹ It would be foolish not to bet more on the former than on the latter. This, again, suggests that what matters for practical reasoning is the subject's probabilities or expected utility and not knowledge. An expected utilities account, again, has an explanatory advantage over the connection thesis here: It can explain why the subject should bet more on one proposition than on the other.

Finally, suppose I have to decide whether to buy or not to buy a lottery ticket. Suppose further that neither do I know that the ticket offered will lose nor, of course, do I know that it will win. On the basis of what am I going to make my decision? Apparently not on the basis of knowledge concerning winning or losing.⁵² My reasoning should be rather based on probabilities. – Apart from all this, Expected Utility can easily explain cases of Gettierization, false belief, lacking knowledge that one knows and false but justified belief that one knows (see sections 2-4 above).

even if all the evidence has been checked, it would still not be appropriate to bet one's life on the trivial conditional.

⁵¹ If one disagrees here, holding that knowledge that p entails a probability of 1 for p , then this example will not work. However, this kind of defense is costly: One has to make very controversial and not intuitively plausible assumptions about knowledge and probability in order to defend a thesis, the connection thesis, which is supposed to be intuitively plausible. It is in general not a good strategy to defend the (allegedly) plausible with reference to the controversial. See also Jeremy Fantl, Matthew McGrath, "Critical Study of John Hawthorne's *Knowledge and Lotteries* and Jason Stanley's *Knowledge and Practical Interests*," *Noûs* 43 (2009): 185.

⁵² Hawthorne, Stanley, "Knowledge and Action," 581, fn.10, mentions this point shortly without discussing it.

(c) Consider a certain kind of epistemological scepticism which does not so much raise doubts about the existence of the external world but rather denies that anyone ever knows anything in particular about the world. Knowledge, according to this kind of scepticism, requires that we meet a condition (e.g., to be able to rule out that we are dreaming at the moment) which we cannot meet, given our actual constitution. Call this “scepticism.” Now, either it is (metaphysically) possible or not possible that scepticism is true. Suppose the defender of the connection thesis accepts that there is a possible world in which scepticism is true. Perhaps we do know lots of things in the actual world but would fall short of the conditions for knowledge in some possible world. Would we then (in that possible world) never be entitled to practical reasoning concerning what to do in the world? This seems very implausible⁵³; however, the defender of the connection thesis would, it seems, have to say exactly that – if he allows for the possibility of scepticism being true. The only alternative is to deny the latter and argue that scepticism is necessarily false. This, however, looks like a very strong thesis in need of much argument, and the connection thesis itself does not provide such an argument (neither does the general account of knowledge Hawthorne or Stanley favour). Hence, if the defender of the connection thesis does not want to go with the first, rather implausible option, they will have to accept a very substantial and controversial thesis concerning scepticism which is very much in need of argumentative support. This does, of course, not show that the connection thesis is false but it reduces its attractiveness drastically. One would first have to decide whether scepticism is necessarily false before one can reach a view about the connection thesis.

Another problem arises with respect to morality. Moral reasoning is an important type of practical reasoning. I will keep my remarks short here, also because the defenders of the connection thesis have said (next to) nothing about this aspect. Consider the following plausible piece of moral reasoning:

- (1) That person is in need of my help
- (2) If someone is in need of my help, then (given certain background conditions), I ought to (better) help that person
- (3) Hence, I ought to (better) help that person.⁵⁴

⁵³ See Dougherty, *Knowledge and Context-Sensitive Norms*, who argues that a subject in a sceptical scenario would have a justified false belief but her practical reasoning would remain unaffected; see also a brief passage in Kvanvig, “Against Pragmatic Encroachment,” 81.

⁵⁴ The background conditions mentioned here are of the following sort: I can help easily, there are no strong reasons not to help that person, etc. We can disregard these complexities here.

Peter Baumann

If the connection thesis also covers moral reasoning and not just instrumental reasoning and if it requires knowledge not just of the factual premises but also of the normative ones, then the connection thesis implies that some form of moral cognitivism must be true: The reasoner in our example needs to know the normative premises (e.g., (2)), too. Normative premises are knowable and truth-apt. If only the moral expressivists had known about the relation between knowledge and practical reasoning! But can a thesis in epistemology really have substantial implications in meta-ethics like moral non-cognitivism? We have good reason to be sceptical here.

To conclude, one should not expect decisive arguments in the debate about the relation between knowledge, practical reasoning and action. The considerations offered here, however, make a strong case against the connection thesis. It remains to be seen whether the defenders of the connection thesis can come up with convincing replies.