

*Knowledge and Context-Sensitive Norms:  
A Defense of Simple Moderate Invariantism*

*Trent Dougherty*

A recent approach to knowledge ties it to pragmatic considerations such as assertion and practical reasoning. In this paper I will first present and discuss what I take to be common arguments offered on behalf of this project. They are explanatory arguments based on several familiar kinds of cases. I will then offer an alternative explanation of the cases which presents a more unified explanation of the cases and relies only upon familiar considerations. I will argue that knowledge-pragmatism should be rejected, because it assumes a form of what I call Norm Invariantism, the idea that norms related to assertion and practical reasoning do not vary across contexts. I argue that Norm Invariantism is false.

Timothy Williamson thinks the quest for a traditional analysis of knowledge in terms of necessary and sufficient conditions is a wild goose chase.<sup>1</sup> Rather than follow so many others on that fruitless path, he prefers the more radical approach of making knowledge the conceptual anchor for other epistemological concepts as well as expanding the role of knowledge in other areas. One important part of the project goes by the slogan “Knowledge is the Norm of Assertion.” This slogan cashes out as the rule: “Assert something, only if you know it.” If Williamson is right, we could at least make assignments of semantic value to knowledge ascriptions in some—perhaps many—cases, for if we took someone’s assertion to be appropriate from the standpoint of conversational norms, we would feel good about ascribing knowledge to them. Being able to assign truth values to knowledge ascriptions would be a handy thing to be able to do. Further normative constraints on knowledge relating it to certain other actions would give us even more direction in evaluating knowledge ascriptions. Understanding the basis for knowledge ascriptions would count as some understanding of the concept itself, even if it falls short of a traditional analysis.

John Hawthorne pursues this strategy in his recent: *Knowledge, Lotteries, and Paradox*.<sup>2</sup> For the purposes of this paper I shall refer to the project of gaining some understanding of knowledge by its normative relation to practical concerns as the Pragmatist Project.<sup>3</sup> I will examine three constraints on knowledge which are part of the

---

<sup>1</sup> He has offered reasons why we shouldn’t even *expect* to be able to have such an analysis in the first place. Timothy Williamson, *Knowledge and its Limits* (Oxford: OUP, 2000), 3, 31-33.

<sup>2</sup> OUP 2004. Jeremy Fantl and Matt McGrath also advance a similar form of pragmatism in “Evidence, pragmatics, and justification”, *Philosophical Review*, 111 (2002), 67–94. Their proposal is subtly different and though I think roughly the same sorts of considerations apply to it, I will not argue that here.

<sup>3</sup> I want to be clear that I am not attempting to use “pragmatism” as a bludgeon in this discussion. I only mean it in the sense suggested by Susan Haak in which pragmatism “insist[s] on the connection of knowledge with action...” (Dancy and Sosa, *A Companion to Epistemology* (Oxford: Blackwell, 1992), p. 351). Like Hawthorne before me, I simply duck the question of what kind of normativity is involved here.

Pragmatist Project, at least as represented by Hawthorne. The first is “knowledge is the norm of assertion,” spelled out by Hawthorne as “The practice of assertion is constituted by the rule/requirement that one assert something only if one knows it”.<sup>4</sup> The second says that knowledge is also the norm of practical reasoning. The third—closely related to the second, since the end of practical reasoning is action—says that knowledge is also the norm of action. In Chapter 4 Hawthorne gives some examples of bad practical reasoning including selling a lottery ticket for a penny and not accepting cheap life insurance (details below). He then says: “In these deliberative settings, then, it is intuitive to suppose that the practical reasoning is flawed and that this is *because* the premise—whether it is an ordinary proposition or a lottery proposition—is not known.”<sup>5</sup>

The “because” can only mean that he is endorsing the reasoning process “*p* is not known, therefore, it shouldn’t be used in practical reasoning.” Lack of knowledge of *p* explains what’s wrong with the reasoning. It does so because lack of knowledge entails lack of appropriateness in employing *p* in practical reasoning. This gives us a norm of practical reasoning paralleling the norm for assertion. However, he goes on to say a few sentences later: “Insofar as it is unacceptable—and not merely because the content of the belief is irrelevant to the issues at hand—to use a belief that *p* as a premise in practical reasoning on a certain occasion, the belief is not a piece of knowledge at that time.... when you are offered life insurance, it would be unacceptable for you to use your belief that you are going to Blackpool as grounds for refusal. So on that occasion you do not know that you are going to Blackpool.” Now here the thought is that reasoning with *p* is not acceptable “so” it’s not knowledge. This is the converse of knowledge as the norm of practical reasoning. Here the idea is that you know something *only if* you can use it in practical reasoning. Putting his two principles together, then, we get that all and only known items are items one can employ in practical reasoning.<sup>6</sup>

I will state these principles as follows.

(AP) S may assert *p* at *t* only if S knows that *p* at *t*.

(RP) S may use *p* in practical reasoning at *t* iff S knows that *p* at *t*.

(PP) S may act as if *p* at *t* iff S knows that *p* at *t*.

---

Hawthorne states: “if someone asserts *p*, it is proper to criticize that person if she does not know that *p*” Hawthorne *forthcoming*, Section 1.3. However, he doesn’t say from what type of normative perspective this criticism may come.

<sup>4</sup> Hawthorne 2004, 23.

<sup>5</sup> Hawthorne 2004, 176 (emphasis added).

<sup>6</sup> It strikes me as unacceptable that knowledge is sufficient for practical reasoning. Suppose I know that O’Leary is Catholic and I’m trying to decide whether to hire her. Surely it is not permissible for me to employ that knowledge in my practical reasoning. I have two suggestions for making this counterintuitiveness go away. First, it is not unlikely that there is equivocation in the normativity involved in my counter example. Hawthorne does not say what kind of normativity he has in mind, but it’s probably weaker than the strong normativity involved in the example. Second, the principle could be modified along the lines as knowledge is *epistemically* sufficient for practical reasoning. In other words, if an item of knowledge cannot be used in practical reasoning, it is not due to epistemic deficiency, but rather to some non-epistemic factor.



some kind. But for what precisely? In both cases, the pragmatist alleges, what explains the wrong-headedness of these pieces of reasoning is that S has flouted the Reasoning Constraint. In the first case, S does not know that she will lose the lottery, so she should not employ the thought which expresses that proposition as a premise in an instance of practical reasoning. To repeat Hawthorne, “In these deliberative settings, then, it is intuitive to suppose that the practical reasoning is flawed and that this is because the premise...is not known”<sup>8</sup> In the second case, S does not know her flight will not crash, so she should not employ (1) as a premise. I have no quarrel with this form of argument. I think that as far as it goes Hawthorne’s explanation is a good one. Thus I think he has provided some support for (RP).

I have an alternative explanation which I think is simpler. What is wrong with the instances of practical reasoning is that they suppress the probabilistic nature of the premises when they ought not to. In the two cases above, S employs the following two premises:

- (P1) If I keep the ticket I will get nothing.  
 (P2) My flight will not crash.

The fact is, we may suppose, S is not certain of either of these propositions. In the context, a lot could ride on the outcome of actions based on them. In effect, not reasoning in a way that pays explicit attention to their probability treats them as certain. When a proposition is held with a very high degree of confidence, there is not much reason to pay attention to its probability, that is, to its divergence from certainty. In fact, from an evolutionary standpoint, there is a slight cost associated with unnecessary precision.<sup>9</sup> However, when confidence in a proposition falls too far short of certainty, it is unwise to treat it as certain. How far is too far will depend on what’s at stake. Since any rational person will harbor some reasonable doubts that they will not die in the next year it is unwise to reason as if one were certain of the premises in the practical deduction, given what’s at stake. This embodies something like the following norm:

- (DC) For any proposition  $p$  and context  $C$ , treat  $p$  as certain only if  $p$  is near enough to certain<sup>10</sup> as is appropriate in  $C$ .<sup>11</sup>

How near is “near enough” will vary by context, depending on what’s at stake in  $C$ . Whether it will be rational for S to sell the lottery ticket or take the insurance will depend, then, on what has the highest expected utility for S. We won’t, of course, expect S to calculate the expected utility as a decision theorist would, but S should weigh his chances

<sup>8</sup> *Ibid.*, 176.

<sup>9</sup> Imagine Ug standing there assessing the likelihood that there’s a tiger in the bush given his evidence.

<sup>10</sup> The conception of certainty I have in mind here is subjective probability 1.

<sup>11</sup> I am only suggesting (DC) as a norm of instrumental rationality. There will of course be moral norms—assert  $p$  only of doing so won’t unnecessarily hurt someone’s feelings—and constitutive norms—make only sincere assertions. My explanatory alternative only requires that I find sufficient reason to explain the inappropriateness of the assertions Hawthorne notes as inappropriate.

against the outcomes the way folks frequently do. The only normative principles in the above cases on my account are rationality principles some form of which most philosophers already accept.<sup>12</sup> There's no need to bring knowledge into the picture at all. This simple view is called by Hawthorne "Simple Moderate Invariantism"<sup>13</sup> (hereafter SMI).

In the Lottery Case, it may not have been rational for S to have purchased the ticket (from a monetary standpoint), but given that she did, keeping it will now have the highest expected utility. Likewise, in the insurance case, it may well be that taking the policy maximizes utility. If there is even a one in a million chance that she will be injured, then it clearly maximizes utility to accept the insurance.

It might be objected that ordinary folks just don't use decision theory and can't be expected to. If the objection is that folks don't use probability theorems and perform careful mathematical calculations, then the objection is probably correct, but it is also irrelevant. All that is required for my suggestion to be viable is that folks reason in roughly decision theoretic terms, which, of course, they do all the time. Mother says, "Take your umbrella to school honey." "But Mom, there's only a 40% chance of rain." "Still, you've got your cashmere coat on and it would be a shame to ruin it." Father says, "Hey, make sure you tape off the chrome when you paint that car." "Dad, when have I ever spilled paint on the chrome?" "Still, this particular paint won't come off chrome no matter what." In many cases a practical deduction will work just fine, its dictates will not significantly differ from what explicitly decision theoretic reasoning would suggest. However, in tricky cases, the decision theoretic nature of the case must be brought out (whether at the common folk level or the rarified philosophical level). Cases involving games of chance and (much the same thing) insurance are particularly amenable to this kind of analysis and particularly unlikely to be represented correctly in a practical deduction.<sup>14</sup>

The common form of the pragmatist's argument in the Lottery Case and the Insurance Case is as follows.

- (1) If folks don't know  $p$ , then they shouldn't use  $p$  as premise (underwritten by (RP)).
  - (2) S doesn't know  $p$ , but uses it as a premise anyway.
- thus
- (3) S is in violation of a norm embodying (RP).

This argument locates the impropriety of S's reasoning in the Lottery and Insurance Cases in the violation of a norm based on (RP). My alternative explanation locates any impropriety in violation of the Decision Principle:

---

<sup>12</sup> Well, maybe moral principles depending on your view of lotteries.

<sup>13</sup> Hawthorne 2004, 148.

<sup>14</sup> I might add here that in my estimation the sort of folk-decision-theory I describe is much more common than anything like the rigidly deductive practical reasoning as Hawthorne presents it. Both, I think, represent basically how folks reason, but it is clear that both are the creations of philosophers.

(DC) For any proposition  $p$  and context  $C$ , treat  $p$  as certain only if  $p$  is near enough to certain for  $C$ .<sup>15</sup>

The parallel argument for my (DC) explanation runs as follows.

(4) If folks aren't sure enough (for a context) of a proposition, they should take its probability into account (from DC)

(5)  $S$  can't be sure enough of  $p$  (for his context), but he doesn't take its probability into account.

thus

(6)  $S$  is in violation of (DC).

My explanation is a good one in that if (DC) is a valid principle, it implies that  $S$ 's reasoning in the Lottery Case and the Insurance case is defective, which is the datum to be explained.

So the first reason I gave for preferring the SMI account is that it adverts to the familiar notion of expected utility rather than sully the apparently purely epistemic notion of knowledge *logical* ties to practical concerns. *Prima facie* there is no reason to think knowledge is *logically* related to practical reasoning.<sup>16</sup> The second reason I think my explanation is preferable is that (DC) has a great deal more initial plausibility than does (RP). (DC) is probably accepted by most philosophers already in some form or other. (RP) is an innovation, a contentious one at that. It is plausible that there is some natural connection between knowledge and practical reasoning. Probably, most of what  $won$  knows could be used in practical reasoning and perhaps the converse is true. This does not, however, automatically lend support to a constraint on the concept. Thus, if (DC) does the job—which it does—then there's no reason to invoke (RP). There's just no work for (RP) to do once we realize that the more plausible (DC) explains what's wrong with  $S$ 's reasoning. This will emerge as a general pattern in what follows, strengthening my argument. That is, the above considerations apply to each of the primary principles of the pragmatist project, so a third consideration in favor of SMI is that it provides a unified theory of what's wrong with the cases, it explains what's gone wrong in all three types of cases with a single principle rather than three similar-sounding but separate principles.

Furthermore, there are already prior reasons to be skeptical about (RC). Almost any case of justified false belief will provide a counterexample to (RC). In any number of possible situations where it is *stipulated* that we don't or can't have knowledge, we can

<sup>15</sup> You might ask: How do I know how to apply this principle, since it is vague? I think the vagueness of the principle is the vagueness of ordinary life, whereas the problems with applying normative constraints based on the concept of knowledge really open Pandora's box: Do I know  $p$ ? How do I know that I know  $p$ ? I'll take every-day vagueness over that sort of thing any day.

<sup>16</sup> Though of course they might have a natural enough connection in human affairs, though even there I have my doubts.

still engage in practical reasoning without any relevant form of criticism. Suppose you're in a demon world where you don't or can't have knowledge. Is it really *wrong* (in any sense relevant to the present context) to reason practically there in just the same way you would if you were in a non-demon world? It is *unfortunate*, but any use of the word "norm" in which otherwise-good practical reasoning is said to be in violation of some norm in a demon world has stretched the word beyond its elastic limit.

Suppose the pragmatist bites the bullet and says Yes, offering as evidence the fact that if it were pointed out to the demon-worlder that she didn't really have knowledge she wouldn't have reasoned as she did. Again we have a simpler explanation: that she wouldn't have reasoned as she did because she discovered her premises were *false*. All ordinary people want to avoid reasoning with false premises, few are concerned to reason only with *known* premises, nor, frankly, should they. My demon world counterpart is not in violation of any kind of norm of practical reasoning when he does so just as I do even though he lacks knowledge.<sup>17</sup>

### The Argument for (PP) and (AP)

Pragmatists have also argued that there are practical constraints on knowledge by offering cases of the following form. Someone asserts something they may plausibly be taken to know. Then the stakes are raised by someone bringing to mind the possibility of error. The one making the assertion takes the possibility seriously enough to take back the claim, take measures to verify the claim, or change their course of action. Pragmatists assume that when the individual seeks to verify what they claimed to know they lose their knowledge (or at least manifest that they have lost it). If (PP) or (AP) are valid principles, then by simple modus tollens, one could deduce that the above situations do not involve knowledge. Thus, their truth would explain the withdrawal of the assertion or the unwillingness to act. Good explanations entail the evidence. If  $(A \rightarrow K)$  is true then  $\sim K$  entails  $\sim A$ . The datum is the withdrawal of an assertion. The explanation is that they are acting in accordance with AC and moving from lack of knowledge to the impermissibility of assertion or if they do not that we may judge them negatively for not doing so. These arguments depend on the assumption that it is in fact correct that knowledge is lost, which I will dispute. Let's consider a couple of cases.

#### *The Bank Case*

Hannah is driving her husband home from work. He says, "We'd better stop by the bank and deposit this check." She says, "It's late, lets just do it tomorrow, the bank is open on Saturday." "Are you sure, banks do change their hours, and if we don't get this check deposited we could be in big trouble." "No, I'm not sure, so we'd better deposit it now."

The pragmatist suggests that even if Hannah originally knew the bank was open on Saturday, she no longer does, since the stakes have raised and she withdraws her assertion and refrains from acting as if it's open.

---

<sup>17</sup> For an expansion of this line of argument see Jon Kvanvig's *The Value of Knowledge and the Pursuit of Understanding* (CUP: 2003), Chapter 1.

*The Oven Case*

Jim and his wife are five minutes down the road to a week's vacation. She says, "Jim, did you shut off the oven." Jim thinks about it a moment and he seems to remember turning off the oven (suppose this is a genuine instance of successful memory and that he knows he shut the oven off). "Yes, I did." "Are you sure, the whole house could burn down." "You're right; we should go back and check."

Here again, the pragmatist's view is that since the stakes have been raised, Jim no longer knows he turned the oven off. The data are supposed to be that Hannah and Jim withdraw their assertion. The proposed explanation of this invokes (PP) or (AP), since their truth would predict the alleged datum, given their lack of knowledge. I propose to undercut this argument for pragmatism by denying that they lack knowledge at all.

In each case the stakes have indeed been raised—by bringing to mind the possible *consequences*. In light of this, the norms of assertion and of action have been raised. As the consequences of a false move go up, so do the strictures on action and assertion. In the bank case, if there were no financial consequences of not getting the check deposited, then it would be fine to act as if the bank were open tomorrow the context-sensitive norm of practical reasoning in this low-stakes context is appropriately low. They could just deposit the check on Monday. The potential consequences didn't raise the standards for *knowledge*, but rather the standards for *action*. In some cases the norm of action is: *act only on what you know*; but in others it will be as strong as *act only on that of which you are practically certain* or as weak as *act on a whim*. It all depends on what the expected utilities are, as with the previous set of examples. Decision theory already gives us a framework to evaluate what's going on in the cases. The people in the cases above are not expected to do a formal calculation of the expected utilities of various actions, but people regularly weigh consequences and confidences in deciding what to do. Informal decision theory is often a part of normal human reasoning and the bank case and the oven case are typical examples.<sup>18</sup>

Intuitively, they've done the right thing in withdrawing their assertions and not acting on their beliefs.<sup>19</sup> The pragmatist explanation is that they are acting in accordance with the rule not to assert or act on what you don't know. the SMI explanation is that they are rightly sensitive to the decision-theoretic nature of the case (acting in accordance with (CD)). The SMI explanation has the advantage that rightness of withdrawing the assertion, the rightness of not acting on the belief, and the wrongness in the lottery reason case above are all explained by the same principle: they are/are not sensitive to the expected utility in a way that fits the context. In contrast, the sensitive invariantist adverts to three separate principles. Now neither the friend of SMI nor the friend of SSI think people are *consciously* acting in accordance with any of the explicit principles considered in this paper, that is, they do not think to them selves "there's this principle

<sup>18</sup> Studies intended to show that people are not good probabilistic reasoners are irrelevant in the present context. All that I need is that people often consider both how likely they think how likely an outcome is to come about and how good or bad it would be for them if it did. I take it to be abundantly obvious that people often do this.

<sup>19</sup> We must assume belief is present or the Knowledge Account is explanatorily idle, since if belief goes knowledge goes with it.

such-and-such which requires that I do so-and-so in this context.” However, I suggest that it is *much* more plausible that people are naturally sensitive to issues pertaining to risk aversion than to subtle differences in shift in their epistemic status. Everyone who’s folded a hand of poker knows when things have gotten too rich for their blood, few have every been conscious of going from a state of knowledge to a state of merely justified true belief. Risk assessment is mundane where epistemology is arcane.

In each case just considered above what has changed are the norms for *assertion and action*.<sup>20</sup> In low-stakes contexts the norms for action are loose. By contrast, in high-stakes contexts, the norms for action are strict. In neither case is an epistemic norm affected, i.e. a norm logically related to knowledge. You know what you know because you have the evidence you have and the world is the way it is. Whether you can rationally act on your knowledge is quite another matter. Decision theory unites the epistemic and the practical in a theory of rational action while keeping them conceptually distinct as one would expect. Perhaps a final case will help make clear that pragmatism misses this truth.

Consider a doctor D who is the worlds leading expert on the always fatal (if untreated) disease X. There is a quick, inexpensive, and sure test a lab can perform to determine whether a subject has X. It takes two minutes, is perfectly sensitive, and has no false positives. D examines a patient P who presents classic symptoms of some other disease Y which non-specialists sometimes confuse with X. On the basis of her examination D is quite sure<sup>21</sup> that P does not have X but rather Y. Let’s assume that P does not have X. Intuitively, D then knows that P does not have X.<sup>22</sup> However, D should not act on that knowledge immediately, for if she is wrong, the patient dies. Acting on her knowledge would mean not performing the test, but the consequences of being wrong are just too high. It would be wrong not to perform the test. I think this is a case where a proposition *p* is known by a subject at a time, yet it would not be appropriate to act as if *p* at that time. So I think it is a counter-example to (PP).

Notice in this case that it would also be inappropriate for D to assert her knowledge. We would probably wince if she said “Congratulations, you don’t have X” before she ran the test. Since the test is simple and precise, she should wait to flat out assert that the patient has no worries until the test results are in.<sup>23</sup> This suggests that the norm of assertion is also context sensitive. In low-stakes cases the norm of assertion can be as low as *assert only what you believe*<sup>24</sup>; in others, such as the case we just considered where the stakes are very high, the norm can be as high as *assert only that of which you are practically certain*. A wide range of cases will fall in between, and things like justified belief or knowledge will be the norm of assertion. Of course, when the results

---

<sup>20</sup> In one trivial respect, of course, there is one invariant norm for action: so act as to maximize expected utility. Clearly, though, we are dealing with a threshold: how much confidence is required for action or assertion. I am arguing that this threshold is not invariant, but contextual.

<sup>21</sup> Assume, just to keep things simple, that her confidence exactly fits the strength of her evidence.

<sup>22</sup> We will charitably assume that S is ungettiered.

<sup>23</sup> Again, we are supposing that this test takes two minutes, is painless and foolproof.

<sup>24</sup> Suppose your friend truthfully tells you he has a quarter behind his back in one hand and asks you which hand you think it’s in. For some reason unknown to you, you have a basic belief that it’s in the left hand (perhaps because you’re a southpaw yourself, but this does not occur to you). You would be in violation of no norms in giving a truthful answer.

come back as expected, S will say, "I knew it all along." The same is true when the couple calls the bank and finds them open on Saturday and when the other couple confirms the stove was indeed off. *They knew it all along*, they just couldn't appropriately *act* on it, nor could they *assert* it. They couldn't act on their knowledge or assert it because the consequences of wrongly doing so were so high that the norms of assertion and action were correspondingly quite strict. Let's call this view Norm Sensitivity. Let's call the opposite view, the view that there is one norm of assertion or action regardless of context, Norm Invariantism. Hawthorne's Sensitive Invariantism has it that norms are invariant but that knowledge comes and goes (with astonishing ease) with the stakes. Norm Sensitivity and Simple Moderate Invariantism go together to show that it is *norms* that are sensitive to context, not *knowledge*. It thus puts things back in their natural order.