



Epistemic Norms: New Essays on Action, Belief, and Assertion

Clayton Littlejohn and John Turri

Print publication date: 2014

Print ISBN-13: 9780199660025

Published to Oxford Scholarship Online: May 2014

DOI: 10.1093/acprof:oso/9780199660025.001.0001

Intellectual Flourishing as the Fundamental Epistemic Norm

Berit Brogaard

DOI:10.1093/acprof:oso/9780199660025.003.0002

Abstract and Keywords

According to the extended knowledge account of assertion, we should only assert and act on what we know. Call this the 'Knowledge Norm'. Because moral and prudential rules prohibit morally and prudentially unacceptable actions and assertions, they can, familiarly, override the Knowledge Norm. This, however, raises the question of whether other epistemic norms, too, can override the Knowledge Norm. The present chapter offers an affirmative answer to this question and then argues that the Knowledge Norm is derived from a more fundamental norm that demands that we do not hinder intellectual flourishing. As the fundamental epistemic norm can come into conflict with the Knowledge Norm, it is sometimes permissible to assert and act on what we don't know. The chapter concludes with a discussion of the consequences of this insight for the extended knowledge account of assertion.

Keywords: virtue epistemology, intellectual flourishing, knowledge, epistemic norm, assertion

According to the extended knowledge account of assertion, we should only assert and act on what we know. Call this the 'Knowledge Norm'. Because moral and prudential rules prohibit morally and prudentially unacceptable actions and assertions, they can, familiarly, override the Knowledge Norm. This, however, raises the question of whether other epistemic norms, too, can override the Knowledge Norm. The present chapter offers an affirmative answer to this question and then argues that the Knowledge Norm is derived from a more fundamental norm that demands that we do not hinder intellectual flourishing.

As the fundamental epistemic norm can come into conflict with the Knowledge Norm, it is sometimes permissible to assert and act on what we don't know. The chapter concludes with a discussion of the consequences of this insight for the extended knowledge account of assertion.

1.1 The extended knowledge account

According to the extended knowledge account of assertion, you should assert and act only on what you know. Call this the 'Knowledge Norm'. The relevant instances of the Knowledge Norm for assertion and action are the following:

Knowledge Norm of Assertion

You should assert that p only if you know that p ¹

(p.12) *Knowledge Norm of Action*

You should treat the proposition that p as a reason for acting only if you know that p ²

In the case of assertion, the Knowledge Norm is motivated by a wide range of linguistic data:

1. If I assert that p , it is reasonable for you to ask "How do you know?" But this question presupposes that I know that p . The knowledge account can explain why this question is apt (Slote 1979, Unger 1975: 263–4). If knowledge is the norm of assertion, then I should not assert p if I don't know it. As we normally assume that people act correctly epistemically, we normally assume that people know what they assert.
2. If I assert that p , and you later discover that I didn't know p , you will likely consider me dishonest and resent my action (Unger 1975: 260–3). The knowledge account can explain this. If I don't know that p , I don't have permission to assert p . So, I have acted wrongly epistemically.
3. If I assert that p and you respond with "Do you really know that?", you are intuitively challenging my assertion (Williamson 2000: 252–3). If the knowledge account is true, we can explain this intuition as follows. By asking "Do you really know that?", you are questioning whether I have the properties I need to have permission to assert.
4. If I ask you "Is p the case?", and you don't know p , it is appropriate for you to respond with "I don't know". The knowledge account can explain why this response is appropriate (Reynolds 2002: 140). If knowledge is the norm of assertion, then "I don't know" informs me that you lack the properties required to assert that p and lack the properties required to assert that not- p .
5. "I don't know" and "I can't tell" or "I can't say" are equally good replies to "Is p the case?" (Turri 2011). The knowledge account can explain this. According to Turri, 'tell' and 'say' both mean 'assert', and

'can' expresses the concept of permission. So, "I can't say" just means that I lack permission to assert. So, if knowledge is the norm of assertion, then "I can't say" just means that I lack knowledge.

6. The knowledge account can explain why the Moorean assertion "p but I don't know p" is inappropriate (Unger 1975: 256–60, Williamson 2000: 253–4). If knowledge is the norm of assertion, then my assertion of "p but I don't know p" is appropriate only if I know that (p and I don't know that p). But 'know' distributes over conjunction. So, my assertion is appropriate only if I know that p, and I know that I don't know that p. As knowledge is factive, my assertion is appropriate only if I know that p, and I don't know that p.

(p.13) As Matthew Benton (2011) points out, the Moorean predicament can also arise in a context that raises a challenge to the speaker. Consider:

A:

It is snowing.

B:

How do you know?

A:

Oh, I don't.

B:

Huh??

A:

Still, it's snowing.

Benton goes on to argue that the Knowledge Norm is further motivated by the fact that this is the only norm that can offer a unified explanation of both challenge questions and the Moorean conjunction.

Similar considerations motivate the knowledge norm of action. John Hawthorne and Jason Stanley (2008) argue that when we act on reasons we don't know, we are subject to negative appraisal. If a chef takes a cake out of the oven without knowing whether it is done, it casts doubt on his competence, regardless of whether the cake is in fact done, and if a doctor uses a needle he does not know is safe, he is likely to be considered negligent. Or consider the following case: Hannah and Sarah are looking for a restaurant, at which they have reservations. Hannah goes on her hunch that the restaurant is down a street on the left. After walking for a while, they realize that they went down the wrong street. A natural way for Sarah to respond to Hannah's decision is to say, "You shouldn't have gone down this street, since you didn't know that the restaurant was here". These 'ordinary folk appraisals of the behavior of others', they argue, suggest

that whether actions are rational depends on whether the agent knows the reasons for which he is acting.

John Turri (2011) has recently defended a variation on the knowledge account, which he calls the 'Express Knowledge Account of Assertion'. According to the Express Account, you should assert S only if S expresses knowledge. The motivation for this variation comes from cases in which someone knows p and asserts p but the assertion nonetheless is defective.

Turri provides two examples to illustrate. In the first, Randy correctly remembers that p. But the information is stored in memory, so Randy is not conscious of it. Randy has a habit of offering arbitrary answers when people ask him something. So, when his friend asks him whether p, Randy doesn't reflect on the question but just asserts p. Randy asserts something he knows but his assertion is defective, as its correctness in no way depends on his knowledge.

In the second example, Spiro is a spiteful guy who plans to tell Lois Lane that Superman has just died. Just before he is about to tell her, he receives a text message saying that Superman has just died. Wishing to cause Lois emotional pain, he tells **(p.14)** her about her fiancé. Spiro asserts something he knows but his assertion is defective, as he would have said the same if he hadn't received the text message.

The reason the assertions in the two examples are defective is that the assertion the agent makes is not grounded in the knowledge he has. To fill this gap, Turri requires that the assertion expresses the knowledge in question.

I think Turri makes a good case for replacing the original knowledge account with the express knowledge account. But nothing I say here hinges on whether we take the Knowledge Norm to be the original one or Turri's variation. So, I will continue to focus on the original version.

Clayton Littlejohn (2010) argues that K is the norm of belief if it's the norm of assertion. He doesn't defend the view that the knowledge norm is the fundamental norm. But if K is the norm of assertion, then it follows that proposition p is a reason for you to believe q only if you know p. What I say here has no direct bearing on theories of belief norms and for good reasons. It may well be that I should know the reasons for which I believe but that it nonetheless is permissible for me to assert things I don't know. So, in this chapter I shall remain neutral on whether or not knowledge is the norm of belief.

An important question about the Knowledge Norm is what the nature of the modal 'should' is. It is easy to see that the modality is not an all things considered 'should'. If you are in my house hiding from an evil dictator, and his guards ring the door bell and ask me if you are there, I should lie to them, all things considered. So, I should assert something I don't know, viz., that you are

not there, all things considered. The modality that features in the knowledge account is not an all things considered 'should' but an epistemic 'should'. If I am epistemically virtuous, I tell the guards the truth, but ethical considerations override my epistemic obligations. So, in these circumstances, I should not be epistemically virtuous, all things considered.

Or suppose you lend your book to someone who mistakenly thinks that your name is Franz Kafka. You know that the person believes that your name is Franz Kafka, and you know that he is not capable of changing his beliefs about names. In the envisaged scenario, it is prudent to inscribe the sentence 'This book belongs to Franz Kafka' inside the book, if you want it returned. So, in these circumstances, it is prudent to assert what you don't know, viz., that the book belongs to Franz Kafka, if you want the book returned.

Here I argue that it is not only ethical and prudential considerations that can override the Knowledge Norm. Conflicting epistemic considerations, too, can override it. This, I argue, is because the Knowledge Norm is a derived epistemic norm, not a fundamental one. So, it doesn't always take priority in decisions about what to assert or not to assert. We sometimes are epistemically permitted to **(p.15)** assert what we don't know. Similarly, we are sometimes epistemically permitted to act on reasons we don't know. So, I argue, the extended knowledge account is mistaken.

1.2 Epistemic norms

Epistemic norms are norms that constrain ways of going about forming and revising beliefs and using beliefs in theoretical reasoning and communication with others. Epistemic norms include principles such as 'Don't believe false propositions', 'Don't form beliefs for which you have no evidence', 'Use valid principles in inferences', and 'Don't impart false information to others'. Most epistemic norms are derived norms. 'Don't believe false propositions', for example, is likely to be a derived norm. How many fundamental epistemic norms there are is the subject of debate (Brogaard 2009). Epistemic value monists hold that there is just one fundamental norm, whereas epistemic value pluralists hold that there is more than one. I shall assume without argument that epistemic value monism is true.

Here I shall defend the hitherto overlooked view that the fundamental epistemic norm is not to hinder intellectual flourishing. Call this norm 'Intellectual Flourishing' for short. Its relevant instances are the following:

Intellectual Flourishing (belief)

You should believe p only if believing p does not hinder intellectual flourishing

Intellectual Flourishing (assertion)

You should assert p only if asserting p does not hinder intellectual flourishing

Intellectual Flourishing (action)

You should treat p as a reason for action only if treating p as a reason for action does not hinder intellectual flourishing

Intellectual flourishing is the epistemic equivalent of Aristotle's eudaimonia (well-being, flourishing, happiness). For Aristotle, eudaimonia requires having a virtuous character, being loved, and having close friends. If we extend this idea to intellectual flourishing, then intellectual flourishing might involve such things as being intellectually virtuous, being respected intellectually, and having good intellectual cohorts.

Just as we cannot flourish, in Aristotle's sense, in solitude, so we cannot flourish intellectually outside of a community. Intellectual flourishing differs in this respect from knowledge acquisition. While a brain in a vat that is not properly connected to other individuals could, in principle, acquire knowledge as well as you and me, it cannot flourish intellectually.

(p.16) Refraining from acquiring knowledge is sometimes needed to flourish intellectually. Suppose A hears of a new proof that God does not exist. A knows that if he sees the proof and the proof is correct, he will become terribly depressed and will spend the rest of his life in isolation from intellectual cohorts. To ensure that he flourishes intellectually, A must refrain from looking at the proof.

Determining what is the right thing to do, intellectually speaking, requires reflection on the intellectual consequences of believing, asserting, and acting in the particular situation in which the belief, assertion or action is to take place. More on that later.

Consider an alternative to Intellectual Flourishing commonly thought to be the fundamental epistemic norm, viz., the Truth Norm:

Truth Norm

For any proposition p , you should believe that p iff p is true

The Truth Norm, as formulated, is problematic. As Lynch (2009) has argued, we cannot believe everything that is true, and it is highly doubtful that our epistemic goal is to be God. Lynch suggests that we should revise the standard formulation of the Truth Norm to take account of our human deficiencies. The Truth Norm, he suggests, may be formulated as follows. For any p that we are able to consider, you should believe that p iff p is true. Lynch admits that it might be difficult to unpack the modality of 'is able to' but adds that whatever it

means, it is not supposed to mean that we should believe only the true propositions that we happen to consider.

In my view, however, the main problem with the Truth Norm is not that it is practically impossible for us to believe all true propositions but rather that it is at odds with Intellectual Flourishing. To see that Intellectual Flourishing and the Truth Norm are in conflict, consider a case that I will call ‘Brain Damage’:

Brain Damage

A has a brain condition that causes him to intend to keep track of truths about leaves. He believes that he can achieve this only if he intends to count the leaves on the trees in his garden every day.

If A does what he believes is necessary for him to intend to keep track of truths about leaves, and he is a good counter, his intention is likely to maximize true belief and minimize false ones. If he didn’t intend to count leaves, he would go about his everyday business forming many more false beliefs than he does if he is just counting leaves all day. But intuitively, intellectual flourishing requires having a wide range of different kinds of true beliefs. So, A hinders intellectual flourishing by intending to count leaves, despite the fact that his intention **(p. 17)** maximizes true beliefs and minimizes false ones. So, the Truth Norm is at odds with Intellectual Flourishing.³

Intellectual Flourishing is an epistemic norm just like principles such as ‘Don’t form beliefs while drunk’, ‘Be open-minded’, and ‘Avoid the gambler’s fallacy’. But there is good reason to think that it, unlike ‘Don’t form beliefs while hallucinating’, ‘Be open-minded’, and ‘Avoid the gambler’s fallacy’, is the fundamental epistemic norm. If we were to give epistemic advice to A in the Brain Damage case, we would instinctively advise him to stop intending to count leaves and start intending to use his time in a less futile way. So, in the envisaged scenario, we would attach more epistemic value to intentions to engage in activities that violate the Truth Norm than we would to intentions to engage in activities that satisfy it.

But if Intellectual Flourishing is the fundamental norm, and it looks like it may be, then all other epistemic norms are derivative. The Truth Norm is good to keep in mind when deliberating because satisfying this principle does not normally hinder intellectual flourishing. Likewise, asserting only what you know is normally a good norm to obey because obeying it does not normally hinder intellectual flourishing. But both the Truth Norm and the Knowledge Norm are derivative epistemic norms, or what some would call ‘*ceteris paribus* laws’, and so can come into conflict with Intellectual Flourishing.

1.3 Intellectual vs. theoretical rationality

Intellectual Flourishing is a good candidate to be the fundamental epistemic norm. However, a puzzle arises here. The following principle is plausible:

The Instrumental Principle

You should, epistemically, [if you intend to do E, and believe that you will do E only if you do M, intend to do M].

(p.18) The Instrumental Principle is an optimized version of Kant's instrumental principle to the effect that if you intend to do E, and you believe that you will do E only if you do M, then you should, all things considered, intend to do M. The Kantian way of reading the instrumental principle, however, runs into trouble. If I intend to smoke and believe smoking requires buying cigarettes, then I should buy cigarettes. Supposing I intend to smoke and believe smoking requires buying cigarettes, it follows that I should, all things considered, intend to buy cigarettes. But I am under no (unconditional) obligation to intend to buy cigarettes.

John Broome (2000, 2001, 2002, 2004) argues that we can bypass this worry by giving the Instrumental Principle a wide-scope reading rather than a narrow-scope reading. This gives us: you should, all things considered, [if you intend to do E, and believe that you will do E only if you do M, intend to do M].

On the wide-scope reading, you can satisfy the wide-scope Instrumental Principle by not intending to do E, by giving up on the belief that you will do E only if you do M or by intending to do M. So, we cannot infer from the wide-scope Instrumental Principle and the premise that I intend to smoke and believe that I will smoke only if I buy cigarettes, that I should intend to buy cigarettes, all things considered.

However, as Kieran Setiya (2007) has argued, Broome's move gets us out of the frying pan and into the fire. This is because 'should' transmits across necessity:

Transmission

If you should do E, all things considered, and doing M is a necessary means to doing E, you should do M, all things considered

I should, all things considered, satisfy the Instrumental Principle. But suppose as a matter of necessity I am incapable of changing my intention to smoke and my belief that I will smoke only if I buy cigarettes. I can then satisfy it only by intending to buy cigarettes. It follows from Transmission that I should intend to buy cigarettes, all things considered. But intuitively, I should not intend to buy cigarettes, all things considered.

Setiya suggests that we can solve the puzzle by reinterpreting the ‘should’ in the Instrumental Principle as the epistemic ‘should’.⁴ Since Transmission does not apply to the epistemic ‘should’, he says, we cannot infer from the claim that the only way to satisfy the instrumental principle is to intend to buy cigarettes that I should intend to buy cigarettes, all things considered.

(p.19) Unfortunately, this move does nothing to solve the puzzle. While it is true that Transmission does not apply to epistemic ‘should’, there is an epistemic version of Transmission that does:

Epistemic Transmission

If, epistemically, you should satisfy epistemic principle P, and doing M is a necessary means to satisfying P, you, epistemically, should do M

Epistemic Transmission is no less plausible than Transmission. If you, epistemically, should assert only what you know, and not asserting that the earth is flat is a necessary means to asserting only what you know, then you, epistemically, should not assert that the earth is flat.

But consider now the earlier case of Brain Damage. In the envisaged circumstance, A, epistemically, should [if he intends to keep track of truths about leaves and believes that he can achieve this only if he counts the leaves on the trees in his garden every day, then he intends to count the leaves on the trees in his garden every day]. He is required to satisfy the Instrumental Principle. But because his brain damage causes him to intend to keep track of truths about leaves, he can satisfy the Instrumental Principle only by intending to count the leaves on the trees in his garden every day. So, epistemically, he should intend to count the leaves on the trees in his garden every day. But this conflicts with Intellectual Flourishing. Intellectual Flourishing requires A to do something other than intend to count leaves on the trees in his garden every day.

What to do? Epistemic Transmission, Intellectual Flourishing and the Instrumental Principle all seem exceedingly plausible. Upon closer scrutiny, it becomes clear that the culprit is a hidden ambiguity in the modality in the three principles. The ‘should’ in the Instrumental Principle and Epistemic Transmission is not the same ‘should’ as that in Intellectual Flourishing.

Theoretical rationality regulates belief and acceptance. A belief or an intellectual action is theoretically rational only if there is evidential and argumentative support in favor of holding it or doing it. This kind of rationality is distinct from what I will call ‘intellectual rationality’. A belief or an intellectual action is intellectually rational only if holding the belief or performing the action does not hinder intellectual flourishing. That is, intellectual rationality requires living the

best possible intellectual life and achieving the most important intellectual goals.

Once we tease apart these two different systems of norms, the theoretical and the intellectual, the puzzle does not arise. Because A cannot avoid forming the intention to keep track of facts about leaves and cannot avoid forming the belief that intending to keep track of facts about leaves requires intending to count the **(p.20)** leaves on the trees in his garden every day, it is indeed a theoretical requirement that he intend to count the leaves on the trees in his garden every day. Intending to engage in this futile exercise is theoretically rational insofar as it does not hinder true belief and does not promote false belief. But the fact that intending to count leaves is theoretically rational is not at odds with Intellectual Flourishing, which requires him to do something other than count leaves on the grounds that not doing something else hinders intellectual flourishing.

Neither the theoretical 'should' nor the intellectual 'should' is an all things considered 'should'. However, in this case, the intellectual requirement wins because it is in A's best interest not to count leaves on the trees in his garden every day. So, in this case, A should not count leaves on the trees in his garden every day, all things considered. But there could be other scenarios in which the theoretical requirement wins, for example a case in which A must kill three people or satisfy the theoretical requirement. Theoretical rationality here trumps intellectual rationality.

1.4 Asserting what you don't know

If there are two different systems of norms for regulating belief and assertion, the theoretical and the intellectual, to which system does the Knowledge Norm belong? The answer to this question is straightforward. The Knowledge Norm is derived from Intellectual Flourishing. The Knowledge Norm is not concerned with true belief. So, there is nothing theoretically mandatory about asserting only what you know. What makes it mandatory to assert only what you know is that, in by far the most cases, this is the best way not to hinder intellectual flourishing.

Because it is a derived epistemic norm, the Knowledge Norm is bound to be overridden by the fundamental epistemic norm in special circumstances. Milgram's classic experiment provides one such case. In the experiment, participants were recruited via ads for a 'study of memory' at Yale. At the site, the experimenter told the participant and a confederate, an actor pretending to be another participant, that they would be participating in an experiment to test the effects of punishment on learning behavior. The participant was led to believe that he was 'randomly' chosen to be the 'teacher', whereas the other was chosen to be the 'learner'. The participant was then given a real 'sample' 45-volt electric shock from the electro-shock generator and a list of word pairs that he would teach the learner. After reading the words, the teacher would read half of

the word pairs and provide the learner with four possible answers. The learner would indicate which second word he believed to be correct by pressing a button **(p.21)** (1 through 4) corresponding to the teacher's choices. If incorrect, the learner would receive 'a shock', increasing by 15 volts for each wrong answer. In each case, the participant was told that he was actually giving shocks to a learner participant, but there was no real shock. At 135 volts the learner would pretend to be in agony. At 300 volts the learner would bang on the wall and complain about his heart condition. After that, the learner would stop complaining and stop answering the questions.

Most participants expressed a wish to stop the experiment and check on the subject at this point. Some would question the purpose of the experiment. Some continued after being assured that they would not be held responsible. Some began to laugh nervously after hearing the screams of pain from the learner. Some went to the end of the shocks (450 volts), but everyone stopped at some point and questioned the experiment. Each time the participant stopped, the experimenter would say "The experiment requires that you continue. Please go on."

Milgram's experiment is one of the most famous experiments involving deceit. Deceit is necessary in experiments that study how people react to stimuli which they would have reacted differently to, had they had knowledge of the protocol. To be successful, experiments involving deceit require asserting falsehoods to the participants. Milgram's experiment, for example, required telling each participant: that he would be participating in an experiment to test the effects of punishment on learning behavior; that he was randomly chosen to be the teacher; and that the other participant was really given electric shocks.

In Milgram's experiment, the experimenter's assertions are genuine assertions, not quasi-assertions of the kind actors make. The experimenter is simply lying to the participant. And lies, albeit theoretically deplorable, are normally genuine assertions.

Despite being theoretically deplorable, the experimenter's assertions do not hinder intellectual flourishing. They are required in order for him to be able to conduct his study of how willing people are to engage in morally bad behavior in the presence of an authority. As asserting what you don't know in this kind of scenario does not hinder intellectual flourishing, it is not required in this kind of scenario that people observe the Knowledge Norm.

Another example of a case in which Intellectual Flourishing trumps theoretical rationality is one in which a journalist purposely comes up with a newspaper headline that is either outright false or else misleading in order to help impart information. Consider a case of a journalist who chooses the following headline for her story:

(p.22) *Scientists discover a way to reverse memory loss in ‘accidental breakthrough’*

In the article in question, it was revealed that, as the researchers were attempting to reverse an obesity problem in a patient using deep-brain stimulation, the patient suddenly had vivid memories of an event that occurred thirty years earlier. The researchers subsequently employed the method to improve the memory of a 50-year-old. The neurosurgeon Professor Andres Lozano was quoted as saying that the treatment could help scientists understand how memory works, and Susanne Sorensen, of the Alzheimer’s Society, was quoted as saying: “We hope the findings can be replicated and eventually lead to new treatments.”

In this case, the headline does not express a truth. There is indicative evidence that deep-brain stimulation can improve memory in people with Alzheimer’s. But scientists have not discovered a way to reverse memory loss. However, the chosen formulation of the headline makes it more likely that readers will read on. Writing the truth in the headline would hinder intellectual flourishing. It is thus epistemically advisable in the envisaged circumstances to disobey the Knowledge Norm.

Consider a further case:

Faculty Meeting

I want to organize a faculty meeting in the first two weeks of August and email my colleagues to inquire about their availability. Faculty member A replies that he is in Rome the first two weeks of August and that we can go ahead without him. I vaguely recall him saying on an earlier occasion that he would be in Rome the last two weeks of August. Knowing A’s attitude toward faculty meetings, I suspect he is lying to avoid going. I also know that if I bring up the issue or convey it to others in the department, it will cause tensions in the department, and A might show up to the meeting acting overly critical and irritated. So, I go ahead with the meeting. When the others ask about A’s absence, I tell them that he is in Rome.

I only vaguely recall A saying that he would be in Rome the last two weeks of August. So, A could be in Rome the first two weeks of August. But I nonetheless don’t know that he is. So, when I assert that A is in Rome, I assert something I don’t know. This is in violation of the Knowledge Norm. But asserting that I don’t know why he is not at the meeting or asserting that he *told* me he would be in Rome rather than asserting that A is in Rome would hinder intellectual flourishing, because it would cause tensions in the department and prevent the rest of us from making fruitful decisions at the meeting. So, intellectually, it is permissible for me to assert that A is in Rome.

(p.23) 1.5 Revisiting the Knowledge Norm of Assertion

If the Knowledge Norm is a derived epistemic norm that we need not always follow, how do we explain the linguistic data that speak in favor of the Knowledge Norm? The most obvious explanation of why the linguistic data speak in favor of the Knowledge Norm is that Intellectual Flourishing does not override the Knowledge Norm in most circumstances. So, we expect the Knowledge Norm to apply even when it in fact doesn't apply.

While this sort of explanation suffices to explain the linguistic data, there is also independent reason to question that the linguistic data support the Knowledge Norm of Assertion. In ordinary speech, we often use 'to know' synonymously with 'to be able to provide the correct answer when asked'. Consider the following case (Hawthorne 2002):

The Capital of the Netherlands

While co-teaching a class you and I are having a chat about how well American students know the capitals of the European countries. I say to you "I bet all our students know the capital of the Netherlands". You disagree. To decide our disagreement I ask each student "What is the capital of the Netherlands?". Each student replies with "Amsterdam". I triumphantly turn to you and say "See, what I said".

In this context, I don't care about whether each student merely asserts, believes, reliably believes or justifiably believes that Amsterdam is the capital of the Netherlands. All I care about is whether they can provide the correct answer when asked. This suggests that 'to know' is used synonymously with 'to be able to provide the correct answer when asked' in some everyday contexts. This, of course, is not to say that being able to provide the correct answer when asked counts as knowledge but only that 'to know' can be used to express other relations than the knowledge relation.

The everyday use of 'to know' as synonymous with 'to be able to provide the correct answer when asked' can explain the linguistic data as well as the knowledge account. 'P but I don't know p' sounds inconsistent, because 'p but I am unable to provide the correct answer to the question of whether p when asked' is self-refuting. The same holds for Moorean predicaments that occur in challenge conversations:

A:

It is snowing.

B:

How do you know?

A:

Oh, I don't.

B:

Huh??

A:

Still, it's snowing.

(p.24) Assuming the folk reading of 'know', A oddly denies that he was able to provide the correct answer to the question of whether it is snowing when asked and then goes on to provide it a second time.

Recall that Turri argued that the knowledge account can explain why "I don't know" and "I can't tell" or "I can't say" are equally good replies to "Is p the case?" According to Turri, 'tell' and 'say' both mean 'assert', and 'can' expresses the concept of permission. So, "I can't say" means that I lack permission to assert. So, if knowledge is the norm of assertion, then "I can't say" just means that I lack knowledge. This explanation, however, does not seem quite right to me. The 'can' in the relevant expression expresses the concept of ability, not permission. "I can't say" means that I am unable to say, which is just another way of saying that I am unable to provide the correct answer.

If I assert that p, it is reasonable for you to ask "How do you know?" But you may well just be asking "How did you come by the correct answer?" I can then reply "John told me" or "I looked it up on the Internet".

Likewise, if I assert that p, and you respond with "Do you really know that?", you are indeed intuitively challenging my assertion. A natural response to the challenge would be: "Yes, I do. John told me/I looked it up on the Internet/..." But your challenge may not be for me to reassure you that I know it and provide the source of my knowledge but rather to reassure you that I am able to produce the correct answer and provide the source of my knowledge.

Similar remarks apply to a case in which I ask you "Is p the case?" and you reply with "I don't know". You could mean that you lack the epistemic properties required to assert that p and lack the epistemic properties to assert that not-p. But you could also simply mean that you are unable to produce the correct answer.

Of course, even when the situation calls for a false assertion, we might ask "Do you really know it?" or "How do you know?". For example, the 'teacher' in Milgram's experiment could ask the experimenter "Do you really know that the learner is receiving electric shocks?" But he would be asking this question aptly only while being deceived into thinking that the situation is one that calls for a true assertion.

Similar remarks apply to Faculty Meeting. As in Milgram's experiment, the situation calls for deception, and it is in the nature of deception that the apparent situation looks quite normal. So, it would be epistemically appropriate for one of my colleagues to ask "Do you really know that A is in Rome?"

Journalism settings are different. Because most people are familiar with the aims of journalism, it would be odd to ask a journalist who tweaks the truth in a newspaper headline to catch the reader's attention "Do you really know it?"

(p.25) If I assert that *p*, and you later discover that I didn't know *p*, you may consider me dishonest and resent my action. However, you are unlikely to consider me dishonest and resent my action if I asserted *p* without knowing it was false or if I had evidence for *p* but *p* turned out to be false. Furthermore, in the special contexts we considered earlier in which it is acceptable to assert something false, the speaker is unlikely to be considered dishonest once it is understood that the situation called for a false assertion. The participants in Milgram's experiment, for example, didn't consider the experimenter dishonest after being told about the deceptive element of the study. They understood that it was needed to carry out the study. We are not likely, either, to consider journalists who tweak the truth in newspaper headlines dishonest. We know that headlines have to be written in a certain way to catch people's attention. Likewise, you would not consider me dishonest if you discovered that I told you a white lie in order to preserve peace in the department and be able to have a peaceful department meeting.

One thought that may arise at this point is whether Williamson's distinction between reasonable and permissible assertions (Williamson 2000: 257) and Keith DeRose's distinction between primary and secondary conformity in rule following (2002: 180) can do some work in explaining why the experimenter in Milgram's experiment and my colleagues in Faculty Meeting are unlikely to make an assessment of dishonesty after being told about the deception. According to Williamson and DeRose, an assertion can be reasonable despite being impermissible. This is so when you believe you know what you assert but you don't. For example, you might have superb evidence for thinking that it is raining and assert that it is raining. If it then turns out that it was not raining after all, your assertion is reasonable—we understand why you asserted what you did—but your assertion is impermissible nonetheless. This is analogous to a case in which you falsely believe that you are parking legally, because someone spray painted the street sign. We understand why you didn't move your car but your action is nonetheless illegal.

Can the distinction between a reasonable and a permissible assertion explain why the participants in Milgram's experiment didn't consider the experimenter dishonest after being told about the deceptive element of the study? It cannot. The case is not at all analogous to the case in which I falsely believe I know it is

raining and say it is raining. We consider the experimenter's false assertions reasonable, not because he falsely believes that he is telling the truth, but because telling the truth would hinder intellectual flourishing. Likewise, I do not falsely believe that I am telling the truth when I say that A is in Rome. I say that A is in Rome because not saying this would hinder intellectual flourishing. So, it is not **(p.26)** because my assertion is reasonable despite being impermissible that my colleagues don't consider me dishonest after being told about the white lie.

The Knowledge Norm of Assertion is thus widely applicable but it can be overridden by the fundamental epistemic norm, which requires us not to hinder intellectual flourishing. Which situations call for assertions that violate the knowledge norm depends on the intellectual consequences of asserting what we don't know.

1.6 Reasons for action

The question that remains is whether there are cases in which the Knowledge Norm of Action conflicts with the fundamental epistemic norm. At first it may seem that our earlier cases are examples of cases in which the Knowledge Norm of Action is at odds with the fundamental norm. But this is not so. In Milgram's experiment, the experimenter must act in accordance with the experiment's set-up. The experimenter must act as if another participant receives painful electric shocks. The actions the experimenter performs are genuine actions, and they are actions he, epistemically, has permission to perform, because they do not hinder intellectual flourishing. Furthermore, the experimenter doesn't know that the learner receives painful electric shocks. But he does not treat the proposition that the 'learner' receives painful electric shocks as a reason for acting. Rather, he acts as he does because the experiment calls for deceit. So, his actions do not violate the Knowledge Norm of Action.

The same kinds of considerations carry over to the journalistic setting and Faculty Meeting. The journalist doesn't know that the scientists have found a way to reverse memory loss. A good headline promotes the imparting of information. So, her assertions and actions are intellectually rational. They are assertions and actions she has permission to perform in the relevant setting. But she does not treat the proposition that scientists have discovered a way to reverse memory loss as a reason for formulating the headline the way she does.

Likewise, when my colleagues inquire about A's absence at the faculty meeting, and I reply that he is in Rome and don't grab the phone to call him and don't contact him to ask him if he can make the follow-up meeting next week, I act the way I do because my actions are intellectually permissible, as not doing what I do would hinder intellectual flourishing. But the proposition that A is in Rome is not a reason for any of these actions. So, I am not in violation of the Knowledge Norm of Action.

The reason that the false propositions cannot serve as reasons in the relevant sense is that the agents in our envisaged scenarios don't *truly* believe them. The **(p.27)** relevant sense of 'reason' here is 'justifying' or 'explanatory reason', not 'personal reason'. Unlike justifying reasons, explanatory reasons for actions needn't be reasons the agent believes. For example, I might choose A as a partner because I believe that A is a good guy when in fact my choice is driven by my attraction to his bad behavior. The explanatory reason for why I chose A is that I was attracted to his bad behavior. My reason was that A is a good guy. But as A is not a good guy, and as my choosing A because of my attraction to his bad behavior makes no good normative sense, there are no justifying, or rational, reasons for my choosing A as a partner.

While the false propositions in our examples do not serve as justifying reasons, there are nonetheless grounds on which to doubt that propositions can be justifying reasons only if they are known. On a common account of internalism about reasons, it is a necessary condition on (justifying) reasons that they are motivational (Williams 1989, Dreier 1990, Korsgaard 1996). We can articulate the requirement as follows:

Internalism about Reasons

The fact that p is a reason for A to do E only if A is capable of being moved to do E by the belief that p

This requirement on reasons explains why false propositions cannot be a reason to act and hence explains why 'the "learner" is receiving electric shocks' cannot be a reason for the experimenter to act as if he is, why 'scientists have discovered a way to reverse memory loss' cannot be a reason for the journalist to formulate the headline as she does and why 'A is in Rome' cannot be a reason for me to act as if A is in Rome.

Hawthorne and Stanley (2008) do not discuss Internalism about Reasons. All they say about the nature of 'their' reasons is that they are 'personal reasons'. However, in the standard sense of 'personal reason', a personal reason is the reason the agent would give for his actions if he were being sincere. As the reasons an agent would give for his actions needn't be true, it cannot be in the standard sense of 'personal reason' that Hawthorne and Stanley's reasons are personal reasons. It is plausible that they take reasons to be motivating (and justifying) reasons in the sense specified by Internalism about Reasons. However, Internalism about Reasons does not require that we act only on reasons we know, and there are many cases that suggest that reasons only need to satisfy Internalism about Reasons and not the Knowledge Norm. Consider the following:

Blood Pressure

A scientist, A, truly believes he has discovered a new drug, D, that can treat high blood pressure. After eight years of studying D, he believes he has good **(p.28)** evidence for that belief. As A has high blood pressure, he starts taking D. Years later another team discovers that there were some serious mistakes in A's protocol.

In Blood Pressure, A is motivated to take the drug by the true proposition that the drug can treat high blood pressure. But A does not know the drug can treat high blood pressure. So, the fact that the drug can treat high blood pressure is a reason for A to take the drug, despite the fact that he doesn't know that the drug can treat high blood pressure.

Note that after discovering the mistake in A's protocol, we wouldn't blame A for taking the drug on the grounds that he didn't know that the drug could treat high blood pressure. It seems that he had a perfectly good reason for taking the drug at the time.

Hawthorne and Stanley (2008) may insist that A's treatment of the proposition that the drug can treat high blood pressure was reasonable but nonetheless epistemically impermissible because it was in violation of the Knowledge Norm. However, the ordinary folk appraisals of the behavior of others that they use to motivate the view that 'the concept of knowledge is intimately intertwined with the rationality of action' do not give us a reason to think that A's treatment of the proposition that the drug can treat high blood pressure was impermissible.

Why then do we mention knowledge in our appraisals of people's actions? It seems plausible that we mention knowledge when it is obvious to us that the agent should have known. If a doctor uses a needle that he doesn't know is safe, we criticize his action because he should have known that the needle was safe before using it. His personal reason for his action, whatever it was, was not a justifying reason for his action. So, we criticize him because there was no justifying reason for his action. But, as Blood Pressure shows, we do not mention knowledge uniformly in our appraisals of actions. So, the linguistic data do not suggest that the Knowledge Norm applies uniformly, even if it's a good rule of thumb.

In their paper defending the Knowledge Norm of Assertion, Hawthorne and Stanley (2008) go on to defend a bidirectional knowledge norm as follows:

The Reason-Knowledge Principle

Where one's choice is p-dependent, it is epistemically appropriate to treat the proposition that p as a reason for acting iff you know that p

The left-to-right direction of the Reason-Knowledge Principle is inspired by Jeremy Fantl and Matt McGrath's Knowledge-Action Principle, which says that if

you know that p , then it shouldn't be a problem to act as if p (Fantl and McGrath 2002).

(p.29) The Reason-Knowledge Principle entails straightforwardly that where your choice is p -dependent and you know that p , it is appropriate for you to treat the proposition that p as a reason for acting. However, the Reason-Knowledge Principle conflicts with Intellectual Flourishing.

Recall the case of A who can't help but intend to keep track of facts about leaves. A believes that intending to count the leaves on the trees in his garden every day is necessary for him to intend to keep track of facts about leaves. When someone asks him why he intends to engage in this futile exercise, he replies with "I do it because counting leaves on the trees in my garden every day maximizes true beliefs and minimizes false ones". A knows that counting leaves on the trees in his garden every day maximizes true beliefs and minimizes false ones. So, by the Reason-Knowledge Principle, it is epistemically appropriate to treat the proposition that counting leaves on the trees in his garden every day maximizes true beliefs and minimizes false ones as a reason for intending to count leaves. But treating this proposition as a reason hinders intellectual flourishing. So, epistemically, A should not treat this as a reason for counting leaves.

Note that it does not help if A cites the cause of his behavior as a reason instead. The proposition that he cannot help himself is not an epistemically appropriate reason for him to intend to count leaves.

Consider another case:

Sexism

In a 20,000 participant, longitudinal study of the location of the hippocampus, Scientist A incidentally discovers that the hippocampus is significantly larger in men than in women. A knows that a larger hippocampus doesn't entail more memory capacity. He also knows that publishing the results will lead to misinterpretations by popular media and will lead grant agencies to fund futile projects on differences in intelligence between men and women, which will leave less funds for more fruitful projects. However, A is up for tenure, and he knows that publishing the results will help his tenure case. So, he publishes the results.

In this case, A knows that publishing the results will help his tenure case. But publishing the results will lead to misinterpretations, and it will likely prevent fruitful scientific projects from getting funded by grant agencies. So, publishing the results hinders intellectual flourishing. If Intellectual Flourishing is more fundamental than the Knowledge Norm, which—as we have seen—is a reasonable assumption, then it is not epistemically appropriate for A to treat the

proposition ‘publishing the results about the differences in the size of hippocampus will help my tenure case’ as a reason for publishing the results.

(p.30) As these cases make evident, the Reason-Knowledge Principle can come into conflict with the fundamental epistemic norm. When it does, the Reason-Knowledge Principle is overridden by the fundamental norm.

In conclusion, actions and assertions are licensed and prohibited not only by ethical and prudential norms but also by epistemic norms. This is the insight underlying the Knowledge Norm. Other things being equal, you should assert only what you know, and other things being equal, you should treat the proposition that *p* as a reason for acting only if you know that *p*. However, when the Knowledge Norm conflicts with the fundamental epistemic norm, the fundamental epistemic norm is the deciding factor in determining whether the action is licensed or prohibited.⁵

References

Bibliography references:

Benton, M. A. (2011). “Two More for the Knowledge Account of Assertion,” *Analysis* 71: 684–7.

Brogaard, B. (2009). “The Trivial Argument for Epistemic Value Pluralism. Or How I Learned to Stop Caring About Truth,” in A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value* (Oxford: Oxford University Press), 284–308.

—(2012). “Wide-Scope Requirements and the Ethics of Belief,” in J. Matheson and R. Vitz (eds.), *The Ethics of Belief* (Oxford University Press, Forthcoming).

Broome, J. (2000). “Normative Requirements,” in J. Dancy (ed.), *Normativity*, (Oxford: Blackwell, 2000), 78–99.

—(2001). “Normative Practical Reasoning,” *Proceedings of the Aristotelian Society* Supplementary Volume 75: 175–93.

—(2002). “Practical Reasoning,” in J. L. Bermudez and A. Millar (eds.), *Reason and Nature* (Oxford: Oxford University Press), 85–111.

—(2004). “Reasons,” in R. J. Wallace, P. Pettit, S. Scheffler, and M. Smith (eds.), *Reason and Value: Essays on the Moral Philosophy of Joseph Raz* (Oxford: Oxford University Press), 28–55.

DeRose, K. (2002). “Assertion, Knowledge and Context,” *Philosophical Review* 111: 167–203.

Dreier, J. (1990). “Internalism and Speaker Relativism,” *Ethics* 101: 6–26.

Fantl, J. and McGrath, M. (2002). "Evidence, Pragmatics, and Justification," *Philosophical Review* 111: 67-94.

Harman, G. (1997). "Practical Reasoning," reprinted in A. Mele (ed.), *The Philosophy of Action* (Oxford: Oxford University Press), 149-77.

Hawthorne, J. (2002). "Deeply Contingent A Priori Knowledge," *Philosophy and Phenomenological Research* 65: 247-69.

(p.31) Hawthorne, J. (2004). *Knowledge and Lotteries* (Oxford: Oxford University Press).

— and Stanley, J. (2008). "Knowledge and Action," *The Journal of Philosophy* 105: 571-90.

Korsgaard, C. (1996). *The Sources of Normativity* (Cambridge: Cambridge University Press).

Littlejohn, C. (2010). "Moore's Paradox and Epistemic Norms," *Australasian Journal of Philosophy* 88: 79-100.

Lynch, M. P. (2009). "The Values of Truth and the Truth of Values," in A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value* (Oxford: Oxford University Press), 225-42.

Reynolds, S. L. (2002). "Testimony, Knowledge, and Epistemic Goals," *Philosophical Studies* 110: 139-61.

Setiya, K. (2007). "Cognitivism about Instrumental Reason," *Ethics* 117, Symposium on Education and Equality: 649-73.

Slote, M. (1979). "Assertion and Belief," in J. Dancy (ed.), *Papers on Language and Logic* (Keele: Keele University Library), 177-91. Reprinted in Slote (2010): *Selected Essays* (New York: Oxford University Press), 94-102.

Turri, J. (2010). "Prompting Challenges," *Analysis* 70: 456-62.

— (2011). "The Express Knowledge Account of Assertion," *Australasian Journal of Philosophy* 89: 37-45.

Unger, P. (1975). *Ignorance: A Case for Skepticism* (Oxford: Oxford University Press).

Wallace, R. J. (2001). "Normativity, Commitment and Instrumental Reason," *Philosophers' Imprint* 1: 1-26.

Williams, B. (1989). "Internal Reasons and the Obscurity of Blame." Reprinted in B. Williams (1995): *Making Sense of Humanity* (Cambridge: Cambridge University Press).

Williamson, T. (2000). *Knowledge and its Limits* (Oxford: Oxford University Press).

Notes:

(1) Defenders of the Knowledge Norm of Assertion include Unger (1975), Williamson (2000), DeRose (2002), Hawthorne (2004), and Turri (2010, 2011).

(2) Hawthorne and Stanley (2008) defend this version of the knowledge account of action.

(3) In correspondence, Clayton Littlejohn has raised the following worry about Intellectual Flourishing. On some accounts, an agent cannot rationally (and wittingly) believe a proposition in awareness that she's in violation of an epistemic norm. That does not appear to be so in the case of Intellectual Flourishing. Whether that is problematic would seem to depend on what the point of norms is.

I deal with this issue further down, where I draw a distinction between theoretical and intellectual aims. On this account, there are two forms of rationality: Theoretical and intellectual. So, an agent can rationally believe a proposition in the theoretical sense of 'rational' but not in the intellectual sense of 'rational' while knowing that she is in violation of Intellectual Flourishing.

Lynch (and other truthers) could say in response to my worries about their norms that their concerns are pitched at a different level than mine. But they would then need to hold that their norms are purely theoretical and hence relatively uninteresting from the point of view of asserting and acting. For related discussion see Brogaard (2012).

(4) Gilbert Harman (1997) and Jay Wallace (2001) have also argued for an epistemic reading of the Instrumental Principle.

(5) This paper was given as the presidential address at the 2011 meeting of the Central States Philosophical Association. Thanks to Jeff Dauer, John Greco, John Hawthorne, John Heil, Clayton Littlejohn, and Michael Neal for helpful comments on an earlier version of this paper.

Access brought to you by: