Can There Be A Knowledge-First Ethics of Belief?

Dennis Whitcomb[1]

1 Stage Setting

What should we believe, and which of our beliefs are justified? These questions, both important, are distinct. You might look at your hands and believe that you have them — but believe this simply on a whim, and not because you see them. Even though you believe what you should, your belief is unjustified. Thus we should distinguish propositions one should believe from justified beliefs. Additionally, we should distinguish beliefs from withholdings of belief. We should sometimes withhold, just as we should sometimes believe; and withholdings are sometimes justified, just as beliefs are sometimes justified.

The job of the ethics of belief is to delineate the cases where we should believe and where we should withhold, as well as the cases of justified belief and justified withholding. In this paper I'll be exploring the first part of this task - the theory of what we should believe and what we should withhold belief on.

I'll explore several knowledge-first theories about these issues, that is to say, several theories that use the notion of knowledge to delineate the cases in which we should believe and in which we should withhold belief. It is an open question to what extent we should take knowledge first in our epistemological theorizing. If a knowledge-first theory of what we should

believe (and what we should withhold belief on) can be made to work, then that is some reason to adopt the knowledge-first approach to epistemology as a whole. If no such theory can be made to work, then that is some reason to reject the knowledge-first approach to epistemology as a whole.

I won't converge on a view about which one of these two options is correct. Instead, I will explore the problems and prospects of the approach. The exploration will proceed via closely engaging numerous attempts to build the knowledge-first theories in question. Most of those attempts don't work, for reasons I will bring out. Others might work, depending on how their details end up; this too I will bring out.

When I inquire into what we should believe, I mean to ask that question in a specifically epistemic sense. To get a fix on this sense, put yourself in the position of an undergraduate who has just learned about Pascal's wager. Imagine you have the following common reaction to the wager: "OK, sure, that shows that believing in God serves my interests. But when I asked whether I should believe in God, I didn't want to know whether that belief would serve my interests. I wanted to know something else....." Undergraduates regularly have this reaction to Pascal's wager. And they are regularly relieved to be told that there is a particular thing they are looking for: *epistemic* (as opposed to practical) reasons to believe in God, and in particular epistemic reasons strong enough that they should, from an epistemic point of view, believe. It is the particular "should" delineated here, that I focus on in the current paper. It is the *epistemic* should.

If you are inclined to deny the existence of this "epistemic" should, then think about it this way instead. There is such a thing as propositional justification. To have propositional justification to believe a proposition is to have good enough epistemic reason to believe it. Similarly, to have good enough propositional justification to withhold on a proposition, is to have good enough epistemic reason to withhold on it. What we should believe, in my sense - what I am calling the epistemic sense - is just what we have propositional justification to believe. Similarly, what we should withhold belief on - in the epistemic sense - is what we have propositional justification to withhold belief on.

In the terminology of "propositional justification", then, the current paper is searching after a knowledge-first theory of what we are propositionally justified to believe and withhold belief on. To my ear, "should believe" talk

expresses the same issues more naturally than "propositional justification" talk, and so I will stick to the former.

2   Identity Theories

Our task is to build a knowledge-first theory of what we should believe and what we should withhold belief on. The most straightforward way to discharge this task is to simply *identify* what we should believe with what we know. Timothy Williamson comes close to advocating such an identification in the following passage:

> ...the fundamental rule of assertion is that one should assert p only if one knows p...more speculatively, we may project the account of assertion back on to its mental counterpart, judgment (or belief). What results is the rule that one should judge (or believe p) only if one knows p.[2]

Similarly, Jonathan Sutton argues at length that

> One must: believe p only if one knows p. [3]

Notice the colon after the "must". It's important; it specifies that the "must" has wide scope ranging over the entire rest of the sentence. Contrast this wide-scope view with

> One must believe p only if: one knows p.

This view gives the "must" narrow scope ranging only over "believe p". To start to see the differences between the two views, consider the moral principle "you should, if you kill an animal, kill it humanely". On at least one reading, this principle is uncontroversial. But we should be clear on the what the relevant reading is. Suppose that we read the principle as giving its "should" narrow scope. Then the principle amounts to:

Narrow Scope Animal Killing Theory

If you kill an animal, then you should kill it humanely

This principle is *very* implausible. Suppose that, in a fit of rage about my constant bragging about my pet dog Frank, you sneak into my house and kill him. It would then be true that you kill Frank. From this and the narrow scope animal killing theory (and some simple logic), we get the result that you should kill Frank humanely. But surely that is a mistake; surely you should not kill Frank at all, humanely or otherwise (see Figure 1).

Figure 1: Frank

So the narrow scope animal killing theory is far from uncontroversial. Yet there is a sense in which "you should, if you kill an animal, kill it humanely" *is* uncontroversial. Plausibly, that sense is what we get when the "should" has *wide* scope:

Wide Scope Animal Killing Theory

You should (kill an animal only if you kill it humanely)

Some theorists think this principle is meaningless as stated, on the grounds that "should" does not meaningfully take sentences as its objects.[4] If you are one of those theorists, then read the wide scope animal killing theory as follows:

> You should be such that the following is true of you: (you kill an animal only if you kill it humanely)

Either reading is fine for our current purposes. Maybe they amount to the same thing anyway. In any case, we can now ask: does the wide scope animal killing theory yield the implausible result that you should kill Frank humanely, in the case described above? No, it does not. To see why not, recall that $(p \rightarrow q)$ is equivalent to $(\neg p \vee q)$. One way to make it true that "if you kill Frank then you kill him humanely", then, is to make it false that you kill him. In such a case you might, consistently with the truth of the wide scope animal killing theory, also be conforming to another true principle, a true principle to the effect that "you should not kill Frank, humanely or otherwise". In other words, there are possible worlds in which it is true *both* that you should not kill Frank, not even humanely, *and* that you should (kill Frank only if you kill him humanely). In some of these worlds you might even kill him - humanely. But it would still be true that you shouldn't kill him, humanely or otherwise. By killing him, even humanely, you would do something you shouldn't do, something that violates the principle "you shouldn't kill Frank humanely or otherwise" even though it does not violate the (also true) principle "You should: kill Frank only if you kill him humanely". This means that, even along with the assumption that you do kill Frank, and even along with the assumption that you do kill Frank humanely, it does not follow from the wide scope animal killing theory that you should kill him humanely. So there is a separation between what is said by the wide scope animal killing theory, and what you should do with respect to killing animals.

As it happens, this separation is quite thorough. Effectively, the wide scope animal killing theory only tells us about the conditions under which you should do the following conjunctive act: killing-an-animal-and-not-killing-it-humanely. It does not tell us anything else. In particular, it does not tell us anything about the conditions under which you should kill an animal, or the conditions under which you should kill an animal humanely. If you don't kill one, then the theory tells us nothing about whether you should kill it,

or whether you should kill it humanely. If you do kill one, then once again the theory tells us nothing about whether you should kill it, or whether you should kill it humanely. It just tells us you shouldn't *kill it inhumanely*, which is something quite different from *killing it*, and also quite different from *killing it humanely*.

To generalize the point, it is this: wide scope theories, which take the form "you should be such that (P only if Q)", do not tell us anything about whether you should be such that P, or about whether you should be such that Q. They only tell us this: you should not be such that (P and not-Q). Putting this point to work, we can now see *why* it turns out that, even in conjunction with the assumption that you kill Frank in the case described above, it does not follow from the wide scope animal killing theory that in this case you should kill him humanely. The reason is this: the wide scope animal killing theory does not say *anything* about whether you should kill Frank humanely in *any* case. *Ipso facto*, it does not say that you should kill him humanely in the case described above.

With the differences between wide scope and narrow scope should-statements now broached, we can return to the question of what we should believe. Recall the theories on offer from Williamson and Sutton. The basic idea behind those theories is that we should believe *only* the things we know. If we should go along with this idea, then it is hard to see why we shouldn't also go along with the idea that we should believe *all* the things we know. Combining these two ideas, we arrive at the view that we should believe *all and only* what we know. It is important to distinguish between narrow scope and wide scope versions of this view. If the "should" has *narrow* scope, the idea amounts to

> **Narrow Scope Identity Theory**
> (You should believe p) iff (you know p)

To be completely explicit about this, the "iff" here is intended to express the material biconditional, and the whole view is intended to be implicitly preceded by universal quantifiers ranging over people, worlds, and times. Stated completely explicitly, then, the narrow scope identity theory amounts to:

6

For any person s, world w and time t: (s should believe p at t in w) iff (s knows p at t in w)

I will consider numerous views in this paper, and it would be cumbersome to state each of them so explicitly. Thus I will often leave implicit the quantifiers over people, worlds, and times, as well as leaving implicit the point that the biconditionals are material.

As for the narrow scope identity theory, it is clearly mistaken. Knowledge entails belief; hence it follows from the narrow scope identity theory that whenever we should believe we do believe. This is tantamount to the claim that we cannot fail to believe what we should, which is absurd. It will help to have a name for this particular absurd result. We'll say that a theory entails "compulsory completeness with respect to what we should believe", or in short, "compulsory completeness", when it entails that we always believe everything we should.

This terminology is borrowed in the obvious way from logic. Just as a logical system is complete with respect to a given property x iff every formula having x is a theorem, one's belief corpus is complete with respect to a given property x iff every proposition having x is something one believes. From an epistemic point of view, it is good for one's belief corpus to be complete with respect to what one should believe. Of course, it is possible for our belief corpuses to *not* be complete with respect to what we should believe. Theories denying this possibility have the problem of compulsory completeness.

The narrow scope identity theory is false because it has this very problem. Of course, there is an easy fix. We can just widen the scope of the "should" (as Sutton does by putting the colon after his "must", and as Williamson means to do with his own principle). This gives us

Wide Scope Identity Theory

You should (believe p iff you know p)

Or, to put it completely explicitly,

For any person s, world w and time t: s should not be such that he believes p at t in w without knowing it, or such that he knows p at t in w without believing it.

7

(The last clause is trivially satisfied, since knowledge entails belief.)

To see why this theory avoids the problem of compulsory completeness, suppose that you should believe p. From the narrow scope identity theory, it follows that you know p (and hence that you believe it). From the wide scope identity theory, it does *not* follow that you know p. For the wide scope identity theory doesn't tell us *anything* about the conditions under which you should believe p, just as the wide scope animal killing theory doesn't tell us anything about the conditions under which you should kill animals. Just as the wide scope animal killing theory speaks only of the conditions under which you should kill animals inhumanely (and says they never obtain), the wide scope identity theory speaks only of the conditions under which you should believe-and-not-know (and says they never obtain). Beyond *that*, the wide scope identity theory says nothing. *Ipso facto*, it says nothing about the conditions under which you should believe things.

From the point of view of ethicists of belief, who are in the business of trying to figure out what we should believe, this is not a helpful way out of the problem of compulsory completeness. What about cases where we don't both believe and fail to know? For that matter, what about cases where we *do* both believe and fail to know? The theory tells us that, in any given case, we should not (believe and not know). Even for the cases where we both believe and fail to know, this doesn't entail that we shouldn't believe - it only entails that we shouldn't *believe and not know*, which is something quite different. (Compare: even for the cases where you both kill Frank and kill him inhumanely, the wide scope animal killing theory doesn't entail that you shouldn't kill him - it only entails that you shouldn't *kill him inhumanely*, which is something quite different.)

In sum, there is a dilemma for theories identifying what we should believe with what we know. On the one hand, those theories can give their "should" narrow scope. So understood, the theories are well enough informative; but the information they give us is egregiously mistaken. On the other hand, they can give their "should" wide scope. Understood like *this*, the theories do not seem to make egregious mistakes. However, it is only by being objectionably uninformative about what we should believe, in fact *completely silent* about the matter, that they avoid such mistakes.

3   Counterfactual Theories

The wide scope identity theory doesn't tell us anything about what we should believe. Can its informational deficiencies be remedied? We might try to remedy them with counterfactuals, for instance by adopting

> **Narrow Scope Counterfactual Theory**
>
> (You should believe p) iff (if you were to believe p, you would know p)

But this theory has the standard problem for theories involving counterfactuals - the "conditional fallacy" problem.[5] To see this, suppose that you have more or less the same evidence all of us currently have about whether Barack Obama is the president of the United States, but due to a pure hatred of Obama you can't quite yet bring yourself to believe it - you think he must be ineligible due to cheating in the vote count, or perhaps to being a Kenyan, and so not really president at all. Of course, you are wrong about that. Obama really is president. Nor are there any Gettier hijinks on the scene; nor would you believe on an improper basis if you were to believe - you'd simply give in to the very same evidence you actually have, and get over your irrational hangups by believing, on the basis of your evidence, that Obama is president.

So far so good: the theory tells us, correctly, that you should believe Obama is president. But now let us add a twist. Suppose that you are being closely watched by a moderately malevolent demon with a taste for irony. If you come to believe that Obama is president, the demon will immediately kill Obama, instantaneously replacing him with a macrovisually indistinguishable Kenyan.

This twist does not make a difference to whether you should believe that Obama is president. But notice that *without* the twist it is true that you would know if you were to believe, whereas *with* the twist it is false that you would know if you were to believe. This means that if the narrow scope counterfactual theory is true, then the twist *does* make a difference to whether you should believe. And, again, the twist does not in fact make a different to whether you should believe. The theory is therefore mistaken.

Is there a way to reformulate the theory so as to avoid the problem? We

might try to do that by changing the relevant counterfactual, somehow building into it a stipulation that no such things as ironic demons are present in the worlds where the theory requires us to know things if in the actual world we should believe those things. Let us say that the a given batch of conditions are "propitious for your knowing p" just in case those conditions don't block your belief that p from being knowledge.[6] Putting this notion of propitious conditions to work, we can consider

> **Revised Narrow Scope Counterfactual Theory**
>
> (You should believe p) iff (if you were to believe p, and the conditions were propitious for your knowing p, you would know p)

This new view avoids its predecessor's problems with the ironic demon. However, the cure is worse than the disease. *Every* proposition is *always* such that, if you were to believe it and the conditions for your knowing were propitious, you would know it. The revised narrow scope counterfactual theory therefore entails that we should always believe everything.

Is there is a more charitable way to revise the narrow scope counterfactual theory? Perhaps we should take it that, in the worlds where the theory requires you to know p if in the actual world you should believe p, no *nonmental* conditions block your belief from being knowledge. Let us say that "the *nonmental* conditions are propitious for knowing p" just in case none of *those* conditions block your belief that p from being knowledge. Using this notion of nonmental conditions that are propitious for knowing p, we can rebuild the view as follows:

> **Re-revised Narrow Scope Counterfactual Theory**
>
> (You should believe p) iff (if you were to believe p, and the nonmental conditions were propitious for your knowing p, you would know p)

This view, unlike its predecessor, avoids the result that we should always believe everything. However, a new problem arises. Suppose that you don't believe some proposition p, but that if you were to believe it, you would believe it for bad reasons. To fill in the details of the case, suppose that you

have more or less all the same evidence most of us have about Obama, plus a small smattering of evidence from conspiracy-theoretic films about the matter. The large preponderance of your evidence indicates that Obama really is president, of course. But due to your hatred of Obama, you tend to pay much more attention to evidence gleaned from conspiracy theory literature, than to the rest of your evidence. You are on the fence between believing that Obama *is not* president (because he is a Kenyan - as many conspiracy theories say) and believing that he *is* president (because he is the son of two American communists - as is maintained by a flashy new documentary film you just watched).[7]

In such a case, the following counterfactual is *false*: if you were to believe that Obama is president, and the nonmental conditions for knowing were propitious, you would know that Obama is president. If you were to believe that Obama is president, and the nonmental conditions were propitious, you would *not* know that Obama is president. Instead of knowing as much, you would believe as much for a bad reason (namely, that Obama had as parents two American communists). As a result, the re-revised narrow scope counterfactual theory entails that you *should not* believe that Obama is president.

But in fact, you *should* believe that Obama is president. After all, you possess more or less all the same evidence as do the rest of us with respect to the issue of whether Obama is president, aside from a bit of conspiracy literature to which you are irrationally attached. The re-revised narrow scope counterfactual theory is therefore mistaken.

Is there another more charitable way to reinterpret the narrow scope counterfactual theory? No - at least, not as far as I can tell. But there is still one more way to try to save the counterfactual approach: we might try to get around the problems by widening the scope of the "should". This gives us

Wide Scope Counterfactual Theory

You should (believe p iff (if you were to believe p, you would know p))[8]

However, it should be clear that the move from narrow to wide scope here won't help, any more than the move from narrow to wide scope identity

11

theories will help. The narrow scope counterfactual theory does not tell us anything about the conditions under which you should believe things. It speaks only of the conditions under which you should be such that (you believe and the embedded counterfactual doesn't hold), or such that (the embedded counterfactual holds and you don't believe). To say something about only these conditions, is to say nothing at all about the conditions under which you should believe things.

4   Knowledge-Minus-Belief Theories

We've seen in two cases now that when narrow scope theories yield bad results, it doesn't help to try to avoid those results by widening the scope of their "shoulds". The pattern generalizes. In general, wide scope theories change the subject matter. For the remainder of our exploration of knowledge-first theories of what we should believe (and withhold on), then, we can leave wide-scope theories aside.

Focusing on narrow-scope theories, we've already seen that identity approaches and counterfactual approaches come up short. In the wake of these two approaches it is natural to try a different approach appealing to the *conditions on knowledge*. On this approach, we should believe a given proposition if and only if we meet all the conditions required for knowing that proposition, other than the belief condition. On this way of thinking about things, we build up the theory of what we should believe by taking knowledge and removing belief; whatever thing is left, that thing is necessary and sufficient for being such that we *should* believe.

The most obvious way of developing this approach takes it that you should believe when, for every condition knowledge requires other than the belief condition, you meet it. Using the box for metaphysical necessity, this amounts to

> **Non-belief Conditions Theory**
>
> (You should believe p) iff (for every condition x such that $\Box(Kp \to x)$ and $x \neq Bp$, x obtains)

But this won't quite do. Consider the conditions *true belief* and *knowledge*.

Each of these is necessary for knowledge, and not identical to belief. Thus, the non-belief conditions theory tells us that every case where you should believe is a case where these conditions hold. But these conditions both entail belief. So the theory tells us that every case where you should believe is a case where you do believe. Compulsory completeness returns.

Trying again, we might consider

Non-Belief-Requiring Conditions Theory

(You should believe p) iff (for every condition x such that $\Box(Kp \to x)$ and $\neg\Box(x \to Bp)$, x obtains)

According to this view, you should believe when, for every condition which is required by knowledge but which does not itself require belief, you meet it. By developing the "knowledge minus belief" approach in this way we nicely avoid the problem of compulsory completeness, because in building up the conditions under which you should believe, we factor out whatever requires you to believe. But there is a different problem. Consider the condition *having a propositional attitude whose content is p*. This condition is required by knowledge, and does not require belief. Thus it is a condition we must meet if we should believe a given proposition, according to the non-belief-requiring conditions theory. The theory tells us that we should believe only those propositions such that we (already) have propositional attitudes with those propositions as their contents. And this is a mistake. To see why it is a mistake, consider Watson, that well intentioned but dim investigator working for Sherlock Holmes. Watson may have never even considered the proposition that Smith is the murderer; but this may only be because he did not properly follow his evidence. We can even imagine Holmes explaining it all to him after the fact; to this explanation, Watson might respond "Now I see - I should have believed it was Smith". In this sort of case Watson speaks correctly; he should have believed it was Smith, even though he had no propositional attitude with that particular content. (If you don't find this case convincing as it stands, alter it so that Watson is not dim, but just intellectually lazy).[9]

In light of the problems with the non-belief-conditions and non-belief-requiring conditions theories, we might try to move from the notion of necessary conditions on knowledge, to notion of conditions *at least partly in virtue of which* we know things.[10] There are many necessary conditions on knowing

p, which are not also conditions at-least-partly in virtue of which we know p whenever we do know p. For instance, *that 1+1=2* must be true whenever anyone knows that he has hands; hence it is a necessary condition on knowing that one has hands. But it is not one of the conditions at-least-partly in virtue of which we know that we have hands whenever we do know as much. To be sure, it is a condition at-least-partly in virtue of which we know that 1+1=2, whenever we know *that* proposition. But still, it is not a condition at-least-partly in virtue of which we know that we have hands, whenever we know we have hands.

Let us say that a condition x is "knowledge-constitutive for p" if and only if: whenever anyone knows p, it is at-least-partly in virtue of x that he knows p. Putting the notion of knowledge-constitutive conditions to work, we can consider some new versions of the knowledge-minus-belief approach to what we should believe. For instance, we can consider a reformulated version of the non-belief conditions theory, to wit:

> **Non-belief knowledge-constitutive conditions theory**
>
> (You should believe p) iff (for every condition x such that x is knowledge-constitutive for p and $x \neq Bp$, x obtains)

This particular approach does not help, though, because it bring back compulsory completeness. Plausibly, whenever anyone knows any proposition p, it is partly in virtue of this condition - *properly forming his belief that p* - that he knows p. What is "properly forming" a belief? Perhaps, forming it in a reliable manner. Or perhaps, forming it on the basis of one's evidence, or on the basis of epistemic virtues, or in some other special way. In any case, it is plausible that *some* sort of proper formation is required, whatever the details of "proper formation" turn out to be. But this means that the non-belief knowledge-constitutive conditions theory entails that whenever we should believe p, we properly form the belief that p. And properly forming the belief that p entails having the belief that p. Thus the theory tells us that whenever we should believe p, we do believe p. We can't fail to believe what we should.

Trying again, we might reformulate the non-belief-requiring conditions theory via the notion of knowledge-constitutive conditions. This gives us

> **Non-belief-requiring knowledge-constitutive conditions**

theory

(You should believe p) iff (for every condition x such that x is knowledge-constitutive for p and $\neg\Box(x \to Bp)$, x obtains)

According to this theory, you should believe p when (and only when) you meet all the conditions which both (a) are conditions at-least-partly in virtue of which anyone who knows p, knows p, and (b) don't require believing p. This view nicely sails around the problem of compulsory completeness. For, much like its predecessor that involved mere necessary conditions instead of knowledge-constitutive conditions, this theory factors out whatever requires you to believe p. This theory also sails around the problem of conditions like *having a propositional attitude with p as its content*. Such conditions are plausibly *not* among the conditions in virtue of which anyone who knows p, knows p. Additionally, this theory sails around the problem of conditions like *forming the belief that p in the right way*. Although that condition is knowledge-constitutive for p, it requires believing p; and for that reason, the theory does not say it must obtain, in order for it to be true that we should believe p.

One might worry that similar conditions still bring up the same problems. For instance, one might worry that the condition *properly forming a propositional attitude with p as its content* is required for us to be such that we should believe p, given the theory we are now considering. However, that worry is off-target. Although *properly forming a propositional attitude with p as its content* is plausibly necessary for knowing p, it is *not* plausibly knowledge-constitutive for p. Although it is at-least-partly in virtue of *properly forming the belief that p* that we know p, it is *not* at-least-partly in virtue of *properly forming a propositional attitude with p as its content* that we know p.

The non-belief-requiring knowledge-constitutive conditions theory finds its way around all of the problems so far discussed for the preceding theories. It is probably the best attempt we've seen so far, among knowledge-first attempts to say what we should believe. But it has its own laundry list of difficulties. Most obviously, it tells us that in every Gettier case, and in every case where p is false, it is false that you should believe p. For instance, it is false that you should believe there is a barn in front of you, when you are looking at one which (unbeknownst to you) is surrounded by barn facades. Similarly, it is false that you should believe that the job-getter has ten coins

15

in his pocket, when the boss says that you will get the job, and you count ten coins in your own pocket, but (unbeknownst to you) someone else who also has ten coins in *his* pocket will get the job. *Whenever* something happens to be false - for instance, whenever it is false that you have hands because (unbeknownst to you) you are a handless brain in a vat being fed perceptions as of the actual world - it is also false that you should believe that thing.

These results seem clearly mistaken, even absurd.[11] Can they be explained away? Perhaps. Numerous theorists argue that one should *assert* p only if one knows p, or that one should *use p as a reason for acting* only if one knows p. These views tell us that we should never assert (or act on) anything false, or anything we are Gettiered about. Such results seem clearly mistaken, and even absurd, in the same ways in which the corresponding results about what we should believe seem mistaken and even absurd. However, there is a standard response to them. The response consists in distinguishing, on the one hand, what we should assert (or act on), and on the other hand, what is reasonable or excusable for us to assert (or act on). This distinction provides some sugar for the bitter pill of saying that Gettier victims and brains in vats regularly assert and act on what they shouldn't. Those characters turn out to be still making reasonable or excusable assertions, and engaging in reasonable or excusable actions, even though they are doing and asserting things they shouldn't. There is something positive to be said of those assertions and actions, then, even though they are not assertions people should make or actions people should do. When theorists find it clearly mistaken or absurd to say that we shouldn't assert (or act on) what we don't know, they are just failing to distinguish what is *reasonable or excusable* to assert or act on, from what we *should* assert or act on. The seemingly absurd results about assertion and action are thus explained away.[12]

Could a similar explanation be applied to the similarly absurd-seeming results about what we should believe? Could such an explanation adequately defend the non-belief-requiring knowledge-constitutive conditions theory from its problems involving falsehoods and Gettier cases?

Starting with the non-belief-requiring knowledge-constitutive conditions theory and then adding such an explanation, we end up with a package of views claiming that (a) we believe what we shouldn't in Gettier cases, and when we are brains in a vats, but that (b) these beliefs are nonetheless reasonable or excusable. Such package views - views packaging fairly demanding knowledge-theoretic claims about what we should believe with fairly unde-

manding claims about what is reasonable or excusable for us to believe - have quite a bit to recommend them.[13] In the end, they may turn out to work. However, if they do work, then they work by passing the theoretical buck from the notion of what we should believe to the notion of what is reasonable or excusable for us to believe. And if there is not something significant also said about the latter notion, then it is hard to shake the idea that the package views relocate our original question instead of answering it.

We start out by asking what we should believe. Package knowledge-first theories answer this question, but their answers lean hard on a new notion, the notion of what is reasonable or excusable for us to believe. In order to fully adjudicate these package knowledge-first theories, we need to ask a second question: what is reasonable or excusable for us to believe? Perhaps there is a good answer. But until we have one, it is hard to shake the idea that what mattered about our first question, or at least a significant part of what mattered about our first question, has been relocated (under the notion of what is reasonable or excusable to believe) instead of being resolved.[14]

Can the package views adequately fill in the details about what is reasonable or excusable for us to believe? Here is a reason to doubt as much. Knowledge-first theories of what we should believe, *if* they are combined with theories of what is reasonable or excusable for us to believe, ought to be combined with *knowledge-first* theories of the latter issue. And it is hard to see how there could be an adequate knowledge-first theory of the latter issue. Each so-far-discussed theory of what we *should* believe - the identity theories, the counterfactual theories, the knowledge-minus-belief theories - could be reinterpreted as a theory about what is reasonable or excusable for us to believe. But the reinterpreted theories seem just as problematic, and for just the same reasons, as do the originals. So, if we are going to find an adequate knowledge-first theory of what is reasonable or excusable for us to believe, then we will have to find it in some place other than reinterpreted versions of the views we've already discussed.

Where else is there to look? One *prima facie* promising strategy is to look at the additional theories (about what we should believe) that I will discuss in the rest of the current paper. However, the same point applies with those additional theories. As I will argue, each of those additional theories has some sort of problem. As for what those problems are, we'll get to them

soon enough. For now, the relevant point is this: each of those problems applies with equal force whether its target theory is about what we should believe, *or* about what is reasonable or excusable for us to believe. This means that we are not going to find - not in the sorts of ideas I explore here anyway - an adequately developed knowledge-first theory of what is reasonable or excusable to believe.

Perhaps we can find one somewhere else. But the prospects do not appear promising. As a result, what we might call "the package strategy" for knowledge-first views about what we should believe - that is, the strategy of saving those views from objections by packaging them with further views about what is reasonable or excusable to believe - does not appear to be a promising strategy. We should therefore stop the attempt to use that strategy to save the non-belief-requiring knowledge-constitutive conditions theory. And without *that* strategy, it is not clear how to save that theory. Nor is it clear how we might come up with another better knowledge-minus-belief theory. Those theories do not appear to work out. We need to look elsewhere.

5   Duplication Theories

Trying again to formulate a knowledge-first theory of what we should believe and what we should withhold belief on, we turn to the notion of *duplication*. A given (possible) person is your x-duplicate at a given time just in case they have all the same x-properties you have at that time. For instance, someone is your *physical duplicate* at t just in case they have all the same physical properties you have at t. Similarly, someone is your mental duplicate at t just in case they have all the same mental properties you have at t.

These notions can help us formulate several *prima facie* plausible knowledge-first approaches to the ethics of belief. For instance, consider

> **Mental Duplication Theory**
>
> (You should believe p) iff (p is known by at least one of your mental duplicates).

According to this view, you should believe a given proposition p at a given

time just in case it is metaphysically possible that there is some person, who has all the same mental properties as you at that time, and who knows p at that time. At first blush at least, this view comports with widespread intuition about what we should believe in numerous cases — for instance Gettier cases and fake barn cases.[15] But whatever its virtues, this view does not give us what we are after. For suppose that you do not believe p. Then, none of your mental duplicates know p, because none of them even believe it. Hence the view entails that whenever you do not believe p, it isn't the case that you should believe p. In other words, it entails that you cannot fail to believe what you should. This is the problem of compulsory completeness, once again leaving a *prima facie* promising knowledge-first theory in its wake.

We could try to save the mental duplication approach by focusing on *past* mental duplicates. Let us say that a (possible) person is your past mental duplicate just in case, at every time in your past, that person was your mental duplicate. Putting this notion of past mental duplication to work, we can consider

Past Mental Duplication Theory

(You should believe p) iff (p is known by at least one of your past mental duplicates).[16]

Interestingly, this theory avoids the problem of compulsory completeness. To see why, suppose you don't believe p. What follows from this supposition? It does *not* follow that none of your past mental duplicates believe p, or that none of them know p. They may believe or know p, even if you don't, owing to the ways in which they currently differ from you despite sharing your mental past. In such cases you might fail to believe p even though you should; thus, completeness is not compulsory. By focusing on the past instead of the present, the past mental duplication theory avoids compulsory completeness.

But there is a new problem. Suppose that S is your past mental duplicate right now, but that he differs mentally from you now in that he now sees that he lacks hands. He knows that he lacks hands; hence the narrow scope past mental duplication theory entails that *you*, right now, should believe that you lack hands. And of course that is a mistake. The general point of this example is that beings with the same mental histories should sometimes

believe different things owing to their current differences, even if one of them knows those things. The past mental duplication theory is inconsistent with this point. Its focus on the past instead of the present extricates us from the problem of compulsory completeness, but it does so by denying the fact that beings with the same mental histories should sometimes believe different things owing to their current differences, even if one of them knows those things.

Since the notions of mental duplication and past mental duplication both turn out to be unhelpful when taken on their own, we might try combining them. Let us say that someone is your *up-to-now mental duplicate* just in case they are your mental duplicate and were your mental duplicate at every time in the past. With this notion we can now formulate another theory, to wit:

Up-to-Now Mental Duplication Theory

(You should believe p) iff (p is known by at least one of your up-to-now mental duplicates).

But this brings back compulsory completeness. For suppose that you don't believe p. Then, none of your up-to-now mental duplicates believe p either. So none of them know p. So it's false that you should believe p. From the narrow scope up-to-now mental duplication theory and the assumption that you don't believe p, then, we get the result that it is false that you should believe p. The theory entails that you can't fail to believe what you should.

To summarize the discussion so far: several initially appealing knowledge-first approaches to the ethics of belief fall short. For various reasons, we end up falling short with identity approaches, counterfactual approaches, various knowledge-minus-belief approaches, and various approaches based on mental duplication. It might be tempting, given all this, to conclude that the knowledge-first ethics of belief amounts to a degenerating research program, and indeed that the ethics of belief shows to be degenerate the whole knowledge-first approach to epistemology. Such conclusions would be very satisfying to many traditional epistemologists, tired of being accused themselves by knowledge-firsters of being the degenerates on the scene. But I don't think these conclusions are warranted, at least not yet.

Let us help ourselves to the notion of "one's evidence", that is to say the

notion of the evidence you have - taking that notion as well enough understood to be usefully employable in the ethics of belief.[17] On the basis of this notion, we can start again the project of building a knowledge-first ethics of belief. Let us say that someone is your evidential duplicate just in case they have all the same evidence you have. Putting this notion to work, we can consider:

Evidential Duplication Theory

(You should believe p) iff (p is known by at least one of your evidential duplicates).

This theory won't quite do. Suppose you have no evidence whatsoever that Dean Martin is a good singer, even though you find yourself believing it. Suppose also that, although you've completely forgotten this, you formed that belief by trusting your mother's testimony, which you knew to be unreliable about such things.[18] Finally, suppose that one of your evidential duplicates formed the same belief on good grounds (which he has forgotten) having to do with listening to Dean Martin, comparing him to other singers, and so on, and without any Gettier hijinks happening. Given these suppositions, your evidential duplicate knows that Dean Martin is a good singer. Thus it follows from the evidential duplication theory that you should believe that Dean Martin is a good singer. But that is a mistake. Trying again, we can conjecture

Past Evidential Duplication Theory

(You should believe p) iff (p is known by at least one of your past evidential duplicates).

This theory, however, fails to properly deal with the relevance of your current evidence to what you should believe. Suppose that one of your past mental duplicates now has, for the very first time, an enormous amount of evidence that he lacks hands. Since you know you have hands, it follows - from the past evidential duplication theory - that he should believe that he has hands. This is a mistake. So neither the evidential duplication approach, nor the past evidential duplication approach, turn out to work.

Trying yet again, we might conjecture:

Up-to-now Evidential Duplication Theory

(You should believe p) iff (p is known by at least one of your up-to-now evidential duplicates).

According to this theory, you should believe exactly those propositions which are known by someone (that is, some possible person) who shares (and always has shared) your evidence. I think that this theory nicely avoids all the problems I've raised so far for the other knowledge-first approaches. For one thing, it avoids the problem of being too uninformative that the wide scope views face. Because its "should" takes narrow scope, it tells us whether we should believe any given proposition in any given case. And in combination with a plausible "uniqueness" principle - a principle to the effect that in any given case and for any given proposition p, either you should believe p, or you should believe not-p, or you should withhold on p - it also tells us whether we should withhold on any given proposition in any given case.[19] It therefore gives us a fully specified account of what we should believe and what we should withhold belief on. Furthermore, it does not have the problem of compulsory completeness, because even if you do not believe something, one of your up-to-now evidential duplicates might still know that thing.[20] In such cases it turns out that you should believe even though you do not believe. Nor does the theory face problems with slow-thinking Watsons or forgotten nefarious histories of belief formation. Our slow-thinking Watson should believe that Smith is the murderer, even though he has never even considered that proposition, because some of his up-to-now evidential duplicates (namely the faster-thinking ones) *know* that Smith is the murderer. And we get an correct account of, and an explanation of why, you should not believe that Dean Martin is a good singer, when you find yourself believing as much but, unbeknownst to you because you've forgotten, you originally formed that belief by trusting your known-to-be-unreliable mother about the matter. Here you should not believe that Dean Martin is a good singer, because none of your up-to-now evidential duplicates know as much.

The up-to-now evidential duplication theory, in combination with the uniqueness principle, amounts to our best option so far for a knowledge-first theory of what we should believe and what we should withhold belief on. But it has an obvious problem: it entails that we should never believe necessary falsehoods. Being necessarily false, such claims are necessarily unknown; in which case nobody's up-to-now evidential duplicates ever know them; in

which case the theory entails that nobody should ever believe them. And that is a mistake, because sometimes we *should* believe necessary falsehoods. For instance, if the entire mathematical community tells us, for several hundred years, that a certain very complicated mathematical claim is true, then we should believe that claim - even if that claim, for some subtle reason, turns out to be false (and hence necessarily false).[21]

Once again, we need to look elsewhere if we are going to find an adequate knowledge-first ethics of belief. Since the appeal to evidence didn't work out, it is natural to turn to the other notion most popular in standard contemporary theorizing about the ethics of belief, namely the notion of a belief-forming process.[22] Reliabilists theorize about the ethics of belief in terms of *reliable belief forming* processes. From a knowledge-first perspective, it makes sense to replace appeals to the notion of a belief-forming process with appeals to the notion of a *knowledge-producing process*. A theory built on this notion could say, perhaps, that we should believe whatever it is that knowledge-producing processes would have us believe. Since "trusting what mathematicians have been saying for hundreds of years" is a knowledge-producing process, this approach may be able to circumvent the problems with necessary falsehoods that refute even the best of our evidence-centered knowledge-first theories.

Of course this theory won't work as stated; it has numerous problems. For one thing, it has the standard "conditional fallacy" problems for theories involving counterfactuals. For another thing, it gives us very few details about the notion of a knowledge-producing process. It would be nice to get some details about (for instance) whether knowledge-producing processes generate knowledge *whenever* they are used, or some of the times they are used, or most of the times they are used, or most of the times they are used in the actual world, or what. And even if these first two problems can somehow be solved, a third problem remains. Consider the process "exercising an infallible ability to know things". Presumably, this is a knowledge-producing process on any adequate characterization of that notion. But then, we are threatened by the result that we should believe every truth, no matter how much (misleading) evidence we might possess against that truth. For if we should believe whatever it is that knowledge-producing processes would have us believe, and "exercising an infallible ability to know things" is a knowledge-producing process, then we should always believe every truth.

Perhaps we can avoid this third problem by drawing on the well-known relia-

23

bilist notion of processes that are "available" to a given believer - processes that are in some sense readily at the believer's disposal.[23]. Since normal human beings do not have processes like "exercising an infallible ability to know things" available to them, our process-focused approach is still a live option. Putting that approach into a precise theory now, and using the notion of an "available" knowledge-producing process, we might conjecture

> **Process Availability Duplication Theory**
>
> (You should believe p) iff (p is believed by at least one of your up-to-now knowledge-producing process availability duplicates).

This evades the problems involving counterfactuals, and the problems involving processes like "exercising an infallible ability to know things". But it still isn't a very good theory. It allows that two people might have dramatically different evidence - one might have all the evidence in the world that he has hands, and the other all the evidence in the world that he doesn't - and still they should believe all the same things if they are knowledge-producing process availability duplicates. This point motivates the return of the notion of *evidential* duplication. Adding that notion to our current approach, we might conjecture:

> **Up-to-now Evidential and Knowledge-Producing Process Availability Duplication Theory**
>
> (You should believe p) iff (p is believed by at least one possible person who is your knowledge-producing process availability duplicate and your up-to-now evidential duplicate).

We still aren't out of the woods, though, because one of these duplicates of yours might believe things for terrible reasons. For instance, one of these duplicates of yours might believe, for terrible reasons, that he lacks hands. But then, our theory entails that you, right now, should believe that you lack hands.

Maybe what is required for something to be what you should believe, then, is that it could be believed on the basis of a knowledge-producing process, by someone who shares and has always shared your evidence. Putting this idea into our standard format, we get:

Evidential Duplication and Knowledge Production Theory

(You should believe p) iff (p is believed, via a knowledge-producing process, by at least one possible person who is your up-to-now evidential duplicate and your process availability duplicate).

This is the last conjecture I'll explore. How good is it? It certainly lacks various details. For instance, it lacks various details about the notion of a knowledge-producing process. Even supposing that we have decided how often a process must produce knowledge to count as a "knowledge-producing" process, it still remains to be settled how those processes are individuated. That is to say, it still remains to be settled what the conditions are under which one use of a knowledge-producing process counts as a use of the same knowledge-producing process as does another use of a knowledge-producing process. Suppose that you and I both listen to a mathematician tell us something of the form "the mathematical community is absolutely certain that p", where "p" is replaced by one sentence for me and a different sentence (expressing a different proposition) for you. Suppose we both trust her, so that I come to form the belief expressed by the sentence replacing "p" in my case, and you come to form the belief expressed by the sentence replacing "p" in your case.

Did we use the same knowledge-producing process in coming to form our beliefs? If we didn't, or if it is in principle possible that we didn't given the setup of the case, then what makes for (or could in principle make for) the difference? Such questions surely deserve answers. They are the analogues of the questions standard reliabilists must answer in order to solve the "generality problem" - the problem of individuating belief-producing processes (as opposed to knowledge-producing processes). It is not clear how we ought to answer these questions.

Additionally, it is not clear whether our final conjecture can deal with all the problems of the other conjectures. In particular, it may not be able to deal with slow-thinking Watsons. When Watson does not realize that Smith is the murderer (or even consider that proposition), and this is only because Watson is too dim to appreciate the fact that his evidence inculpates Smith, does Watson have "available" to him any knowledge-producing processes which, compatibly with possessing the evidence he actually possesses, can lead to knowledge that Smith is the murderer? If so, then our final conjecture

deals with the Watson cases; if not, not.

So, does Watson have available to him the relevant sorts of knowledge-producing processes? I suppose that, in order to make our final conjecture work, we need "available" knowledge-producing processes to include (in Watson's case) something that would or could get him to know that Smith is the murderer, and to exclude (in the case of everyday people) processes like "exercising an infallible ability to know things". I am not sure how to define the notion of availability so as to yield these results.

6   Conclusions

We've found, in the evidential duplication and knowledge production theory, a conjecture that serves as a knowledge-first theory of what we should believe and (given the uniqueness principle) what we should withhold on, and that does not make any obvious mistakes. But it avoids those mistakes by leaving numerous important details for another day: details about how often a process must produce knowledge to count as "knowledge-producing", about how it is that knowledge-producing processes are individuated, and about exactly which knowledge-producing processes are "available" to us, in the sense of availability conjured up by the theory. Perhaps these details can be filled in; perhaps not.

Zooming out, and looking at the foregoing dialectic as a whole, what we see is that in a series of attempts to build up a knowledge-first ethics of belief, the detailed attempts have clear problems, and the attempts without clear problems lack details. This should not be at all surprising. It is the pattern we see over and over again in attempts to give necessary and sufficient conditions for things. It is the *standard* pattern.

Although the discovery of this pattern here is not surprising, it *is* theoretically interesting - both for the ethics of belief, and for contemporary epistemology more generally. It is theoretically interesting for the ethics of belief, because it shows where (or at least, some of the places where) the knowledge-first approach needs more detail in order to be viable. And it is theoretically interesting for contemporary epistemology more generally, because it brings a measure of disconfirmation to a certain idea that is popular in certain circles.[24] In particular, it brings a measure of disconfirmation to

the idea that by taking the notion of knowledge first, we can escape the frustrating patterns - that is, the patterns of informative theories facing counterexamples and counterexample-free theories that are uninformative - which we find in other more traditional approaches.

Notes

[2](Williamson, 2000, 11)

[3](Sutton, 2007, 44)

[4]See Schroeder (2004) for illuminating discussion of these issues

[5]See Shope (1978).

[6]The notion of conditions that are "propitious for knowing" is appropriated from (Bird, 2007, 85-86). Bird does not define that notion, but he seems to mean something like the above by it.

[7]See Gilbert (2012).

[8](Sutton, 2005, 373-374) seems to suggest this view; also see (Sutton, 2007, 56). For useful discussion see Littlejohn (2012) and Coffman (2010).

[9]Thanks to Esa Diaz-Leon for alerting me to this problem. It may also be a problem for a principle advocated by (Smithies, 2012, 268, 284) depending on how that principle is best interpreted. On Smithies' view, it is "correct" for you to believe p iff you meet all the "epistemic, as opposed to psychological" conditions on knowing p. He calls this principle "the K rule for correct belief". It is not fully clear whether what one should believe is identical to what it is "correct" for one to believe, in Smithies' sense of correctness. Nor is it fully clear whether, on Smithies' use of these terms, the "epistemic, as opposed to psychological" conditions on knowing amount to the conditions which are required by knowledge but which do not themselves require belief. If these identities both hold, though, then the problems for the non-belief-requiring conditions theory are also problems for Smithies' "K rule for correct belief".

[10]Thanks to Ryan Wasserman for making this insightful suggestion.

[11]Of course, these problems aren't unique to the non-belief-requiring knowledge-constitutive conditions theory. They apply to many other theories as well, for instance the narrow scope identity theory. For useful discussion see Littlejohn (2012).

[12]See Williamson (2000), DeRose (2002), Sutton (2007), Hawthorne and Stanley (2008).

[13]I once argued for versions of them concerning the satisfaction of curiosity (see Whitcomb (2010)), but I'm now unsure about those arguments.

[14]For a similar point about the similar package views about assertion and action, see Gerken (2011).

[15]Actually, on reflection it is not so clear that this theory adequately deals with all of these cases. Suppose that you, while looking at a particular barn which nobody else is looking at, believe that *that* thing is a barn. On some views of content (and modality and transworld identity), it turns out that the content of your belief is not identical to the content of any other possible person's belief. Any other possible person would be looking at a different barn (in a different possible world - which is why it would be a different barn), thinking of *it* that it is a barn. And that person's belief, being the belief that *that* thing is a barn, would have a different content than your belief, the belief that the thing *you* are looking at is a barn. Supposing that differences in content are sufficient for differences in belief, the mental duplication theory mistakenly entails that when you are in fake barn country, you should not hold the belief you would express by uttering "that is a barn in front of me". This is a problem for the mental duplication theory. Similar problems arise for theories that involve modal conditions on knowledge - safety conditions, for instance. Perhaps we should try to solve these problems by advocating some sort of internalistic theory of content, or by denying the sufficiency of differences in content for differences in belief, or by building up some sort of counterpart-theoretic treatment of belief identity, or by building up some non-world-theoretic treatment of modality. Or perhaps we shouldn't try to solve them at all, but instead take them to refute the theories they target. I will leave these issues - which arise for most any duplication theory - aside. For useful discussion see Manley (2007).

[16]Bird (2007) develops a somewhat similar approach to the notion of *doxastically justified belief* - as opposed to my own target notion of what we should believe (that is, the notion of what we have propositional justification to believe). For discussion of Bird's views, see McGlynn (2012). For another similar knowledge-first approach to doxastically justified belief, see Reynolds (forthcoming).

[17]Of course, there is a great deal of recent controversy over the notion of "one's evidence". See Feldman (1986), Williamson (2000), Kelly (2006), Goldman (2009), and Neta (2008).

[18]This is a well-known sort of scenario; for discussion see Goldman (2001), Conee and Feldman (2001), Greco (2005), and Feldman (2005).

[19]The uniqueness principle, though plausible, is not uncontroversial. See White (2005), Ballantyne and Coffman (2011), and Matheson (2011).

[20]This assumes, *pace* Williamson (2000), that the following conjunction is false: knowledge is identical to evidence and knowledge entails belief. Should that conjunction hold, the up-to-now evidential duplication theory would have the problem of compulsory completeness.

[21]Following Bird's (2007) approach to doxastic justification, we could invoke the notion of "corresponding propositions" to get around this problem of necessary falsehoods. The strategy here would be drop the idea that we should believe p iff p is known by some of our up-to-now evidential duplicates, and replace it with the idea that we should believe p iff *some proposition corresponding to p* is known by some of our up-to-now evidential duplicates. In cases where p is necessarily false, then, we should nonetheless sometimes believe p, because propositions *corresponding to p* are known by our up-to-now evidential duplicates. Perhaps this line of thought can be made to work, but it leans hard on the new notion of a "corresponding proposition". Until we can say something informative about that new notion, the approach seems underspecified.

[22]Here is one more attempt to make the appeal to evidence work out. Combining some standard evidentialist ideas from Conee and Feldman (1985) with some knowledge-first ideas from Williamson (2000), we might (a) identify one's evidence with one's knowledge, and (b) conjecture that one should believe (or withhold on) p iff believing (or withholding on) p *fits* one's evidence. However, there are problems with this view. Evidence may not be identical to knowledge. And even if it is, there are still other problems. Without an account of "fit", the view is very underspecified. And it is hard to see what an adequate account of "fit" would look like. The most obvious accounts of fit define that notion in terms of probability, so that believing p "fits" one evidence just in case p is rendered probable enough by one's evidence. But then how probable is "probable enough"? Presumably (but see Whitcomb (forthcoming) for some ways to relax this assumption), there is some value x such that for any proposition p, p is rendered "probable enough" by your evidence iff its probability given your evidence is at least x. What then is the value of x? If $x = 1$, the view becomes too skeptical, entailing (for instance) that we should believe relatively few things on the basis of testimony. For example, suppose that you are irrationally unwilling to accept the testimony of others, perhaps because your father constantly lied to you during your childhood. A wholly trustworthy person might tell you that the store has pumpkins on sale. You might decline to believe this person, coming to believe only that *he said* pumpkins are on sale, and not that pumpkins *are* on sale. If your evidence is identical to your knowledge, and you should believe only those propositions having probability 1 given your evidence, then you are making no mistake in this case, because in this case you *should not* believe that the store has pumpkins on sale. But in fact, in this case you *are* making a mistake; you *should* believe that the store has pumpkins on sale. So the view is too skeptical if $x = 1$. But if $x < 1$ the view is too credulous, entailing that we should believe that our lottery tickets are losers, even before the drawing takes place. Thus we are left with either an underspecified theory or a false theory - a theory underspecified due to the underspecified notion of fit, or a theory false due to being either too skeptical or too credulous. This is the standard theoretical pattern: we get either a counterexample-free theory that is uninformative, or an informative theory that has counterexamples. We could try to get out of the pattern by replacing the notion of fit with some sort of shiny new Bayesian machinery. But this strategy is unpromising: see (Williamson, 2000, 184-237) for some steps in its direction, and Whitcomb (2008) for a case that these steps can't get us out of the pattern.

[23]See (Goldman, 1979)

[24]See Williamson (2000), Sutton (2007), and Hossack (2007).

References

Ballantyne, Nathan and Coffman, E.J. 2011. "Uniqueness, Evidence, and Rationality." *Philosophers' Imprint* 11(18).

Bird, Alexander. 2007. "Justified Judging." *Philosophy and Phenomenological Research* 74(1):81–110.

Coffman, E.J. 2010. "Is Justified Belief Knowledge? Critical Notice of Jonathan Sutton, *Without Justification.*" *Philosophical Books* 51/1:1–21.

Conee, Earl and Feldman, Richard. 1985. "Evidentialism." *Philosophical Studies* 41:8:15–34.

—. 2001. "Internalism Defended." In Hilary Kornblith (ed.), *Epistemology: Internalism and Externalism.* MIT Press.

DeRose, Keith. 2002. "Assertion, Knowledge, and Context." *Philosophical Review* 111:167–203.

Feldman, Richard. 1986. "Having Evidence." In *Essays Presented to Edmund Gettier.* Kluwer.

—. 2005. "Justification is Internal." In Matthias Steup and Ernest Sosa (eds.), *Contemporary Debates in Epistemology.* Blackwell.

Gerken, Mikkel. 2011. "Warrant and Action." *Synthese* 178:529–547.

Gilbert, Joel. 2012. "Dreams of My Real Father: A Story of Reds and Deception." Documentary Film.

Goldman, Alvin. 1979. "What is Justified Belief?" In George Pappas (ed.), *Justification and Knowledge.* D. Reidel.

—. 2001. "Internalism Exposed." In Hilary Kornblith (ed.), *Epistemology: Internalism and Externalism.* MIT Press.

—. 2009. "Williamson on Knowledge and Evidence." In Patrick Greenough and Duncan Pritchard (eds.), *Williamson on Knowledge.* Oxford University Press.

Greco, John. 2005. "Justification Is Not Internal." In *Contemporary Debates in Epistemology*. Blackwell.

Hawthorne, John and Stanley, Jason. 2008. "Knowledge and Action." *Journal of Philosophy* 105(10):571–590.

Hossack, Keith. 2007. *The Metaphysics of Knowledge*. Oxford University Press.

Kelly, Tom. 2006. "Evidence." *Stanford Encyclopedia of Philosophy* .

Littlejohn, Clayton. 2012. *Justification and the Truth Connection*. Cambridge University Press.

Manley, David. 2007. "Safety, Content, Apriority, Self-Knowledge." *Journal of Philosophy* 104(8):403–423.

Matheson, Jonathan. 2011. "The Case for Rational Uniqueness." *Logos and Episteme* 2(3):359–373.

McGlynn, Aidan. 2012. "Justification as 'Would-Be' Knowledge." *Episteme* 9, 4:359–374.

Neta, Ram. 2008. "What Evidence Do You Have?" *British Journal for the Philosophy of Science* 59:89–119.

Reynolds, Steven. forthcoming. "Justification as the Appearance of Knowledge." *Philosophical Studies* .

Schroeder, Mark. 2004. "The Scope of Instrumental Reason." *Philosophical Perspectives* 18:337–364.

Shope, Robert. 1978. "The Conditional Fallacy in Contemporary Philosophy." *Journal of Philosophy* 75(8):397–413.

Smithies, Declan. 2012. "The Normative Role of Knowledge." *Noûs* 46(2):265–288.

Sutton, Jonathan. 2005. "Stick to What You Know." *Noûs* 39:3:359–396.

—. 2007. *Without Justification*. MIT Press.

Whitcomb, Dennis. 2008. "Williamson on Justification." *Philosophical Studies* 138/2:161–168.

—. 2010. "Curiosity Was Framed." *Philosophy and Phenomenological Research* 83/1:664–687.

—. forthcoming. "One Wage of Unknowability." *Synthese* .

White, Roger. 2005. "Epistemic Permisssiveness." *Philosophical Perspectives* 19:445–459.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.