

## Externalism for Doxastic Agents?

### Abstract:

*There is a tension between externalism and a central obligation of doxastic agency—the obligation to be involved in deliberately shaping one’s beliefs in light of one’s best judgments about what the reasons support. I consider the familiar case of Norman, the reliable clairvoyant, and argue that Norman’s case reveals a deliberative problem for externalism that arises from within a person’s first-person point of view. Externalism leaves open the possibility that sometimes the only way to aim at knowledge is to reject the aim of having reasonable belief, and a common distinction between various epistemological aims or concerns does not dissipate this tension. Despite the fact that the tension arises from within a person’s deliberations, it cannot be settled by deliberation. Rather, it is a higher-order worry about whether pursuing knowledge can direct a person qua doxastic agent to eschew her own doxastic agency. As a result, externalists need to recognize that the onus is on them to defend the significance of doxastic agency within an externalist framework.*

There is a tension between knowledge externalism and a central commitment of doxastic agency—the obligation to be involved in deliberately shaping one’s beliefs in light of one’s judgments about what the reasons support. I say “tension”, which is a weaker term than incompatibility, because I do not think that this tension counts against externalism as a theory of knowledge. But there is a point at which an externalist theory of knowledge and the commitments of doxastic agents uncomfortably press against one another. My aim is both to identify this tension and to argue that a common externalist strategy of dealing with first-person obligations does not satisfactorily address this worry.

The kind of externalism that concerns me here is an epistemological view that allows for the possibility that the requirements that suffice to make a true belief count as knowledge, whatever they may be, can obtain without the subject’s awareness that these features obtain or even the possibility that the subject might be able to become aware that these features obtain. Thus S can know that p even though S is unable to identify any considerations at all in favor of her belief that p.

She need not recognize or be able to recognize any reasons for believing  $p$ , and she need not even have the capacity to understand what it is to have a reason for one's beliefs. Neither reflection nor any sort of first-person perspective is required for knowledge. Internalists deny this possibility. According to internalism, knowledge requires the capacity to make some sort of positive assessment of one's epistemic status from one's own point of view. Depending on the version of internalism, it may require a great deal more.

There are, of course, weaker versions of externalism. One might, for example, think that *some* of the features that justify one's belief are external to one's own perspective without granting that a person can have knowledge while having *no view whatsoever* of her epistemic state. I am concerned here with the stronger sort of externalism for several reasons. First, I agree with Michael Bergmann that while there are many ways one might draw the internalism/externalism distinction, the key difference between internalists and externalists is that "internalists think there is an awareness requirement on *knowledge* and externalists think there isn't." (2006, 7).<sup>1</sup>

Rehearsing in detail the arguments for thinking that this is the proper place to locate the distinction is outside the scope of this project, but it might be helpful to mention two. First, externalists often take quite seriously the thought that animals and small children can and often do have knowledge, despite the fact that in many cases they do not even have the capacity to have a view about the epistemic status of their beliefs.<sup>2</sup> Locating the distinction here preserves this

---

<sup>1</sup> Those who are familiar with Bergmann's work might note that there is a difference in emphasis between my formulation of the identifying feature of the internalism/externalism debate and the one Bergmann offers. I have formulated the distinction between internalism and externalism according to whether or not a particular theory allows that it is possible for  $S$  to know that  $p$  even if whatever justifiers there might be for  $S$ 's belief are entirely unavailable to  $S$ 's own point of view. While I do not have an argument for it, as far as I can tell, my characterization is equivalent to Bergmann's insofar as it entails Bergmann's view and is entailed by it. If  $S$  knows that  $p$  while whatever justifiers there might be for  $S$ 's belief are entirely unavailable to  $S$ 's own point of view, then there is not an awareness requirement for knowledge. And if  $S$  cannot know that  $p$  while any justification for  $S$ 's belief is entirely unavailable to  $S$ 's own point of view, then there does seem to be an awareness requirement for knowledge. One reason I prefer my formulation of the difference is that it adds more precision to the notion of what "access" requires.

<sup>2</sup> For example, Fred Schmitt (2001) suggests that children can have justified perceptual beliefs even though they lack the concept of appearance. Without even the concept of appearance, it seems unlikely that a child could have a view of the reliability of her perceptual appearances.

possibility and I think accurately reflects a commitment most externalists readily affirm. Second, many of the traditional objections to externalism are aimed at refuting exactly this position. For instance, Laurence Bonjour (1980) and Keith Lehrer (1990) have offered counterexamples to externalism that focus on a subject's inability to recognize any support for his or her beliefs.

The notion of doxastic agency that I have in mind is not particularly robust. I propose the following minimal requirement:

Minimal Requirement for Doxastic Agency:

If X is a doxastic agent, X must have a capacity such that X's judgment *that X ought to believe that p* results in it being the case *that X believes that p*.

This is a familiar notion within a conception of agency that Nishi Shah identifies in which "the capacity to be moved by an appreciation of reasons is that in virtue of which agents exercise control over their beliefs" (2002, 442). This minimal requirement does not involve taking a stand on the degree to which one's influence over what one believes is in some sense *voluntary*. But it does locate the capacity for doxastic agency in a person's rationality insofar as one's judgments about what the reasons support play a constituting role in forming or sustaining one's belief. A doxastic agent is able to make judgments about what she ought to believe, and her beliefs can be formed in accordance with those judgments. A doxastic agent can also pursue knowledge by addressing to herself the deliberative question "*What am I to believe?*" and believing whatever the evidence supports. Stronger still, a doxastic agent has an *obligation* to be involved in deliberately shaping her beliefs in light of her judgments about what the reasons support. In this context, I will not address the source or the nature of this normative requirement. It is enough for my purpose here that some obligation of this sort exists, whatever its source. Philosophical and scientific inquiry certainly require attending to the reasons there are for adopting one view rather than another, and these pursuits assume that an agent is able to adjust her beliefs in accordance with her judgments about what is reasonable to believe.

I consider what I hope is not too tired a case—that of Norman, the reliable clairvoyant. I am interested in the Norman case not because it reveals any deep intuitions about the requirements for knowledge, for on my view it doesn't. Rather, I am interested in it because it illuminates a close connection between our assessment of a person's doxastic agency and our willingness to attribute knowledge and puts us in a position to see how the externalist faces a deliberative problem that arises from within a person's first-person point of view. The problem is that externalism leaves open the possibility that a person's pursuit of knowledge may direct that person to eschew her doxastic agency. It may turn out that sometimes the best way to aim at knowledge is to be unconcerned with having reasonable belief. Unlike the usual analysis, I contend that the interesting problem here is about the fit between externalism as a theory of knowledge, on the one hand, and our view of ourselves as doxastic agents, on the other.

#### I. Norman, Reliable Clairvoyant

Bonjour originally proposed the Norman case as a counterexample to the externalist claim that having a belief that is produced by a reliable process is sufficient for epistemic justification. By "epistemic justification" I mean here simply whatever sort of justification is necessary and, when combined with other conditions such as truth, is also sufficient for knowledge. Here is the case:

*Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable (62).*

Bonjour argues that Norman does not know that the President is in New York City even though Norman's belief resulted from the exercise of a completely reliable cognitive power—his reliable faculty of clairvoyance. This purported counterexample to externalism aims at exactly the sort of externalism that concerns us here, one in which Norman's knowledge does not require that he is even capable of articulating any reasons in support of his belief that the President is in New York.

Although Bonjour has since modified his own position, he originally claimed that the lesson is that externalist accounts of epistemic justification are mistaken. But it may be helpful to note that I (and other externalists) have found it difficult to see how Bonjour's counterexample amounts to a critique of externalism. Bonjour claims that the difficulty with externalism is that "a person may be ever so irrational and irresponsible in accepting a belief, when judged in light of his own subjective conception of the situation, and may still turn out to be epistemically justified" (1980, 59). Far from serving as a critique, it appears that this statement merely indicates that Bonjour has indeed understood an externalist theory of knowledge. In an externalist theory of knowledge, a person's own subjective conception of what she knows, and even whether she is able to have any such conception at all, is not required for knowledge. If Bonjour thinks that Norman does not know that the President is in New York, it can only be because he has already rejected the central thesis of externalism. So I think that an externalist need not be concerned that the Norman scenario poses a counterexample to externalism.

But my aim here is not to revisit Norman's case in the context of arguments about the nature of epistemic justification. I raise the issue in order to highlight that the tension I see between externalism and doxastic agency is *not* a matter of the correct analysis of knowledge. With that end in view, I wish to reconsider the case, this time asking whether or not there is something objectionable about Norman's situation, despite his epistemic success; that is to say, despite the fact that he does have knowledge. If the answer to that question is affirmative, then a closely related question has to do with how we ought to characterize what it is that is objectionable.

Suppose that we add to the mix a contrast between Norman's case and that of Irma, his reliably clairvoyant twin sister. Irma is in a similar circumstance as our original Norman with respect to her reliable clairvoyant faculty, but unlike Norman, Irma is aware of her faculty of clairvoyance and of its reliability in matters such as determining the President's location. One day, Irma comes to believe that the President is in New York through an exercise of her clairvoyant

faculty. When she considers why she believes that the President is in New York, she concludes (correctly) that she believes that the President is in New York as a result of an exercise of her power of clairvoyance. Both Norman and Irma's beliefs were produced by their reliable faculties of clairvoyance, and they both know that the President is in New York. But Irma is in a position to positively assess the epistemic status of her belief, while Norman is not in a position to do so.

What should an *externalist* say about these two cases? The externalist will certainly not follow Bonjour's lead and attribute the difference to the fact that Irma has knowledge while Norman does not. At the same time, it also seems that we should not say that there is nothing undesirable about Norman's position. The most common response here is to rely on a sort of division of epistemological labor. The basic thought here, and it is one with which I am generally sympathetic, is that we ought to keep the theory of knowledge distinct from what some call "first-personal" questions, such as whether or not we do, in fact, know anything at all, or whether a particular subject is being epistemically rational or has a justified belief. For instance, John Greco (2010) claims that recognizing a plurality of epistemic projects and epistemic goods can dissolve the conflict between internalism and externalism:

*If knowledge is understood as just one intellectual good among others, however, then the concept of knowledge need not do so much work. In particular, it now becomes possible to accommodate deep-seated intuitions motivating internalism, evidentialism and coherentism while rejecting them as intuitions about knowledge per se.*

If the question "What is knowledge?" belongs to the theory of knowledge and the question "What am I to believe?" belongs in a different inquiry, then the externalist is in a position to suggest that there are two, distinct evaluative criteria that apply to our assessment of Norman. One set of criteria concerns the requirements for knowledge. The other concerns some sort of requirements for reasonable or epistemically responsible belief.<sup>3</sup> Norman's failure is not epistemic (he does

---

<sup>3</sup> The notion of epistemic responsibility I employ here is an internalist one that places demands on an agent given how things look from her point of view. While it is certainly possible to have a thoroughgoing externalist account that extends to epistemic responsibility, to rely on an externalist notion of epistemic responsibility here will not do. The kind of epistemic responsibility I have in mind will be robust enough to

know), but he does display a failure to satisfy an obligation of doxastic agents, the obligation to be involved in deliberately shaping his beliefs in light of his best judgments about what the reasons support. For Norman is clearly *not* involved in shaping his beliefs in light of his best judgments about what the reasons support. Norman is not considering reasons at all. In fact, I think that this is precisely the sort of response an externalist should give. But it is also the case that prying apart knowledge and first-person concerns creates a problem.

## II. Norman, Doxastic Agent?

Before proceeding to consider this problem further, however, I wish to explore a bit further the source of the intuitions that suggest to some that Norman does not have knowledge. Although Bonjour claims that Norman fails to be epistemically responsible and attributes this to an intuition about what knowledge requires, there is a question about Norman's doxastic agency that is prior to the question of his epistemic responsibility. Norman fails to be epistemically responsible in in this case only if his doxastic agency extends to this particular belief.

The account of doxastic agency I have offer here takes it to be a minimal condition for agency that a person's judgments about what he ought to believe result in it forming his beliefs in accordance with those judgments. For most, if not all, ordinary (non-ideal) doxastic agents, the extent of the effectiveness of this capacity falls somewhere on a spectrum. At one end of the spectrum lies the *non-agent*. If X lacks this capacity, X is not a doxastic agent at all; X does not have the ability to form or maintain X's beliefs in this particular way. At the other end we find what we might identify as the *ideal* doxastic agent, a being whose judgments about what she ought to believe *always* result (or, perhaps, would result) in the formation or persistence of her beliefs in accordance with those judgments. Ordinary human doxastic agents fall somewhere between. We have the capacity for doxastic agency toward many of our beliefs, but, of course, not toward all of them. Most of us are all too well aware that sometimes our judgments about what we ought to

---

distinguish between Norman and Irma's situation and to maintain the commitment to shaping one's beliefs in light of one's best judgments about what the reasons support, and an externalist account cannot do either.

believe come apart from what we actually do believe.<sup>4</sup> Being a doxastic agent does not mean that one's capacity for doxastic agency is *always* effective, for a person may be a doxastic agent and yet still have particular beliefs fall outside the purview of his rational adjudication.<sup>5</sup> A belief that is insensitive to rational adjudication is, for that very reason, a belief over which a person's doxastic agency is not effective. There are, then at least these two ways of assessing agency. One is an assessment of whether one has the capacity for doxastic agency at all, and the other is whether one's agency extends to a particular instance, such as a particular belief.

When we consider Norman's situation, we are told that Norman believes the President is in New York despite the fact that he has no evidence for or against it. Since Norman has no reasons for believing what he does, Bonjour concludes that Norman is epistemically irresponsible, and as he understands it, being epistemically responsible involves evaluating things like whether or not a person has done her epistemic duty. In order to censure him for epistemic irresponsibility according to this internalist notion, we must assume that his doxastic agency is in effect.<sup>6</sup> But this is not something we can assess given the scenario as it stands. We don't know whether or not Norman has considered the matter and, crucially, we don't know what would happen if Norman were to consider the matter. Suppose Norman reflects on his belief and realizes that he doesn't have any evidence for believing that the President is in New York. It *may* be appropriate to consider Norman to be epistemically irrational, for it may be that when Norman realizes he doesn't have any reasons

---

<sup>4</sup> Sometimes our judgments about what we ought to believe do not result in the formation of beliefs that accord with those judgments. Cases of motivated irrationality provide many examples. Consider a scientist who finds himself in the uncomfortable position of being unable to give up his hypothesis despite the fact that it has been repeatedly falsified. The scientist may find himself an odd position of saying things such as: "The evidence shows that the hypothesis is false, but I cannot shake the belief that it is true."

<sup>5</sup> I say "at this time" because I want to leave open the possibility that circumstances may change (or may have changed) such that she would have the capacity on another occasion. Being a doxastic agent toward a particular belief likely depends on the presence or absence of other conditions. The love of one's son can be a deep motivation for believing one's son is innocent. Nevertheless, there may come a point in the future where these motivations are no longer sufficient to limit the influence of one's rational judgments.

<sup>6</sup> I take it that our capacity for doxastic agency legitimates the robust account of epistemic responsibility I have in view here. Without the ability to have one's judgments about what the reasons support result in the formation or perseverance of one's beliefs, it is hard to see how there could be any bite to the charge of failing to do one's epistemic duty.



for believing the President is in New York, his response is something like this: “Who cares? I don’t have any evidence, but I’m convinced the President is in New York.” If Norman’s response is of this kind, then we might be in a position to conclude that he is violating some demand of epistemic rationality after all.

There is also another possibility. Suppose that Norman realizes he doesn’t have reasons for believing the President is in New York, but this time he concludes as a result that he ought not to believe it. However, despite his own judgment that his belief is unreasonable, his belief persists. He can’t identify a shred of evidence to support his belief that the President is in New York, and yet he can’t shake his conviction that it is so. In this case, Norman’s belief is a recalcitrant datum about himself over which his best judgments about what is reasonable exert no influence. Were he to express his belief, he would feel compelled to qualify his own commitment, perhaps by saying “I cannot help but believe that the President is in New York” or “I cannot shake the conviction that the President is in New York.” What Norman would be indicating by these qualifications is the persistence of his belief despite his best judgment about what he has reasons to believe.

The suggestion I now want to put forward is that if Norman’s judgment about what he has, or doesn’t have, reason to believe fails to result in an adjustment in what he actually believes, then we should think of this particular belief as being outside the scope of his doxastic agency and, therefore, outside his epistemic responsibility. For when a person’s judgments about what he ought to believe are irrelevant to what he actually does believe, his relation to that particular belief is that of subject, not doxastic agent. If it turns out that Norman’s belief that the President is in New York is one over which he has no rational influence, the fact that Norman does not have reasons for what he believes becomes irrelevant to our assessment of his epistemic responsibility. Here I am addressing the possibility that the source of Norman’s epistemic failure is due to the fact that he has no reasons for believing what he does. But the point still stands even if we locate Norman’s supposed epistemic irresponsibility elsewhere. For instance, one might suggest instead that the

failure is due to the fact that his judgments about what is reasonable would not cohere with what he believes were he to consider the matter. We might say in this regard that his belief does not cohere with his all-things-considered judgment,<sup>7</sup> keeping in mind that such a judgment has not yet been made. But again, if his judgment about what it is reasonable to believe would have no effect on what he actually believes, then there is no cause to censure him for failing to make such a judgment, nor is there cause to censure him for failing even to try.

With this last point, I want to note that any general normative requirement there might be for having reasonable belief depends on the capacity to have one's judgments about reasons influence one's beliefs in the particular sort of way that constitutes doxastic agency. When we view someone as a doxastic agent, we usually assume that their capacity for doxastic agency is intact with regard to any given belief unless we are given a reason to think otherwise. Thus it may seem as though Norman is epistemically irresponsible in the context of his status as a doxastic agent. But this general expectation disappears when it becomes apparent that a particular belief is outside the effective range of his agency. It may be tempting to think that because a doxastic agent is subject to a requirement that he has reasons for his beliefs that this norm still applies in deviant cases where one's doxastic agency is not in effect. But this is precisely what I am arguing we must not do.

If we recognize that we have reached a limitation in Norman's doxastic agency, our expectation that he ought to have reasons for this belief disappears. It seems likely that the unease we feel about Norman's situation is not simply due to the fact that Norman is epistemically irresponsible, but that it may be that he is in a situation that threatens to undermine our conception of him as being a doxastic agent at all. If one does not concern oneself with what the reasons support or to makes an error in reasoning, one's failure is *as a doxastic agent*. But when one's own judgments about what the reasons support have no bearing on what one actually believes, within

---

<sup>7</sup> Davidson (1980) uses this phrase in regard to his proposed solution for the problem of *akrasia*. The notion he employs there is similar to what I have in mind here, although he uses it in a different context.

the context of that particular belief one is not a doxastic agent at all. And when a person's own judgments about what she ought to believe *regularly* fail to influence what she actually does believe, our conception of that person as a doxastic agent begins to erode.

I have been suggesting that there is an alternative to viewing Norman's predicament as a case of epistemic irrationality or irresponsibility. We ought to hesitate to agree with Bonjour that Norman is clearly epistemically irresponsible because it is possible that Norman's belief that the President is in New York persists regardless of Norman's own best judgment that there is no reason to believe it. Despite the fact that Norman has no reason for believing as he does, it may be that Norman is doing all that could be reasonably expected of him in this regard. In Norman's situation, and indeed any circumstance in which a person knows that *p* (and thus believes that *p*), but can find no reason for so believing, a great deal depends on the way in which a person's judgments are in a position to influence his belief. It might turn out that Norman is epistemically irresponsible, but it is also possible that he might come to reluctantly acknowledge his belief as a datum about himself with which he must come to terms. Norman's epistemic responsibility is in question here only insofar as his doxastic agency is not.

Considering the Norman case outside the context of our intuitions about epistemic justification helps us to recognize the relevance of certain background assumptions regarding Norman's doxastic agency. It also suggests that any censure that is appropriate for Norman should be understood as existing within the norms that apply to doxastic agents. Furthermore, while I cannot develop this thought in detail here, I suspect that if we do not view Norman as a doxastic agent, then we will also be a good deal less hesitant to ascribe knowledge to Norman. Most importantly, though, it directs our attention toward the relevance of Norman's doxastic agency in our evaluation of him. If Norman's doxastic agency does extend to this belief, his failure is *as a* doxastic agent. But if Norman's doxastic agency does not extend to this belief, then he has a failure *of* his doxastic agency. In this second case, the failure is not attributable to him as agent.

Nevertheless, Norman's circumstance is still problematic insofar as there is a limitation on the effective exercise of his capacity for doxastic agency.

### III. A resolution: different aims?

I wish now to consider in more detail the thought that there are distinct evaluative criteria that apply to our assessment of Norman: one concerning the requirements for knowledge, and another arising from his doxastic agency. The general idea is that any prima facie discomfort we might feel when we assess Norman's epistemic situation can be relieved by sorting out the various sorts of evaluative criteria that are in play. Externalism is a theory of knowledge that addresses questions about what knowledge is, but certainly does not entail that the knowledge-evaluation is the only evaluative standard that is of philosophical interest here.<sup>8</sup> Assessing Norman as a doxastic agent belongs to a different project that is concerned with something we might call epistemological guidance.<sup>9</sup> When a person asks a first-person deliberative question such as *What am I to believe?*, she is asking a deliberative question from the standpoint of her own perspective. There is different game afoot.

As I indicated, by and large I am in favor of acknowledging that various evaluative criteria are relevant to a case such as Norman's. The problem arises with a further step, although as far as I can tell few if any of the philosophers who make this move recognize it to be two steps rather than one. The thought goes like this: when we distinguish between various epistemological aims, we identify *distinct* evaluative criteria. Thus we can conduct our inquiry into what knowledge is without unnecessarily complicating things by also having to contend with first-personal concerns,

---

<sup>8</sup> I will not here take a position on whether or not the knowledge evaluation is the only *epistemic* evaluation that applies.

<sup>9</sup> Some think that an externalist can still address questions about epistemological guidance. For example, reliabilism tells us that our beliefs are epistemically justified just when they result from reliable processes. We can easily adapt this to be a recommendation for epistemic guidance: if p is the result of a reliable process of belief formation, then you ought to believe p. If p is not the result of a reliable process of belief formation, then you ought not believe p. Of course, this answer will not satisfy someone who is looking for an answer to the question "What am I believe?" that starts from a Cartesian position in which every belief is simultaneously suspect. But an externalist is not required to accept the premise that one has to be able to defend one's beliefs from a Cartesian position in order to have any means of epistemological guidance whatsoever.

and vice versa, we can proceed to address first-person deliberative questions without getting bogged down in quibbles about the nature of knowledge. Hilary Kornblith (2012, 2013), for example, argues at length that externalists need not concern themselves the first-person view.

*An externalist should not hold his beliefs hostage to the dictates of introspectively available evidence...It is important to recognize that externalists are not externalists about first-order belief only while holding that second-order belief must meet internalist standards. Externalism is a view about justified belief of whatever order (2012, 131-132).*

Picking up on Kornblith's hostage metaphor, we might conclude that eschewing first-person concerns amounts to something of a liberation for the externalist, a final deliverance from the internalist intuitions that have been so tenacious. This view is certainly understandable, especially when the focus is, as it has been for so long, on the nature of knowledge. But it is a threat to doxastic agency and to what I regard as a more fundamental commitment than the most deeply held externalist theory: the obligation to be involved in deliberately shaping one's beliefs in light of one's judgments about what the reasons support. In the remainder of my time, I want to try to bring out a problem: externalism leaves open the question of whether a doxastic agent ought to involve herself in pursuing knowledge at all. And I don't think that, despite Hilary Kornblith's protestations to the contrary, that is a question that can be left open.

It is widely acknowledged that doxastic agency plays a limiting role in our knowledge. This does not mean that the actions or perspective of an agent are necessary for knowledge. Rather, the claim is that a doxastic agent's self-assessments play an important role in forming her beliefs, and one can only know something if one believes it. One's activity as a doxastic agent matters to what one knows because the judgments of a doxastic agent matter to what she believes. When I, as a doxastic agent, judge that I ought to believe that p, if my doxastic agency is in effect with regard to this particular belief, I will, then, end up believing that p. Similarly, when I, as a doxastic agent, judge that I ought not believe that p, if my doxastic agency is in effect with regard to my belief, I will no longer believe that p. As a result, even if we regard the evaluation for knowledge as entirely

distinct from our evaluation *of* doxastic agents, our evaluations *as* doxastic agents will not be distinct from what it is we in fact can know. (This is so regardless of the quality of a person's judgments about what she ought to believe, for certainly a great deal can go awry. One might make mistakes about what follows from what, make inferences from false beliefs, or be subject to motivated irrationality. The account given in one's deliberation need not be true, although hopefully it is).

Once we turn our attention to how our evaluations *as* doxastic agents govern what we believe, we also cannot escape the realization that the exercise of our doxastic agency may not always work out so well. Consider that when Norman judges that he has no reason to believe that the President is in New York but still believes it anyway, this limitation of his deliberative agency actually enables him to have knowledge. For suppose that upon reflection Norman realizes he has no reasons for believing the President is in New York. When he recognizes that he has no reasons for believing the President is in New York, he no longer believes it. Once Norman no longer believes that the President is in New York (as a result of his own judgment that he has no reasons to believe so), of course he cannot know that the President is in New York. Here Norman's own involvement in managing his beliefs has resulted in him not knowing something that he would have known had he not considered the matter. If Norman is aiming for knowledge of the president's whereabouts, then on this occasion he can best satisfy that goal by *not* involving himself in considering reasons. Instead, he should follow his strongest impressions about the location of the President, despite the fact that to do so requires temporarily setting aside the commitment to be engaged in deliberately shaping his beliefs by judging what he has good reasons to believe. The effective exercise of Norman's deliberative agency would result in a loss of knowledge because his own judgment that he has no reasons for believing that the President is in New York will result in his no longer believing it. Here the persistence of his belief without regard to his own judgment is an epistemic

advantage. Sometimes one can best serve the aim for knowledge by intentionally not pursuing reasonable belief.

This is a first-person deliberative question that cannot be resolved by noting that there are distinct third-personal evaluations. Suppose I, as an externalist and a doxastic agent, want to do what I can to aid my procurement of knowledge and to decrease the errors I might make. Once I realize that what I believe, and therefore what I know, is influenced by my judgments about what there are good reasons to believe, I now face a question: To what degree ought I qua doxastic agent to be involved in evaluating reasons if my aim is knowledge? Ought I to proceed by pursuing careful, reasoned belief, or should I attempt insofar as it is possible to allow my beliefs to be formed, rejected, and sustained with as little of my own assessments or deliberation as possible?<sup>10</sup>

#### IV. The deliberative situation

It may be thought at this point that this is a deliberative question, but one that is not too pressing. After all, there is an empirical matter about the reliability of our own rational activity, and while it is certainly true that on occasion our own deliberations may prevent us from having knowledge, this hardly entails that these judgments are *generally* unreliable. After all, reliability does not require infallibility. As long as the circumstances in which our own assessments lead us away from knowledge are infrequent enough, then there is no cause for concern. We might even note a practical recommendation which goes something like this: One ought to be committed to deliberately shaping one's beliefs according to what one has good reasons for believing, for doing so is a generally reliable way to obtain knowledge. But, in circumstances where doing so would result in losing knowledge, the aim of knowledge outweighs the obligation to try to shape one's beliefs according to what the reasons support; that is to say, it cancels out the obligation to be a

---

<sup>10</sup> Within this context, I am concerned with a strictly epistemic answer to this question. This is not to ignore the possibility that a person could attend to pragmatic or aesthetic reasons for belief.

responsible doxastic agent.<sup>11</sup>Perhaps so, but note that this appeal to the general reliability of our rational agency doesn't help in the particular situation where I am asking myself whether or not, in this particular instance, my own judgments about what it is reasonable to believe interfere with my pursuit of knowledge. This practical recommendation fails to take into account the fact that this sort of judgment is inaccessible to a person's own, first-person point of view.

Here's why. A person can never be in circumstances where he would be able to tell that his own judgments about what is reasonable would interfere with his having knowledge because if a person were in a situation where he could tell *that*, then his decision not to consider reasons would itself be a reasonable decision. Consider again how Norman might proceed. Given what we know about Norman's clairvoyant capacity, we might from our third-person vantage point recommend that in this case if Norman is aiming for knowledge he ought to accept his inexplicable conviction that the President is in New York without concerning himself with what he does or does not have reason to believe. But Norman is in no position to make this judgment since he is entirely unaware of his reliable clairvoyant capacity. This assessment is only available to a third-personal view, to someone who is aware that Norman is clairvoyant. If Norman were aware of his clairvoyance, then he would be in a position to judge that his belief was justified after all.

We might think that what Norman needs here is the intervention by a third-party, someone who does have the third-person perspective available to make such an assessment and who is in a position to give a third-personal recommendation to Norman. This presumably reliable expert tells Norman that he should not concern himself with having reasons for believing that the President is in New York. Instead, the expert recommends that Norman should trust his unexplained but strong impressions about the location of the President. But if Norman takes the expert's advice, it is now no longer the case that Norman has no reasons at all for believing what he does. He now has

---

<sup>11</sup> Some might object to the thought that the value of being involved as an agent in shaping one's beliefs is derived entirely from its reliability as a means of procuring knowledge. Perhaps not, but I cannot see how other reasons for valuing doxastic agency would be epistemic. And if they are not epistemic, then it becomes puzzling why we shouldn't invoke criteria that aim for those values, rather than reasons that aim at the truth.



testimonial evidence that he ought to believe his strong impression about the location of the President, and therefore has some reason to think that these impressions are reliable. Furthermore, insofar as Norman trusts the expert's testimony, he is not abdicating his commitment to have reasons for his beliefs. Accepting expert testimony does not entail rejecting altogether the commitment to believe what one has good reasons to believe, even if the testimony we accept tells us that our own judgments about what is reasonable to believe are misleading. On the contrary, accepting this kind of testimony from others is a very important way that we check the reliability of our own judgments. Often enough, we are not in the best position to tell whether or not our own judgments are reliable. Editors point out bias to journalists, or a therapist may help a patient to realize he is engaging in self-deception. If one has good reasons to doubt the reliability of one's own judgments about a particular matter, depending on expert advice may be the most reasonable course of action to take. Yet in none of these cases would it be correct to say that a person is rejecting the commitment to deliberately shaping one's beliefs in light of one's judgments about what the reasons support. In fact, quite the opposite is true. The reasons may support not trusting our own judgments in a particular case or range of cases, but if we have a reason not to trust our own judgments, then we are not rejecting the obligations that arise from our doxastic agency. We are fulfilling them.

This is as we might expect when we consider that from within the first-person point of view, there is no way to pursue knowledge other than by aiming to believe what one has good reasons for believing. The concept of knowledge and the concept of reasonable belief are of course distinct. But identifying knowledge and reasonable belief as separate aims is a luxury of the third-person point of view, a view a person can recognize as an abstract possibility but cannot assess from the inside, so to speak. An analogy may be helpful here. Richard Moran (2003) points out that when a person avows a belief such as "I believe it's raining out" that person is able to recognize two commitments that he claims are both "unavoidable" and at the same time that "pull in different directions." One

commitment is to the state of the weather being a certain way. But Moran goes on, “At the same time, however, I must acknowledge myself as a finite empirical being, one fallible person in the world among others, and hence acknowledge that my believing something is hardly equivalent to its being true. And even when a person’s fallibility is not the issue, anyone must recognize that his believing *P* is nonetheless an additional fact, distinct from the fact of *P* itself” (74). Moran offered this remark in the context of describing the commitments that a person undertakes in believing, but the notion of dual commitments that, at the same time, pull in different directions, is also an apt way to describe a certain feature of the situation here.

It is straightforward enough to see how one could pursue reasonable belief and let the chips fall where they may with regard to knowledge, but it is not so easy to figure out how to aim for knowledge apart from attempting to believe what one has the best reasons to believe. In order for Norman to be in a position to make such a choice, he would need some way to tell when his involvement in managing his beliefs through reasoning about what he ought to believe helps him aim at knowledge and when it does not. But in order for Norman to be in a position to see that, he would need a reason to think that he is more likely to attain knowledge by not reflecting; that he is epistemically at his best in this circumstance if he in a sense *goes with* his conviction that the President is in New York, rather than concerning himself with having reasons for thinking so. And as we have already noted, if Norman has a reason to think that he is more likely to obtain knowledge by not considering the reasons for his belief, then his judgment that he does not need reasons is itself reasonable.

We can, of course, recognize the possibility that we may be in circumstances in which our judgments about what it is reasonable to believe would actually prevent us from having knowledge, just as we can recognize that the things we believe may not be true. However, we cannot be in a position to tell we are in those circumstances, any more than we can be in a position to tell that something we currently believe is false. (Odd cases of incoherence aside, once one realizes *p* is false,

then one no longer believes p.) So while the potential solution may be correct from the third-person point of view, it does not offer a resolution from the first-person point of view. From within the first-person point of view, one cannot abdicate one's doxastic agency on the basis of reasons without thereby casting a vote in favor of reasonable belief. What we have here is a deliberative problem that cannot be settled by rational deliberation, for to deliberate about the matter is already to attend to my own judgments about what the reasons support. The decision to eschew one's doxastic agency is not a decision one can make for reasons. The thought here is about how the aim of pursuing knowledge presents itself to a person; that is to say, how a person can proceed in pursuing knowledge. And my claim is that the only way a person can participate in pursuing *knowledge* qua doxastic agent is to aim for rational belief, and this holds true even if having reasonable beliefs ultimately fails to amount to knowledge. This doesn't amount to skepticism about rationality. In fact, I have a very strong commitment to the view that we should consider what the reasons support and believe accordingly. But I also think that commitment is stronger than any commitment to an externalist theory of knowledge.

#### V. Conclusion

To sum, there is something misleading in a pervasive externalist thought that there is no tension between an externalist theory of knowledge and the interests and obligations of doxastic agents because these things belong in different epistemological baskets, so to speak. In even the weakest internalist theories of knowledge, a subject must be able to make some kind of positive assessment of her belief from her own point of view in order to know. As a result, a person's own judgments about what she ought to believe are indispensable to knowing. But for an externalist, there is no such union between the aim of knowledge and a doxastic agent's aim of reasonable belief. Externalism allows for a possibility that internalist theories of knowledge do not—the possibility that the pursuit of knowledge would, at least in some cases, best be served by doing as little as possible to exercise rational influence over one's beliefs. If an externalist is to concern

herself at all with obtaining knowledge, even insofar as considering the reasons for being an externalist, she cannot escape the worry that her own involvement may actually keep her from knowing what she otherwise would know. This possibility places a special pressure on externalists to explain how the theory fits with the commitment to be involved in deliberately shaping one's beliefs in light of one's judgments about what the reasons support.

## References

- Bonjour, L. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy*, Vol. 5, No. 1, pp. 53-74.
- . 2002. "Internalism/Externalism." In *The Oxford Handbook of Epistemology* (2003), edited by Paul K. Moser. Oxford: Oxford University Press.
- Davidson, D. 1980. [Essay first published 1969]. "How is Weakness of the Will Possible?" *Essays on Actions and Events*. Oxford: Oxford University Press.
- Foley, Richard. 1993. "What Am I to Believe?" in *Naturalism: A Critical Appraisal*, Ed. Steven J. Wagner and Richard Warner. University of Notre Dame Press, Notre Dame, IN, p. 147.
- Greco, J. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. New York: Cambridge University Press.
- Kornblith, H. 2012. *On Reflection*. Oxford: Oxford University Press.
- . 2013. "Naturalism vs. The First-Person Perspective." *Proceedings of the American Philosophical Association*, Romanell Lecture, 87, 107-126.
- Lehrer, K. 1990. *Theory of Knowledge*, Boulder, CO: Westview.
- Moran, R. 2001. *Authority & Estrangement: An Essay on Self-Knowledge*. Princeton: Princeton University Press.
- Prichard, D. 2005. *Epistemic luck*. New York: Oxford University Press.
- Schmitt, F. 1993. "Epistemic Perspectivism," *Rationality, Morality, and Self-Interest: Essays Honoring Mark Carl Overvold*, ed. John Heil (Savage, Maryland: Rowman and Littlefield, 1993), pp 3-27. Reprinted in *Epistemology: Internalism and Externalism*, ed. Hilary Kornblith (Oxford: Basil Blackwell, 2001), pp. 180-206.
- Shah, N. 2002. "Clearing Space for Doxastic Voluntarism." *The Monist*, Vol. 85, No. 3, Controlling Belief, pp. 436-445.