The Philosophical Quarterly, Vol. 56, No. 225 *October 2006*
ISSN 0031–8094

Winner of *The Philosophical Quarterly* Essay Prize 2005

A NEW ARGUMENT FOR EVIDENTIALISM

By Nishi Shah

When we deliberate whether to believe some proposition, we feel immediately compelled to look for evidence of its truth. Philosophers have labelled this feature of doxastic deliberation 'transparency'. I argue that resolving the disagreement in the ethics of belief between evidentialists and pragmatists turns on the correct explanation of transparency. My hypothesis is that it reflects a conceptual truth about belief: a belief that p is correct if and only if p. This normative truth entails that only evidence can be a reason for belief. Although evidentialism does not follow directly from the mere psychological truth that we cannot believe for non-evidential reasons, it does follow directly from the normative conceptual truth about belief which explains why we cannot do so.

I. INTRODUCTION

When we deliberate whether to believe some proposition, e.g., whether to believe that it is snowing outside, we feel immediately compelled to look for evidence of its truth: we look outside. When our thinking is aimed at arriving at a belief, we treat the truth or falsity of the proposition as the only relevant question to be answered. Richard Moran puts the point in this way:

> The transparency of belief places certain constraints on the sort of active stance I may take towards my belief, and hence qualifies the 'practical' nature of practical questions. Transparency means that I must treat the practical question 'Shall I believe *p*?' as the impersonal theoretical question about *p*, and this means that the reasons I may have for adopting the belief are restricted to reasons connected with the truth of *p*.[1]

Following Moran, I shall call this feature *transparency*, although it may misleadingly suggest an even tighter relation between the two questions than Moran actually describes. To be clear, the feature that I call transparency is this: the deliberative question *whether to believe that p* inevitably gives way to the factual question *whether p*, because the answer to the latter question will

[1] R. Moran, 'Making Up Your Mind: Self-Interpretation and Self-Constitution', *Ratio*, NS 1 (1988), pp. 135–51, at p. 146.

Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK, and 350 Main Street, Malden, MA 02148, USA

determine the answer to the former. In the sense I have in mind, deliberating whether to believe that *p* entails intending to arrive at belief as to whether *p*. If my answering a question is going to count as deliberating whether to believe that *p*, then I must intend to arrive at belief as to whether *p* just by answering that question. I can arrive at the belief just by answering the question whether *p*; however, I cannot arrive at the belief just by answering the question whether it is in my interest to hold it.

In the last sentence quoted above, Moran moves from the fact of transparency to the conclusion that the only *reason* I can have for the belief that *p* is evidence of *p*'s truth. I shall call this position, that only evidence can be a reason for belief, *evidentialism*. Richard Foley claims that whether a consideration is a reason for belief depends solely upon its relation to the furtherance of an agent's ends. According to this view, if believing that his wife is faithful would better serve the cuckolded husband's ends than believing that she is having an affair, then it would be irrational for him to believe that she is having an affair, even if this is the proposition best supported by the evidence. Foley is a *pragmatist* about reasons for belief: he is committed to the existence of at least some non-evidential reasons for belief. Yet Foley, like Moran, seems to accept that only evidential considerations move us in deliberating what to believe. In fact, he cites this phenomenon to explain why we do not normally evaluate the rationality of our beliefs in terms of how well they promote our non-intellectual goals:

> Offering you a million dollars to believe that the earth is flat may convince you that you have a good economic reason to believe the proposition, but in itself it won't be enough to persuade you that the earth is really flat.
>
> By contrast, becoming convinced that you have good intellectual reasons to believe something – in particular, good evidential reasons – ordinarily is enough to generate belief. A belief is a psychological state that by its very nature, in Bernard Williams' phrase, 'aims at truth'.

A paragraph later, he writes

> Similarly, in your own deliberations about what to believe, you ordinarily don't consider what practical reasons you might have to believe something, and part of the explanation is similar to the third-person case. Deliberations concerning your practical reasons are customarily inefficacious and hence pointless.[2]

In the latter quotation, Foley does use the terms 'ordinarily' and 'customarily' to qualify his acceptance of transparency. However, in the former, he claims to accept Williams' thesis that it is in the nature of belief to aim at truth, which according to Williams expresses the fact that it is *impossible* to acquire a belief that *p* by deliberating on considerations which have no

[2] R. Foley, *Working Without a Net* (Oxford UP, 1993), p. 16.

bearing on whether *p* is true.[3] Although Williams is correct that I cannot deliberate from non-evidential considerations so as to arrive at a belief, non-evidential considerations can influence my beliefs by way of non-deliberative processes such as wishful thinking. In order to square Foley's agreement with Williams with what he says in the last quotation, I shall assume that the latter, rather than expressing a qualification of his acceptance of transparency, is intended to leave open the possibility that practical considerations may have this kind of non-deliberative influence on belief.

It looks as though Moran thinks that transparency has an obvious normative implication, while Foley does not think so. Moran seems to think that the fact of transparency alone implies that only evidence can be a reason for belief. After all, he does not cite anything other than transparency in support of evidentialism. But Foley seems to think that it is just as obvious that transparency has no normative implication: it is merely a psychological barrier to an agent's ability to believe for the right reasons when his ends would best be served by a belief which is not supported by his evidence. Attempting to persuade someone to change an opinion by citing non-evidential considerations is an ineffective strategy, but according to Foley the inefficacy of such considerations does not undermine their reason-giving force.

I think that the resolution of the disagreement between evidentialists like Moran and pragmatists like Foley turns on the correct explanation of transparency. I aim to show that acknowledging the psychological phenomenon of transparency puts the pragmatist in an unstable position, and that the best explanation of this phenomenon supports evidentialism. I shall proceed as follows: I shall first describe an argument for a deliberative constraint on reasons which is a generalized version of an argument Bernard Williams gives.[4] When combined with transparency, this deliberative constraint on reasons implies evidentialism. This means that if the deliberative constraint on reasons is correct, then the inference which Moran makes from transparency to evidentialism is sound. I shall then argue that the pragmatist who acknowledges transparency is faced with a dilemma. He must either accept or reject the deliberative constraint on reasons. If he accepts it, he must give up his pragmatism. But if he rejects it, he must deny transparency, because the best explanation of transparency entails the deliberative constraint on reasons. Acceptance of what is required to explain transparency thus entails commitment to evidentialism. Having made my

[3] B. Williams, 'Deciding to Believe', in his *Problems of the Self* (Cambridge UP, 1973), pp. 136–51, at p. 148.

[4] Williams, 'Internal and External Reasons', in his *Moral Luck* (Cambridge UP, 1981), pp. 101–13.

case for evidentialism, I shall finish by explaining how to analyse the most famous purported counter-example for evidentialism, Pascal's wager.

My defence of evidentialism should put it on a more secure basis than that provided by its moralistic advocate W.K. Clifford, who famously argued that it is a sin against mankind to believe anything upon insufficient evidence.[5] Clifford's argument was effectively rebutted by pragmatists such as William James, who pointed out cases in which believing in accordance with one's passions when one has insufficient evidence would have either morally neutral or even morally salutary results.[6] The argument I shall give for evidentialism, however, does not rely on dubious claims about the moral consequences of belief, but follows from the best explanation of transparency, a widely acknowledged psychological phenomenon.

For my purposes, it is important to start with a description of transparency that is neutral between the pragmatist and evidentialist. While I shall argue later that close examination of transparency shows it to be the reflection of a *conceptual* truth about belief, I shall not assume from the outset that both sides accept this claim. But I shall assume that both pragmatists and evidentialists agree that transparency is an *unalterable psychological fact* about our doxastic deliberation, where 'our' includes at a minimum everyone reading this paper. While there may be those who deny even this weak form of transparency, I shall be satisfied if I can show that if transparency is an unalterable fact about our doxastic deliberation, then evidentialism is correct.

II. REASONS AND REASONING

In this section I shall give an argument supporting a deliberative constraint on reasons. In this, I take myself to be following a path laid down by Bernard Williams in his essay 'Internal and External Reasons', although I do not have space to give textual support for this claim here. The argument is as follows:

1. R is a reason for X to ϕ (where ϕ ranges over actions and beliefs) only if R is capable of being a reason *for which* $X\,\phi$s
2. R is a reason for which $X\,\phi$s only if it is possible for X to treat R as counting in favour of ϕing in X's deliberation whether to ϕ.

Treating a consideration as counting in favour of ϕing in deliberation about

[5] W.K. Clifford, 'The Ethics of Belief', in his *Lectures and Essays*, Vol. II (London: Macmillan, 1877), pp. 339–63.

[6] W. James, 'The Will to Believe', in his *'The Will to Believe' and Other Essays in Popular Philosophy* (New York: Dover, 1956), pp. 23–8.

whether to ϕ involves being disposed towards ϕing by one's recognition of the (apparent) normative force of that consideration. Therefore premise (2) can be restated thus:

2*. *R* is a reason for which *X* ϕs only if *R* is capable of disposing *X* towards ϕing through *R*'s role in *X*'s deliberation whether to ϕ.

Premise (2*) does not require that an agent's reason must actually enter into his reasoning about whether to ϕ; it just requires that the reason must be eligible for entrance into his reasoning. It is important to allow for this possibility, because it might plausibly be denied that in order for a consideration to be a reason for which an agent ϕs, it must actually enter his reasoning about whether to ϕ. The point expressed by (2*) is that in order for a consideration to be an agent's reason for ϕing, it must be possible for the agent to be moved towards ϕing by recognizing that the consideration normatively favours ϕing. Also, when I say that *R* must be capable of disposing *X* towards ϕing through its role in *X*'s deliberation about whether to ϕ, I mean to be pointing to the characteristic motivational force that a consideration possesses when it functions as a premise in deliberation. I am not currently in a position to undertake the philosophical task of describing this characteristic motivational force in such a way as to render the description immune to counter-examples. As the seemingly endless cycle of counter-examples and repairs contained in the literature on deviant causal chains attests, this is an unenviable job. I shall merely note that it does seem to be uncontroversial that considerations can move an agent to believe or act by way of their functioning as premises in reasoning, and that there are both characteristic and deviant ways in which this functioning can eventuate in belief or action. So, putting these thoughts together, one can arrive at the conclusion which I shall call 'the deliberative constraint on reasons':

3. *R* is a reason for *X* to ϕ only if *R* is capable of disposing *X* to ϕ in the way characteristic of *R*'s functioning as a premise in deliberation whether to ϕ.

The sense in which *R* must be capable of disposing *X* to ϕ is that there are no unalterable features of *X*'s psychology that prevent *R* from disposing *X* to ϕ in the way characteristic of a consideration's functioning as a premise in deliberation about whether to ϕ.

The line of thought expressed by the argument is this: claiming that a consideration is a reason for an agent to ϕ implies that it is capable of being a reason for which the agent ϕs. Claiming that a consideration is a reason for which an agent ϕs in turn implies that the consideration guided the agent in its capacity as a reason. A consideration could not guide an agent to ϕ in

its capacity as a reason unless the agent were capable of ϕing on the basis of his recognition of the consideration as a reason to ϕ. Deliberation, or reasoning, is the process in which agents recognize reasons, and then ϕ on the basis of this recognition. So something could not be a reason for an agent to ϕ unless it was capable of swaying him towards ϕing in his deliberation about whether to ϕ.

This argument is a way of spelling out the familiar thought that it is the function of reasons to guide agents. Agents are not mere bystanders detached from the causes of their beliefs and actions, but have the power to believe and act on the basis of their reflective appreciation of reasons. It appears to me that the stick is bent, but on reflection I realize that it is in water; distrusting my eyes, I come to believe that the stick is straight. I desire to eat the cake, but I realize that it will thwart my goal of losing weight, so I abstain. These garden-variety cases illustrate the fact that by reflecting on reasons, agents can sometimes directly change the causal flow from perception to belief and desire to action. But even when agents 'go with the flow', reflection on reasons can contribute to what they believe and do. I might come to believe that the stick really is bent because upon reflection I trust my eyes; or I might decide to eat the cake because upon reflection I judge the pleasure of eating the cake to outweigh the health benefit of losing weight. Guiding agents in this way is crucial to a consideration's being a *reason for which* an agent acts or believes. Deliberation, or reasoning, is connected to the nature of reasons by being that through which agents are guided by reasons. This function of guiding agents distinguishes reason-judgements from other types of endorsement. Telling someone that he has a certain reason to pay back a loan is advice. Its status as advice would be undermined if it turns out that he is incapable of heeding it, that is, if he is incapable of being motivated to pay back the loan for that reason. Telling someone that it would be wonderful if he paid back the loan need not be a piece of advice, however; it might be true that it would be wonderful were he to pay back the loan even if in fact he is incapable of paying it back. Telling someone it would be wonderful if he jumped over the moon makes sense; telling someone he has a reason to jump over the moon does not.

Here is the argument when 'believe' is substituted for the variable 'ϕ':

B1. *R* is a reason for *X* to believe that *p* only if *R* is capable of being a reason for which *X* believes that *p*

B2. *R* is a reason for which *X* believes that *p* only if *R* is capable of disposing *X* towards believing that *p* in the way characteristic of *R*'s functioning as a premise in doxastic deliberation

B3. Therefore R is a reason for X to believe that p only if R is capable of disposing X towards believing that p in the way characteristic of R's functioning as a premise in doxastic deliberation.

The sense in which R must be capable of disposing X to believe that p is that there must be no unalterable features of X's psychology that prevent R from disposing X to believe that p in the way characteristic of R's functioning as a premise in doxastic deliberation.

So why cannot the pragmatist accept the deliberative constraint on reasons? The conception of reasons to which pragmatists all either explicitly or implicitly subscribe is that R must make ϕing in some way *attractive* to X if R is to be a reason for X to ϕ.[7] And it is in the light of an agent's desires that options appear attractive or unattractive. In the end, an agent's reasons are thus determined by his set of desires. If we focus exclusively on reasons for action, this pragmatist thesis about reasons appears consistent with the deliberative constraint on reasons. That an action would be either intrinsically or instrumentally attractive in some way certainly seems to be the kind of consideration capable of moving an agent in his practical deliberation about what course of action to undertake. However, although the attractiveness of an option can engage one's practical deliberation as a consideration in favour of pursuing it, the attractiveness of believing that p cannot similarly engage one's doxastic deliberation as a consideration in favour of believing that p. This is because the attractiveness of a belief does not tell for or against the truth of p, and the question of p's truth occupies the sole focus of our attention in doxastic deliberation. When we ask ourselves the deliberative question whether to believe that p, this question gives way to the question *whether p is true*, and so the only way for us to answer the former question is by answering the latter. This is the phenomenon of transparency with which I began the paper.

Transparency, when combined with the deliberative constraint on reasons, thus rules out pragmatic considerations from being our reasons for belief. Crucial to this deliberative constraint on reasons is the claim that if there are unalterable features of X's psychology that prevent a consideration from entering X's deliberation about whether to ϕ, the consideration is not a reason for X to ϕ. Both pragmatists such as Foley and evidentialists such as Moran seem to agree that transparency reflects a psychologically unalterable fact about us. Therefore pragmatists such as Foley must deny the deliberative constraint on reasons. And it seems that Foley would deny it.

[7] See Foley, *The Theory of Epistemic Rationality* (Harvard UP, 1987), ch. 5, and *Working Without a Net*, ch. 1; P. Railton, 'Moral Realism', *Philosophical Review*, 95 (1986), pp. 163–207, at pp. 166–7. Both writers explicitly endorse this instrumental conception of reasons.

After all, he cites something like transparency as a mere psychological barrier that interferes with our believing for the right reasons. Rather than seeing transparency as telling us something about the nature of reasons for belief, Foley seems to think that it is an obstacle to be overcome or accommodated in attaining rational belief. Thus he is bound to reject a principle which allows one to draw a normative conclusion from transparency about reasons for belief. In the spirit of Hume, Foley might say that the deliberative constraint on reasons illegitimately permits one to draw a normative conclusion from a non-normative premise.

III. THE PRAGMATIST'S DILEMMA

I shall now argue that the pragmatist cannot both accept transparency and deny the deliberative constraint on reasons. If this is right, pragmatists such as Foley who acknowledge the phenomenon of transparency face a dilemma. If they accept the deliberative constraint on reasons, then they must give up their pragmatism. But, as I shall argue, the best explanation of transparency entails the deliberative constraint on reasons. Thus if pragmatists reject the deliberative constraint on reasons, then they must also deny transparency. I shall first briefly summarize this explanation of transparency, which I have given elsewhere, and then argue that pragmatist-friendly attempts to explain transparency fail.[8]

An account of transparency must explain why it is that when we ask ourselves whether to believe that *p*, this question is answered by, and only by, answering the question whether *p* is true. It is not as though when we ask ourselves whether to imagine or suppose that *p*, we answer those questions by figuring out whether *p* is the case. Why does the question whether *p* is true *hegemonically* impose itself on our doxastic deliberations but not on these other kinds of deliberations? This is the question that an account of transparency must answer.

My hypothesis is that the concept of belief includes a standard of correctness.[9] When one deliberates whether to have an attitude conceived as a belief that *p*, one deliberates about an attitude to which one already applies the standard of being correct if and only if *p* is true, and so one is already committed to considering it with an eye exclusively to whether *p* is true. When one deliberates whether to have an attitude conceived as an

[8] See my 'How Truth Governs Belief', *Philosophical Review*, 112 (2003), pp. 447–82, and N. Shah and J.D. Velleman, 'Doxastic Deliberation', *Philosophical Review*, 114 (2005), pp. 497–534, for a fuller explanation.

[9] Also see P. Boghossian, 'The Normativity of Content', *Philosophical Issues*, 13 (2003), pp. 31–45.

assumption or fantasy, one does not yet apply any particular standard to it, and so one does not yet have any commitment as to how one will go about considering it.

My explanation of transparency conceives of deliberation as a norm-governed activity. Deliberation is reasoning aimed at issuing in some result in accordance with norms for results of that kind. Deliberating about whether to ϕ is reasoning aimed at issuing or not issuing in ϕing, in accordance with norms for ϕing. The deliberative constraint on reasons naturally flows from this conception of deliberation. Deliberating whether to ϕ is not merely one amongst other causal routes to securing results of kind ϕ; it is the most explicit route by which we are guided by the application of norms in reaching results of kind ϕ. Reasons for ϕing are considerations which indicate whether ϕing would be correct according to the norms for ϕing. Deliberation whether to ϕ thus is precisely the place to look if we want to understand what can be a reason for ϕing, and not just a mere cause of ϕing. We uncover what sorts of considerations can be reasons for ϕing by seeing which norms are applied in deliberating whether to ϕ. If there is a sole norm governing all activities of type ϕ, this norm will determine what sorts of considerations can be reasons for ϕing. The deliberative constraint on reasons says that only considerations which can be deliberated from to conclude in ϕing are eligible to be reasons for ϕing. According to my explanation of transparency, this is because only such considerations can determine whether ϕing would be correct according to the norms governing ϕing.

To deliberate whether to believe that *p* is to engage in reasoning aimed at issuing or not issuing in a belief that *p* in accordance with norms for such a belief. The sole norm for belief is this: believing that *p* is correct if and only if *p* is true. Therefore considering whether *p* is true constitutes reasoning aimed at issuing or not issuing in one's believing that *p* in accordance with the norm for believing. Because the norm is contained in the concept of belief, and doxastic deliberation is framed by the question of whether to believe that *p*, anyone engaging in doxastic deliberation accepts that this is the relevant norm.

This does not mean that deliberation about what to believe cannot be influenced by non-evidential considerations; it just means that such deliberation cannot explicitly treat such considerations as relevant to the question what to believe. Any influence that such considerations exert must be unacknowledged. But this is just as it should be. Transparency is a conscious phenomenon: we cannot consciously acknowledge considerations which are irrelevant to the truth of *p* as determining whether to believe that *p*. However, as we know very well, what cannot be consciously acknowledged often has a powerful influence none the less.

My explanation of transparency also leaves room for the possibility that beliefs can be influenced by non-evidential considerations in non-deliberative contexts, because it entails only that one is forced to apply the standard of correctness in situations in which one exercises the concept of belief. Not all belief-forming processes involve the deployment of the concept of belief on the part of the believer, so not all beliefs are causally regulated by the believer's application of the norm of truth. In such cases belief may be influenced by non-evidential factors.

It is a conceptual fact about belief that truth is its sole norm. It is a conceptual fact about deliberation that it is an activity governed by norms. Deliberation about whether to believe that *p* thus must be governed by considerations relevant to determining whether *p* is true, on pain of not counting as doxastic deliberation at all. The claim that truth is normative for belief is not a descriptive claim about how beliefs are causally regulated, nor does it imply any such claim. My explanation of transparency thus allows for the fact that passions can influence belief, at least in non-deliberative contexts.

This explanation of transparency is unavailable to the pragmatist, since it entails the deliberative constraint on reasons, and, as I argued earlier, the pragmatist who acknowledges transparency must deny the deliberative constraint on reasons. The pragmatist thus must either deny transparency or give an alternative (and better) explanation of transparency which does not entail the deliberative constraint on reasons.

As far as I can see, there are two ways available to the pragmatist of trying to account for transparency. First, he might claim that we all have a desire to have true beliefs, or better, that we all have a desire to have true beliefs about those matters of which we desire to have beliefs at all. This latter formulation avoids the objection that there are many insignificant matters about which we do not trouble to form opinions at all, much less true opinions. We may not take care to maximize the number of true beliefs we have, but we do want the beliefs we have to be true. This hypothesis would attempt to explain transparency as follows: we ask ourselves whether to believe some proposition *p* because we have an interest in forming an opinion about *p*, and in answering the question whether to believe that *p*, we focus on the question whether *p* is true because we have an interest in having a true opinion about *p*.

While this explanation might account for why truth is relevant to doxastic deliberation, it cannot account for the fact that truth is hegemonic with respect to doxastic deliberation. Our desire for truth is only one of our desires. In circumstances such as the cuckolded husband's, in which believing what is true is less desirable than believing what is false, according

to the pragmatist the question whether *p* is true ought to give way to other questions, such as whether believing that *p* is likely to save one's marriage. But determining whether *p* is true is not just relevant to determining whether to believe that *p*: it is the *only* question relevant to determining this. While it might be plausible to suggest that we all have an interest in having true beliefs concerning those matters about which we care to form opinions, it is not plausible to claim that everyone has an overriding interest in having beliefs that are true. Some might disagree with the cuckolded husband who wants to believe that his wife is faithful whether or not she is, but nobody finds his desire unintelligible. Whether it is more desirable to have a true belief than a false belief which is comforting is something that people disagree about, and certainly most of us would agree that there are circumstances in which believing what is false is more desirable than believing what is true. But transparency holds of everyone's doxastic deliberation, not just those who have an overriding desire for truth. Therefore transparency cannot be accounted for by the assumption that everyone has a desire for having true beliefs (concerning those matters about which they desire to form an opinion). Depending upon how we interpret the importance accorded to this desire, the claim is either too weak to support the exclusive focus on truth in doxastic deliberation, or too strong to be true.

The pragmatist might claim instead that as a matter of psychological fact human beliefs are determined solely by evidence. Given this fact about human psychology, if I fail to focus solely on evidence for and against *p*, my deliberation will not deliver a belief about *p*. So if I want my deliberation to conclude in belief, which is the point of doxastic deliberation, I must take this fact into account.

I find this account of transparency phenomenologically off-key. It is not as though, in deliberating about whether to believe that *p*, the reason why one focuses on whether *p* is the case is that one has noticed that as a matter of psychological fact one has come to believe only what one has ascertained to be the case. This would involve an inferential step: 'Should I believe that *p*? Well, I shall end up believing that *p* if and only if I ascertain that *p* is true, so I had better consider whether *p* is true.' But there is no such inferential step involved in moving from the question whether to believe that *p* to the question whether *p* is true. When I ask myself whether to believe that it is raining, the question whether it is raining becomes immediately and solely relevant. I recognize immediately that the only way to answer the former question is to answer the latter.

I suppose that a pragmatist might argue that the immediacy is the result of the obviousness of the fact that our beliefs are solely caused by evidence, so that the inference need not be explicitly made, and the move from the

one question to the other will be psychologically seamless, and will appear non-inferential. Even if this were the case, there is a much worse problem with the pragmatist's explanation of transparency: it relies on a false premise. In order to explain why the question whether *p* is true is solely relevant to answering the question whether to believe that *p*, the pragmatist has to claim that human beliefs are as a matter of fact solely caused by evidence. This flies in the face of the platitude that evidentially insensitive processes such as wishful thinking sometimes influence beliefs. In our practice, we routinely criticize one another for having beliefs which are the result of prejudice or preference, and a philosophical account of belief ought at least to allow that such criticisms are intelligible. But if it were a fact that our beliefs are only affected by evidence, then such criticisms would never make sense. The lesson is that a correct account of transparency must focus on features that are unique to doxastic deliberation, as my account does, not on features possessed by all belief-forming processes. Evidence exerts hegemony over doxastic deliberation, but not over all belief-forming processes.

Might the pragmatist retreat to the claim that while non-evidential considerations can influence belief outside deliberation, it is just a brute fact that only evidence is efficacious in doxastic deliberation? To say this would not be to explain transparency, but to claim that there is no explanation to be had. Although it is true that the pragmatist cannot give a satisfactory explanation of transparency, there is, as I have shown, a satisfactory explanation of transparency available to the evidentialist.

What is the scope of my defence of evidentialism? Is evidentialism true only of creatures like us, or does it apply to all possible rational agents? If transparency were an accident of human psychology, then it would be possible for other creatures to have pragmatic reasons for belief; it would just be an unfortunate fact about us that we are confined to evidential reasons. Transparency would limit us to evidential reasons for belief in the same way as our physical abilities limit our reasons for action. It is an unsurprising fact that people who have different abilities have different reasons for action. Even though I would very much like to be an NBA player, my utter lack of any of the physical abilities required to play in the NBA makes it the case that I have no reason to try out for an NBA team; but those fortunate few with such desires and the requisite physical abilities may have reasons to do just that. Similarly, if there are creatures that have the ability to hear sounds that we cannot, they may have reason to create kinds of music that we humans have no reason to produce.

If transparency were merely a feature of human psychology, then it would be of a piece with the kind of physical limitations that constrain our reasons for action. It would be surprising if the central debate in the ethics of

belief could be resolved by such a mundane consideration. Although I started with the assumption that transparency is only a fact about *our* doxastic deliberations, it turns out that the best explanation of transparency shows that it is a fact about doxastic deliberation as such. Transparency is best understood as the first-personal deliberative recognition of a conceptual truth about belief. But if it is a conceptual truth about belief, then anyone who asks himself whether to believe that *p* must recognize it. What he recognizes is that his belief about *p* will be correct if and only if it is true, and this sets the standard for what can count as a reason for or against believing that *p*. The most plausible explanation of transparency thus implies that evidentialism is not merely the correct view about human reasons for belief, but is the correct view about reasons for belief as such.

Last, this explanation of transparency assuages the worries of those who are suspicious of the Williams-style argument for evidentialism because it appears to derive a normative conclusion from non-normative premises. According to my account of transparency, the evidential nature of doxastic reasons is due not to the fact that we are psychologically incapable of believing for non-evidential reasons (although this is true), but to the fact that belief is subject to the standard of correctness of truth.

To sum up, I first argued for the deliberative constraint on reasons. This deliberative constraint, when combined with transparency, entails evidentialism. From this I concluded that the pragmatist who accepts transparency must deny the deliberative constraint. I then argued that the best explanation of transparency entails the deliberative constraint. The pragmatist who acknowledges transparency thus has no hope of arriving at a non-deliberative conception of reasons. Any conception of reasons which fails to imply the deliberative constraint would be inconsistent with a full acknowledgement of transparency. Therefore although we cannot infer evidentialism from the fact of transparency alone, as Moran seems to do, I have argued that the best explanation of transparency in fact does entail evidentialism.

IV. PASCAL'S WAGER

How should evidentialism deal with the counter-examples that philosophers have generated against this view? Although I have not based my case for evidentialism on its ability to deal with such examples, I shall finish by explaining how I think the evidentialist should deal with the most famous alleged counter-example, Pascal's wager. I hope this will clarify the strategy for undermining other purported examples of practical reasons for belief.

People often claim that religious belief is based upon faith, not evidence. The evidentialist claims that there are people who come to believe that God exists via deliberation, and that there are people who come to believe this via non-evidential processes such as faith, but he denies that non-evidential considerations can be reasons for believing that God exists, because he denies that such considerations can function as premises in reasoning which concludes in one's believing that God exists.

Pascal's wager is intended to show that even if the probability of God's existence is very low, it is in one's interest to believe that God exists. This is because, Pascal claims, if God exists then God will punish those who do not believe with an eternity of pain and horror and reward those who do believe with the infinite bliss of heaven. Therefore, to put it in modern decision-theoretic terms, as long as there is a non-zero probability that God exists, the expected utility of believing that God exists is higher than not believing that God exists. Pascal claims that if you accept his argument but are still unable to believe, 'it is because of your passions, since reason impels you to believe and yet you cannot do so'.[10]

Certainly there have been people, some of my students, for example, who claim that Pascal's argument provides a good basis for believing that God exists. But have any of them come to believe that God exists directly on the basis of Pascal's practical argument? If not, is Pascal correct in saying that this is because the influence of *passion* blocks the operation of *reason*?

I do not think so. Once we have a correct description of the ways in which people can actually use Pascal's wager in coming to believe that God exists, it is possible to see that in fact the wager is not the reason for which anyone believes that God exists.

Bob is convinced by Pascal's argument. Does this mean that he is persuaded that God exists? No, only that he is convinced that it is in his interest to believe that God exists. After being convinced by the wager, Bob might judge that it would be *desirable* for him to believe that God exists, but he will not infer from this that God exists. Given transparency, the only way he can conclude his doxastic deliberation in the belief that God exists is if his deliberation puts him in a position to judge that God exists. It is thus impossible for Bob to come to believe that God exists directly on the basis of appreciating the practical argument expressed by the wager, since to do so would violate the very constraint that is constitutive of our doxastic deliberation.

Even Pascal displays recognition of this as an obstacle to be overcome, when he suggests (p. 152) to someone who claims to accept the wager, but is unable to believe that God exists, that the other should self-induce the belief by these alternative means:

[10] B. Pascal, *Pensées*, tr. A.J. Krailsheimer (New York: Penguin Press, 1966), p. 124.

> Learn from those who were once bound like you and who now wager all they have. These are people who know the road you wish to follow, who have been cured of the affliction of which you wish to be cured: follow the way by which they began. They behaved just as if they did believe, taking holy water, having masses said, and so on. That will make you believe quite naturally, and will make you more docile.

Pascal is suggesting that if one acts as if one believes that God exists, eventually this will bring about the belief. There are at least three ways in which this might happen. The simplest would be straightforward conditioning. By habitually acting as though he believes, Bob may eventually find himself with the actual belief in question. This process of conditioning might be successful in a community that rewards behaviour that conforms to belief in the articles of faith and punishes behaviour that does not. So putting himself in a religious community and mimicking the behaviour that is rewarded might eventually cause Bob to believe that God exists. On this model of conversion, Bob comes to believe that God exists because he is rewarded for behaviour of a type that would be produced by those who believe, and so it is appropriate to say that practical considerations are directly responsible for his conversion. However, in this case the practical considerations are not working through Bob's deliberation. Bob does not, nor could he, employ the fact that he is being rewarded for such behaviour in deliberation that concludes in his believing that God exists, since he knows that such rewards are no evidence that God exists. There might be a mechanism in Bob that is sensitive to the benefits of belief and is able to cause beliefs that would be rewarded (i.e., his unconscious is sensitive to practical considerations in this way), but Bob is not moved by taking these considerations as reasons in his own deliberation about whether to believe that God exists, although he may very well take them as reasons in his deliberation about whether it would be *desirable* to believe that God exists.

The second way in which placing oneself in a religious community and acting as if one is a theist might bring about religious belief is by putting oneself in a position to gain direct knowledge that God exists. Just as one has to fill one's head with theory and look through a microscope in order to gain knowledge about amoebae, one might need to be in the right environment with the right frame of mind before God is willing to reveal himself. This idea is expounded by William James (p. 325):

> Just as a man who in a company of gentlemen made no advances, asked a warrant for every concession, and believed no one's word without proof, would cut himself off by such churlishness from all the social reward that a more trusting spirit would earn – so here, one who should shut himself up in snarling logicality and try to make the gods extort his recognition willy nilly, or not get it at all, might cut himself off forever from his only opportunity of making the gods' acquaintance.

James is suggesting that only by making oneself trusting, or as Pascal puts it, 'docile', will one make the acquaintance of God. And if being acquainted with something is a way of knowing it, as it presumably is, then knowledge of God can be secured only by practical means.

There are at least two interpretations of what is meant by 'acquaintance' here. On one interpretation, putting oneself in a particular frame of mind will lead God to grant one *perceptual* experience of him. But if this is really a case of perception, then one's perception (or the content of the perception), rather than any consideration having to do with the benefits of believing in God, could enter as a reason in one's doxastic reasoning. Were Bob to have such a perceptual experience, he would be in a position to reason his way to believing that God exists.

If acquaintance were not a way of perceiving God, then being acquainted with God would involve nothing more than causation of belief (of course, by God). On this interpretation, God might *cause* Bob to believe that he exists if Bob makes the good-faith effort of putting himself in a religious community and takes part in all the activities of the faithful. But if this is the idea, then it is not a case of coming to believe that God exists on the basis of reasons at all, much less of coming to believe on the basis of non-evidential reasons. Bob will have been caused to believe that God exists in such a way that his belief is true and caused by the thing that makes it true, but nevertheless his belief will not be due to the cause's impinging on him in the form of a reason. So even if this were to count as a case of knowledge, it would not count as a case of belief on the basis of reasons.

This case is essentially the same as a science-fiction case in which a pill induces Bob to believe that God exists. Pascal's wager persuades Bob that it would be in his interest to believe that God exists; he also believes that the pill will induce the belief, and by a means–end piece of reasoning he concludes that he ought to take the pill, which either constitutes or causes an intention to take the pill. However, this would not be a case in which Bob reasons his way to believing that God exists, since the intention to take the pill, not the belief that God exists, would be the conclusion of his reasoning.

Similarly, in James' case, Bob can conclude deliberation with the intention of making himself docile so that he is receptive to making God's acquaintance (if God exists). The fact that the pill induces belief more rapidly than the routes that James and Pascal suggest does not make an essential difference to the structure of these deliberations. In both cases, the endpoints of deliberation are not beliefs, but intentions or actions that may induce belief, either right away or eventually. They thus do not refute the claim that practical considerations cannot move an agent within the context of doxastic deliberation, that is, deliberation which concludes in belief.

The third way in which performing the actions that Pascal prescribes might create belief is by surreptitiously influencing Bob's doxastic deliberations. The person to whom Pascal is prescribing these actions does not think that there is very good evidence that God exists, so let us assume that Bob is in this position. By partaking in the activities of the religious, he might hope to come to a different estimation of the weight of the evidence. After having lived amongst the religious, if things go well, Bob may come to think that what he once thought weak evidence, or what he had never thought of as evidence at all, might now appear to be good evidence for God's existence. If this method were ultimately successful, his perception of the evidential force of these considerations would lead him to judge that it is true that God exists and thus to believe that God exists.

But might not someone mistakenly take a practical consideration to be an evidential consideration, with the result that it is a justifying premise in the agent's doxastic deliberation? For example, let us imagine that Bob accepts the general principle that if something is good for him then it is probably true (i.e., maybe he thinks that there is a law of nature that ensures this connection). This would not falsify the claim that doxastic deliberation only weighs evidential considerations. Bob, in virtue of accepting this principle, takes facts about his good to provide him with reasons to believe. He is therefore guilty not of mistaking a practical consideration for an evidential one, but of accepting an unwarranted evidential principle. We, taking an external third-person perspective, might say that he has based his belief on considerations which are not in fact evidence, and thus that his belief is not based on evidential reasons. But we are not in a position to say that Bob is basing his belief upon a practical consideration, merely that he is mistaken about whether the desirability of a belief indicates that the content of the belief is true.

V. CONCLUSION

The relevance of transparency to the ethics of belief has been obscured by failure to distinguish the deliberative question of whether to believe that *p* from questions that arise within perspectives external to doxastic deliberation, such as whether believing that *p* would be prudent or virtuous or desirable in some other way. Failure to distinguish these questions is a symptom of the dominance of the model of practical deliberation in philosophical thinking about reasons. Transparency does not structure practical deliberation about what to *do*, about what action to perform, even when that question is whether to do something that will bring about one's

believing that *p*. If this were the question at issue for the ethics of belief, then the pragmatist would win hands down. I know of no compelling argument that truth is always more desirable than other values such as happiness or virtue, or even that attaining true belief never conflicts with attaining other values. But the question *what belief to bring about* is distinct from the question *what to believe*. Answering the former question issues in an action or intention, and thus is determined by practical considerations, such as whether it would be immoral or imprudent to bring about the belief, whereas answering the latter question issues in a belief, and thus is determined by reasons which speak to the truth of the proposition to be believed.

Although answers to the former question might be couched in terms of reasons *for* belief, in the generic sense that they are judgements which express some kind of positive evaluation of belief, they do not guide one in deliberation whether *to* believe that *p*, because they fail to engage the normative perspective from within which this question is raised. It may be desirable for the cuckolded husband to believe his wife faithful, but this fact does not speak to the question whether it would be correct to believe it. The only way for him to answer this latter question, and thus to conclude his deliberation, is by determining whether his wife is deceiving him. He cannot arrive at the belief that his wife is faithful just by deliberating whether it would be desirable to believe it: he can only do so by deliberating whether his wife is faithful.

Conflating the doxastic question whether to believe that *p* with the practical question whether to bring about the belief that *p* has misled philosophers into thinking that evidential and practical considerations issue in competing answers to the same question, and that therefore we must decide which has priority. If I am right, only evidence for and against the truth of *p* is relevant to answering the doxastic question whether to believe that *p*, whereas only the desirability of believing that *p* is relevant to answering the practical question whether to bring about the belief that *p*. We do not need a substantive prudential or moral argument to secure the truth of evidentialism, because evidentialism is built into the very nature of doxastic deliberation.[11]

Amherst College, Massachusetts

[11] Thanks first and foremost to Casey Perin for extensive discussion of all aspects of this paper. Thanks also to Jennifer Acker, Jonathan Adler, Stephanie Beardman, James Bell, Justin D'Arms, Andrew Dole, Matt Evans, Jyl Gentzler, Alexander George, Elizabeth Harman, Daniel Jacobson, Jeffrey Kasser, Niko Kolodny, Sharon Street, Jonathan Vogel, an audience at Dartmouth College, and the members of a graduate seminar at the University of Massachusetts. Amherst, for comments on earlier drafts. Work on this paper was supported by an Amherst College Trustee Faculty Fellowship.