

teorema

Vol. XXVIII/3, 2009, pp. 113-130

[BIBLID 0210-1602 (2009) 28:3; pp. 113-130]

A Quasi-Pragmatist Explanation of Our Ethics of Belief*

Veli Mitova

RESUMEN

La mayoría de nosotros acepta normas epistémicas: nos guían a la hora de formar, mantener y revisar nuestras creencias. Un síntoma de esta aceptación es nuestra aprobación de la crítica epistémica. ¿Pero por qué aceptamos normas epistémicas? Se argumenta que lo hacemos como parte constitutiva de tener intenciones. El argumento es que, constitutivamente, la intención requiere curiosidad sobre el mundo; y la aceptación de normas epistémicas se sigue de esta curiosidad.

PALABRAS CLAVE: *normas epistémicas, aceptación de normas, ética de la creencia, intención, pragmatismo.*

ABSTRACT

Most of us accept epistemic norms: we are moved by them when we form, maintain, and revise our beliefs. One symptom of this acceptance is our deference to epistemic criticism. But why do we accept epistemic norms? I argue that we do so as a constitutive part of having intentions. The argument is that intention constitutively requires curiosity about the world; and acceptance of epistemic norms comes for free with this curiosity.

KEYWORDS: *Epistemic Norms, Norm-Acceptance, Ethics of Belief, Intention, Pragmatism.*

Among the myriads of questions which impel the ethics-of-belief debates, one is surprisingly seldom asked: Why is it that we have an ethics of belief? This *explanatory* question is not to be confused with its *justificatory* cousin, which is asked often enough: Why *should* we have an ethics of belief? Despite the historical popularity and undoubted importance of the justificatory question, sometimes we would do better to ask the less glamorous explanatory question. One area to which explanatory questions seem particularly well suited is the general puzzle about how normativity fits in a world of exclusively naturalistic furniture. Only descriptive claims can help us with this puzzle in a non-circular way. And, crudely put, justifications are normative, while explanations are descriptive. So only an answer to the *explanatory* question about our belief-ethics can deliver the epistemic piece of the general normativity puzzle.

The explanatory question can take many shapes, of course. The particular version which will occupy me here is: Why do we accept epistemic norms? The answer I defend in this paper is that we accept these norms as a constitutive part of having intentions.

I. SETTING IT UP

Before defending this claim, I say a bit more about why the account advanced here is appropriately viewed as an explanation of our belief-ethics (I.1). I then map some possible alternatives to my explanation (I.2).

I.1 *Epistemic norms and acceptance*

Our beliefs can be evaluated from every point of human concern. They can be deemed conservative, bourgeois, politically incorrect, morally deplorable, socially irresponsible, aesthetically repugnant, old-maidish, or plain false. All these evaluations target the *contents* of the belief in question, and are unconcerned about its justificatory status. The jury is still out on whether such evaluations form part of our ethics of belief. At any rate, in this paper I focus away from contents-evaluations and onto *procedural* norms as the more uncontroversial constituents of our belief-ethics. Such norms discipline *epistemically correct* belief formation, maintenance, revision, and evaluation. They prescribe ways of living up to belief's formal aim, truth. These prescriptions are our *epistemic norms*, the core of our belief-ethics. The least controversial epistemic norms are neighbours of the (overly stringent) evidence-norm that W. K. Clifford championed when he introduced the notion of an ethics of belief [Clifford (1879)]. Some examples are: you ought not to form beliefs on insufficient evidence; you ought not to have inconsistent beliefs; you ought not to base your beliefs on suspect inferences.

An agent *accepts* a norm only if, when he thinks the norm applicable, he does one of two things – he either tries to observe it or feels in the wrong when he does not try. That we accept epistemic norms in this sense is obvious enough. Consider the epistemic criticisms which we trade daily. We say things like ‘But you don’t have any evidence for (believing) this’; and things like ‘That flat out contradicts everything you believe about him’; and things like ‘I’ve never given you reason to believe this of me’. When we are criticised in these ways, we rarely shrug and move on. We go, instead, to great lengths either to rationalise the criticism away or to revise the belief in question. The seriousness with which we both flourish and suffer epistemic criticism betokens our acceptance of epistemic norms.

I.2 *A Map*

So *why* do we accept these norms? Three possible answers can be harvested from available answers to the justificatory question: a pragmatist, a

moral, and a conceptual. (This inference from justificatory to explanatory answers may be problematic. But no matter: I am just mapping possibilities here.)

The pragmatist view – if it buys the truth-aiming epistemic norms which concern me here (Stich (1990), for example, doesn't) – would be roughly this: we accept epistemic norms because observing these norms is our best means of getting true beliefs; and true beliefs are necessary for the attainment of our goals [Heil (1992), Kornblith (1993)]. The moral view would be that we accept epistemic norms either as a means to, or as a constitutive part of, being morally virtuous [Zagzebski (2004) and arguably Clifford (1879)]. According to the conceptual view, acceptance comes for free with the concept of belief which is itself construed as normative [Wedgwood (2002)].

Now, I think that none of these suggestions works. To oversimplify: the first two fail to take seriously the idea that belief is not the sort of attitude to be *primarily* evaluated by its conduciveness to things we want (whatever they are), but rather by whether it is true or false.¹ The conceptual view honours this insight, but asks too much of the concept of belief. In particular, it asks the concept to play a robust *motivating* role which a cognitive concept is unlikely to sustain. I will not defend these misgivings in the paper. The point of mapping these views is merely to situate the view I will advance here as quasi-pragmatist. I do this at the end.

II. THE OVERARCHING ARGUMENT

So, the claim that I defend in this paper, which to my knowledge has not been explored in the literature, is that *having intentions constitutively requires that one accepts epistemic norms*.

II.1 *The Argument*

The outline of the argument for this claim is this:

- (P1) Having an intention constitutively requires that one is curious about how things are in the world.
- (P2) Being curious about the world constitutively requires that one accepts epistemic norms.
- (C) Therefore, having intentions constitutively requires that one accepts epistemic norms.

The argument works purely from the concept of intention. It is neutral on the question of whether, in order for one to have an intention, one needs to apply – or even be capable of applying – the concept of intention to one's attitude. Nonetheless, there are conceptual constraints which one needs to meet in or-

der to be credited with having an intention. In particular, the argument is meant to establish that it will be incorrect for us, who do have the concept, to ascribe an intention to someone who does not accept epistemic norms.

I will have little to say about premise 2 of the argument (and I will say it in a moment). The original and interesting claim here is premise 1. So the paper is dedicated to it. I defend it in two stages. First, I establish a less ambitious claim in sections III and IV (I call it ‘the Constitutive Claim’). I then extend this claim to premise 1 (section V). But first, a quick word on the two central concepts in the argument – intention and curiosity.

II.2 *Intention and curiosity*

An intention, as I will understand it, is a favourable attitude to bringing about something in the world.² An agent *S* intends to bring about *X* only if:

- (a) *S* has a favourable attitude to *X becoming* the case in the world; and
- (b) *S* has a favourable attitude to *S causing* *X* to come about.

Clearly there is much more to something’s being an intention, and I add some more characteristics shortly (section III). But the important point for now is that the current argument works solely from this agent-centred notion of intention to which condition (b) is necessary. The argument remains silent on whether this picture *exhausts* our concept of intention. The concept may well extend to pro-attitudes to certain states of affairs obtaining, without a commitment to action on the part of the agent. But the argument has nothing to say about whether that is a legitimate sub-class of intentions, whether the class involves curiosity about the world, or whether it is constituted by acceptance of epistemic norms.

I use ‘curiosity’ as a term of art.³ It denotes *a robust favourable attitude to finding out or knowing about* the object of curiosity. Calling it robust is supposed to muffle two of curiosity’s usual overtones: first, a kind of frivolity sometimes attaching to the notion, and second, the possibility that the curiosity is idle in the sense of never being acted upon.⁴ Curiosity, then, is a *serious* attitude, which *motivates* (but, of course, not infallibly) the agent to find out about the world. Serious and robust as it is, though, curiosity cannot be quite as strong a favourable attitude as an *intention*, for then the account would be circular.

Armed with this concept of curiosity, I can now say the little I have to say about premise 2, the claim that this curiosity constitutively requires acceptance of epistemic norms. I take it as a truism that not accepting epistemic norms – being deliberately unmoved by evidence, by consistency-considerations – amounts to a kind of wilful blindness to the world. *Wilful* blindness is not compossible with curiosity about the world understood as a robust attitude. So, premise 2 is at least *prima facie* plausible.

III. THE DISTINCTNESS OF CURIOSITY ABOUT THE OUTCOME

I now isolate the claim which I shall use to support premise 1. The claim, which I call ‘the Constitutive Claim’, is that *an intention is partly constituted by the agent’s curiosity about whether the intended outcome will come to obtain once he has acted on the intention.*

III.1 *The Constitutive Claim*

Suppose I form the intention to finish writing this paper. The Constitutive Claim is that if I really intend to finish this paper, I cannot be indifferent to the *result* of my endeavours to do so. In particular, simultaneously with forming the intention, I automatically become curious about whether two states of affairs will come to obtain once I have attempted to finish it:

- (a) whether the paper is complete; and
- (b) whether it is I who made it so.

The automatic curiosity about these states of affairs arises directly from the definition of intention (II.2). The paper’s being complete and complete due to me are the two conditions under which the intention to finish the paper is fulfilled, and so under which it is typically relinquished.

Now, this curiosity is often obscured from us in practice, because the success of many of our ordinary intentions is pretty much guaranteed by the familiar means which we take to fulfilling them. The curiosity about the outcome becomes more vivid as our intentions grow more complicated and long-term. Nonetheless, the claim to support premise 1 is that this curiosity is always there because it is *constitutive* of intention.

III.2 *Two other attitudes*

The first step to defending this claim is to distinguish curiosity from two sets of attitudes arguably also enmeshed in having an intention: what I will call ‘the Davidson set’ and ‘the instrumental set’. First, as the standard Davidsonian story has it, an intention involves beliefs concerning the concepts featuring in the content of the intention. For instance, I cannot form the intention to finish this paper unless I have at least the following beliefs:

- (i) that some of the paper has been written;
- (ii) that the paper is incomplete;⁵ and
- (iii) that I am capable of completing it.

This is an example of intending to perform a non-basic action. But the Davidson

set equally accompanies intentions to perform basic actions. The impossibility of forming the intention unless I had these beliefs⁶ stems exclusively from the concepts featuring in the content of the intention. I could not intend to finish – rather than start – the paper, if I did not believe that I had at least written some of it; I could not intend to finish it – rather than edit it – if I did not believe it to be incomplete; and I could not intend to finish it *myself* – rather than have a clever friend finish it for me – if I did not believe myself capable of doing so. Without having this particular set of Davidson beliefs, I cannot have *this* particular intention with *this* particular content.⁷

The second set of attitudes arguably enmeshed in the concept of intention is the instrumental set. When I form the intention to finish my paper, I am forming the intention to undertake some actions conducive to a complete paper. So I need to have some beliefs about the means which I will take to this end – that by fixing those paragraphs, and scrapping that argument, I would finish the paper.

Curiosity about the intended outcome is distinct from both the Davidson and instrumental sets, on two counts: in its direction of fit and in its object. First, the Davidson and instrumental sets are cognitive attitudes, they are beliefs, while curiosity is a *conative* attitude to finding out about the outcome. Second, the *object* of the Davidson set is how things are before *forming* the intention. The instrumental set's object is how things are before and during the *execution* of the intention. The object of curiosity about the outcome is how things are *once the intention has been acted on*, although the curiosity, of course, arises simultaneously with forming the intention.

In the next section I develop a positive argument for curiosity's special fitness for supporting premise 1. Here I will offer some negative remarks on why neither the Davidson nor the instrumental set necessarily involves curiosity about the world, and so why neither can support premise 1. The positive argument does not depend on these remarks, but they reinforce the distinctness of curiosity about the outcome.

III.3 *Why These Attitudes Won't Do*

The Davidson set is unsuitable for supporting premise 1, because all that is required for the possibility of forming the intention is that I have the appropriate Davidson beliefs, where appropriateness is understood as a relation *not* between the belief and how the world is, but as *consistency* between the belief and the intention. Assuming that I could somehow self-induce the beliefs that this paper exists and it needs finishing, I can form the intention to finish it, regardless of the way I formed the belief or of its truth value. So, whatever interest goes with the Davidson set, is not an interest in whether *p*, but an interest in whether *my belief about p is consistent* with having a particular intention. The Davidson set of beliefs, then, cannot lend support to premise 1.⁸

The instrumental set fails to support premise 1 in a different way – due

to its not being constitutive of intention. Two considerations show this. First, the instrumental set is required only by intentions to perform *non-basic* actions. No instrumental beliefs are requisite for the intention, say, to raise my arm. Second, the instrumental set is not required for the *formation* of an intention but for its *execution*. I cannot execute my intention without having some beliefs about what actions to undertake. But I can still *form* the intention and only then start thinking of ways of executing it. Forming the intention to finish this paper does not commit me to any particular instrumental beliefs.⁹ I may form the intention without knowing what is wrong with the paper and so without knowing what actions the fixing and finishing would involve. The instrumental set of beliefs, then, is not constitutive of intention, and so is infelicitous for supporting premise 1. (Of course, it could provide an *independent* and pragmatic explanation of our ethics of belief, as I mentioned with scepticism in section I.)

So far, I have isolated curiosity about the intended outcome as a conative attitude distinct from two other attitudes involved in having an intention. I have also briefly argued that these two attitudes cannot ground premise 1. I now defend the Constitutive Claim as the apt suitor for supporting premise 1 (section IV). I then suggest how curiosity about the outcome can be extended to curiosity about the world in general (section V).

IV. THE CONSTITUTIVE CLAIM: CURIOSITY ABOUT THE OUTCOME IS CONSTITUTIVE OF INTENTION

Can we form an intention and remain indifferent to how things will be after we have acted on it? Plainly no: by the definition of intention (II.2), for me to intend something means, in part, that I have a favourable attitude to the intended outcome *obtaining*. The task for the present argument, though, is to establish the further Constitutive Claim that this favourable attitude, when it characterises intention, is at least in part a favourable attitude to *finding out* whether the outcome obtains. I now do this, by considering two versions of a thought experiment. The first version fails to support the Constitutive Claim, but fails instructively (IV.1). The second version does support the claim (IV.2).

IV.1 *Charity I*¹⁰

Suppose Charity claims that she intends to eliminate world hunger. At her declaration, the genie of world hunger conjures two doors before her, and explains that going through the one door prevents world hunger, while going through the other induces the belief that one has eliminated world hunger. The belief, he promises, would never be controverted by the world, and would be phenomenologically indistinguishable from the belief Charity would get if she had really eliminated hunger (add a pinch of memory-

cleansing to taste). Before she makes the choice of which door to go through, he can tell her which door is which. Does she want to know, he asks.

Let me time-index the envisaged scenario so as to distinguish it more clearly from the second version of the thought experiment: Charity will make her choice between the doors at t_3 , and will go through a door at t_4 ; at t_1 the genie offers her the choice to [know at t_2 which door is which]. Suppose she turns down the offer of knowing which door is which, but goes through one of the doors anyway.

Before we get to the punch line, let me assure the vigilant critic that nothing weirder than the genie himself is going on in the envisaged situation. In particular, two weird things are not going on. First, Charity's choice is not constrained in any way. She does not *have to* make a choice between the two doors. She can choose, if she wishes, to go through neither door and pack up the genie straight back into his lamp. Second, Charity does not have any deviant beliefs concerning the genie's reality and trustworthiness, or the freedom of the choice, or the doors, or the consequences of going through them. She has true and well-behaved beliefs about all these matters.

Charity has rejected both the opportunity of *ensuring* that world-hunger ceases and the opportunity to *know* that it has. Is it correct to describe her attitude to eliminating world hunger as an *intention* (rather than, say, as a wish, or hope, or something suitably feeble)? The answer that the experiment ought to yield is 'no'. It is supposed to dramatise the intuition that we cannot intend to bring about X without thereby being curious about whether X really obtains once we have acted. The thought is that since Charity flouts the option of knowing whether the world is saved, she cannot intend to save it. But sadly, the thought experiment cannot sustain this intuition, because two other explanations of her action contend with the desired one:

- (i) She *has* the intention to save the world, but it is an *irrational intention*, because she has failed to take the appropriate means to the intended end, namely, to discover which door to go through, so that world hunger really ends.
- (ii) Alternatively, her attitude is *not an intention*, because she fails to meet a more basic condition than curiosity about the outcome. An intention, by definition is a favourable attitude to the outcome obtaining. But since Charity did not choose to go through the world-saving door, she cannot have such a favourable attitude, and so cannot have the intention.

The two alternatives may shade into each other, depending on one's preferred view of instrumental rationality. But their very availability in one form or another hinders the thought experiment from doing work for me. The reason is

that we are looking here for an example of someone who claims to have an intention but (a) fails to have it, *and* (b) fails to have it *for no other reason* but lack of curiosity about the outcome. But in the envisaged situation the availability of (i) – she has an irrational intention- means that we are not forced to think (a); and the availability of (ii) - she has not met another condition for having an intention - means that we are not forced to think (b).

IV.2 Charity 2

I have dwelled on these alternative explanations, in order to make precise what must be ruled out by an argument showing that curiosity about the outcome partly constitutes intention. So let us rewrite the experiment in a way which bars these alternative explanations. Suppose now that everything is as before, except that the genie says that Charity must arbitrarily pick a door, and she must choose now, whether once she is on the other side she learns which door was which, and so learns whether world hunger has ceased. The time-indexed description now is: Charity will make her choice at t_2 , and go through a door at t_3 . At t_1 the genie offers her the choice to [know at t_4 which door was which].

Suppose that Charity again remains indifferent to finding out whether her action will have succeeded in eliminating world hunger. Now we have ruled out the above two explanations. We have ruled out (ii) – the possibility that she doesn't have an intention for some reason other than lack of curiosity – because her going through the door, given her limited (albeit unconstrained) choice, suggests that she does have a *favourable attitude* to preventing world hunger. Or at the very least her actions do not rule out such an attitude as they did in the earlier version. We have also eliminated (i) – the possibility of her having an *instrumentally irrational* intention – because she did take whatever means was available to her at the time, namely, going through either door. She has not now flouted the option of ensuring that she saves the world, because she did not have this option. But she has flouted the option of *knowing* whether she has or not. Can we call whatever attitude moved her to go through one of the doors an intention to rid the world of hunger?

I think not. The most plausible explanation for Charity's going through the door, if it is to feature an intention at all, would invoke the intention to get the belief or sensation that she has saved the world. That she will get this belief is the only thing in common between the two doors, and if she is indifferent between them, then she is indifferent to all of their other features, including significantly, the individuating feature of a hunger-free world. But if she is indifferent to this individuating feature, then she cannot intend to eliminate world hunger.

The verdict, then, is that Charity does not have the intention to save the world. The verdict is informed by the intuition that *finding out* that one's actions have succeeded is the typical condition under which one relinquishes an

intention, since one's actions *having* succeeded is the *only* condition under which the intention gets fulfilled. If Charity really intended to save the world, chose to find out which door she had gone through, and learnt that it was the wrong one, presumably she would retain her intention to save the world. Since, however, in the imagined case, she is indifferent to finding out, we can only ascribe to her a different intention, which she now thinks will be fulfilled once she has gone through a door. Unless we resort to this intention, we cannot explain her action by appeal to intentions at all. So, we finally do have an example of someone who fails to have an intention because, and only because, she has failed to be curious about the intended outcome.

V. HAVING AN INTENTION CONSTITUTIVELY REQUIRES CURIOSITY ABOUT THE WORLD

If these thoughts are on the right track, then we have shown that intending is partly constituted by curiosity about whether the *intended outcome* comes to obtain. But premise 1 claims more than just curiosity about intended outcomes. It claims curiosity about the world at large. It is important to establish this further claim, or else the overarching argument would merely get us acceptance of epistemic norms with respect to beliefs about the outcomes of our intentions.

V.1 *From Curiosity About Outcomes to Curiosity About the World*

Some brief considerations from the holism of the mental should suffice to extend the Constitutive Claim to general curiosity about the world. (We have already bought something along these lines in the discussion of the Davidson set.) Beliefs are not isolated mental items but come in networks, the members of which partake in complex inferential patterns. There will be a two-way inferential traffic between one's beliefs about intended outcomes and one's other kinds of beliefs. So beliefs about outcomes will be *inferentially dependent* on other kinds of beliefs one has. And given the complexity of belief systems, one will often not know in advance which of one's beliefs about the world will turn out to be relevant to one's beliefs about outcomes. So, one cannot want to know how things stand with intended outcomes without also wanting to know about *whatever* bits of the world are relevant to one's belief about the outcome. So, one cannot be curious about outcomes without being curious about the world.

This suggestion is confirmed when we consider what it would mean to accept epistemic norms with respect to beliefs about outcomes alone. It is impossible for my belief that I have rid the world of hunger to be justified, without its inferential predecessors (which are not about outcomes) being likewise justified *by the same standard*. Epistemic norms capture this stan-

dard. So, either one accepts norms for all of one's beliefs or for none. Since curiosity about outcomes explains acceptance for beliefs about outcomes, we have explained acceptance for all beliefs.

V.2 *A problem*¹¹

No doubt much more needs to be said on this attempt to plug the gap between curiosity about intended outcomes and curiosity about the rest of the world. But I will only say something about the most pressing problem for the attempt. The problem is that this strategy seems to overreach itself. It seems to have the unintended consequence that if we can explain our acceptance of epistemic norms for *any* subset of our beliefs whatsoever, then we have an explanation of acceptance for all of our beliefs. Any such subset – say one's fishing or cooking beliefs – would trade in a two-way inferential traffic with one's other beliefs; and would be unjustified unless those other beliefs were likewise justified. But if this is right, then explaining acceptance of epistemic norms for one's cooking or fishing beliefs will do the trick just as well as acceptance for beliefs about intended outcomes. So we needn't have bothered with the somewhat baroque story about intentions.

I have no sure-fire fix here. But I do have a suggestion that should at least alleviate this worry. To begin with, intention – the concept from which my argument works – is more basic to agency than is our being curious about some thing or other.¹² So my explanation of our ethics of belief cuts deeper than the alternatives. Now, it may not seem that such depth is entirely relevant from an epistemic point of view. But as it turns out, the depth of connection to agency translates into epistemic depth. The fundamentality of intention makes curiosity about intended outcomes superior (for our purposes) to other types of curiosity, in at least two respects: the *strength* and the *scope* of our acceptance of epistemic norms.

Firstly, we cannot abandon having intentions in the way we can abandon particular types of intellectual curiosity, or indeed the entire enterprise of having intellectual pursuits at all. To be sure, an agent without any intellectual curiosities is an extremely impoverished agent (not to mention a dreadful bore). But he is an agent nonetheless. A creature without intentions, by contrast, is no agent whatsoever, not even an impoverished boring one. Intentions are the stuff of our agential life; intellectual pursuits are a luxury (an often sought-after luxury, but luxury nonetheless). So, intentions are more important to us *qua* agents, and so whatever pro-attitudes they secure will have greater force for us. Correspondingly, the *acceptance* of epistemic norms that we secure by appeal to intention is going to be stronger than whatever acceptance we get through appeal to other interests we have.

Second, the curiosity involved in intention brings larger tracts of our other (unrelated to intended outcomes) beliefs under epistemic norms. Fishing and cooking beliefs concern a rather limited portion of the world. Of

course, other types of curiosity – say, about quantum physics or history – require beliefs about larger portions of the world. But they are still *topically* restricted. Because intentions concern a wide range of activities, on the other hand, the beliefs one needs for beliefs about outcomes will not be confined to single topics. Beliefs about outcomes, then, will partake in a two-way inferential traffic with a much greater number and topical variety of other beliefs. So, curiosity about outcomes involves curiosity about vaster portions of the world than would curiosity about fishing or quantum physics.

The upshot, then, is that although we *could* get acceptance of epistemic norms from acceptance with respect to any subset of our beliefs, the acceptance we get in this way will be neither as strong nor as pervasive as the acceptance we get by appeal to intention.

VI. AN OBJECTION AND A REPLY

This completes the main argument of the paper. We accept epistemic norms because we are curious about the world; and we are curious about the world because we are creatures with intentions. In this section, I secure this result by considering a purported counterexample to the Constitutive Claim. The claim, which helped to establish premise 1, was that intending is partly constituted by curiosity about the outcome.

There are some intentions, the objection goes, whose outcomes one cannot find out about, such as an atheist's intentions in writing a will or his intention to commit suicide.¹³ But if one cannot find out about these outcomes, and one believes that (since one believes that there is no afterlife), then it does not make sense to want to find out. The existence of these intentions, then, impugns the Constitutive Claim.

VI.1 *Two Non-Replies*

I first offer some thoughts on two non-replies, because I think that they help to sharpen the notion of curiosity about the outcome, and to forestall some obvious misconstruals of the Constitutive Claim. I then develop a reply to the will case as an exemplar of my strategy for handling such counterexamples in general.

The first non-reply I have in mind is to deny the assumption on which the objection rests: that if one believes that something is impossible to find out about, then it does not make sense to be curious about it. Denying this assumption would dispense with this entire group of counterexamples in a much neater way than the one I opt for: we can say that even though one cannot find out about these intended outcomes, one still wants to find out, and so the Constitutive Claim remains unthreatened.

This is a non-reply, its neatness notwithstanding, because of the robust-

ness and seriousness of curiosity (II.2). The reply would attenuate the notion of curiosity about the outcome to the point at which if we had any smidgen of a favourable attitude to finding out about the outcome, the attitude would count as curiosity. The reply would do that because it would dissociate curiosity from the stronger pro-attitudes (such as desires), which are as much subject to Davidson requirements as I said intention was (II.2). A genuine desire requires at least that I do not believe myself incapable of doing the desired thing. This is what distinguishes desires from wishes. So if we denied the objector's assumption, we would reduce curiosity about the outcome to a weak pro-attitude, which does not need to meet the Davidson requirement, such as wishing or hoping that one can find out. But such attitudes hardly evince a very robust connection to how the world is, and so will be uncongenial for delivering acceptance of epistemic norms. So we must grant the objector's assumption that if I believe that it is impossible to find out about something, I cannot be curious about it.

The second non-reply I have in mind is to mould the Constitutive Claim into the conditional claim that intention is partly constituted by [curiosity about the outcome *if it is possible to find out*]. I brand it a non-reply, though it is plausible enough, because it unnecessarily etiolates the Constitutive Claim. Most of our intentions are such that we can, and do want to, find out about their outcomes. If this were not the case, we would relinquish our intentions only due to other, conflicting intentions. But as a matter of fact, the typical condition under which we relinquish an intention is when the intended outcome has obtained and we believe that it has. So, although I could comfortably take shelter under this conditional version of the Constitutive Claim, I think we can do better against the present objection. The following is a sketch of a reply which dilutes neither curiosity nor the Constitutive Claim. The unconvinced are welcome to revert to the conditional Constitutive Claim.

VI.2 *The Reply*

Let us distinguish three intention-candidates which may be involved in a will:

- (a) the intention to *leave my money to S*;
- (b) the intention (absolutely) to *ensure that S gets the money*;
- (c) the intention *that my money ends up in S's hands*.

The three intentions have different fulfilment conditions, and so differ with respect to the possibility of finding out whether their outcomes obtain. I can find out the outcome of intention (a) – I can go to several lawyers to ensure that the will is legally water-tight, that it eliminates competing claims, I can destroy all my old wills, etc. If I do these things, then I will have succeeded in leaving my

money to S. But these are things that I can, and can want to, find out.

Intention (b), by contrast, looks like an apposite counterexample to the Constitutive Claim. If ensuring means ‘absolutely guaranteeing’, then of course I cannot find out the intended outcome and so I cannot be curious about finding out. The bad news for the objector, though, is that the attitude to this guarantee cannot be one of intention, not because of any reasons to do with the impossibility of *finding out* whether the outcome obtains, but because the attitude breaches a Davidson requirement – that the agent *believe* such a guarantee to be possible. The guarantee is impossible, because there is no action that I can perform in order to ensure in the absolute sense that S gets the money. After I die, I cannot act.

Intention (c) – if intention it be – is again an intention whose outcome I cannot find out about, and so is eligible for a counterexample. Sadly for the objector, however, intention (c) does not involve a commitment on the part of the agent to act. (Call these ‘non-agent intentions’.) Yet I restricted my argument right at the beginning (II.2) only to those intentions which do involve such a commitment. (Call these ‘agent-intentions’.)

So the reply to the objection is that, appearances notwithstanding, intentions involved in a will do not constitute a counterexample to the Constitutive Claim, because either (a) they are intentions whose outcome we can, and do, want to find out about; or (b) the attitudes are not intentions for Davidson reasons; or (c) they are non-agent intentions, about which the argument remains silent.

VI.3 *Two Worrying Distinctions?*

Now this reply seems to invite two worries, corresponding to the two distinctions on which the y hinge – first, between intentions (a) and (b) and second, between agent- and non-agent intentions. First then, the strategy may be thought to introduce an implausible distinction between intending to *do one’s best* and intending to *succeed*. My opinion of this distinction is that, implausible as it is in general, it is perfectly in order in cases in which the agent cannot, in principle, guarantee success. But what is important for the current argument is that the distinction was not foisted on us by *my claim* about the relationship between intending and being curious about intended outcomes. Rather, we were inveigled in the distinction by Davidson considerations.

The second worry is that my reply seems to bail out the Constitutive Claim at the price of generating a mystery: what is the relevant difference between agent and non-agent intentions, and why should it be quite such a hefty difference, that the former but not the latter necessarily involve curiosity about the outcome? The answer to this question is disappointingly simple. When I form the intention *to do* something, I thereby commit to seeing it through, so to speak. And I do not think that this expression is an accident: ‘seeing it through’, ‘seeing to it’, and ‘seeing it out’ are expressions capturing

a crucial feature of intentions: that they involve – as the seeing metaphor suggests – *seeing whether* I have seen it through or out. They involve, that is, curiosity about the outcome of my actions. Non-agent intentions, by contrast, intentions that something becomes the case, do not necessarily (though they may) involve a commitment to seeing it through, or out, or to it, because they do not involve such a sturdy commitment to action on the part of the agent. I think this difference between the two types of intention amply accounts for their different relations to curiosity about the outcome.

So intentions involved in a will and other such, give us no reason to fret about the plausibility of the Constitutive Claim.¹⁴

VII. CONCLUSION

We accept epistemic norms, then, by virtue of having intentions. Why is this a *quasi-pragmatist* explanation of our ethics of belief? The pragmatist element is the intuition that since our acceptance of epistemic norms is a matter of motivation, of being moved by these norms, it must have something to do with our having goals and projects. The account accommodates this sound pragmatist tenet by construing acceptance as constitutive of intention. But it does so without falling prey to the main anti-pragmatist worry, mentioned in the introduction. The account pays homage, that is, to the intuition that truth – rather than conduciveness to our goals – is the primary standard for evaluating our beliefs. This is the quasi-element in the account: it is not the pragmatic, but the *epistemic*, constituent of intention which does the explanatory work in sourcing our acceptance of epistemic norms. We do not accept these norms as a *means to fulfilling* the intention (as the pragmatist has it), but because, in order to *have* an intention at all, we must want to know how the world is. This quasi-element allows us to distance epistemic norms from pragmatism sufficiently to give them autonomy. So the pragmatism allows us to keep our normative feet on the naturalistic ground of intentions. The quasi-part is what makes our belief-ethics an *ethics* – rather than a theory of practical rationality – and an ethics of *belief*, rather than an ethics of action.

Instituto de Investigaciones Filosóficas
UNAM, Zona Cultural
Circuito Mario de la Cueva, Ciudad Universitaria
Coyoacán, C.P. 04510, México, D.F.
E-Mail: velimitova@filosoficas.unam.mx

NOTES

* Many thanks for invaluable discussions of the ideas presented here to Simon Blackburn, Edward Craig, Jonathan Dancy, Jane Heal, Ward Jones, Hugh Mellor, Thad Metz, the Moral Sciences Club, the Rural Sciences Club, my colleagues at the Rhodes Philosophy Department, and to the participants in the *4th Meeting on Pragmatism: Truth and the Ethics of Belief* (held in 2008, Murcia). Thanks to the Andrew Mellon Post-Doctoral Fellowship and Rhodes University for their financial support.

¹ For an argument along these lines in the context of the justificatory question, see Mitova (2008a).

² Intentions, of course, are not only about positively bringing about changes. They can also be about the maintenance or prevention of states of affairs [Blackburn (1996), p. 196]. But we needn't get addled in detail here: both conditions (a) and (b) are to do with the world. And the world includes the agent himself. So, self-directed intentions are still about the world. This connection to the world is all that needs to seep into the rest of the argument. So from now on I will only talk of intentions to bring about states of affairs, but the argument equally extends to intentions to maintain or prevent states of affairs.

³ Thanks to Jane Heal for the term.

⁴ A less frivolous and idle term for this attitude would be 'interest in finding out'. But 'interest' is sadly ambiguous in that it can mean that the thing I am interested in is beneficial for me. This ambiguity may make the argument slide around, so I use the term 'curiosity'.

⁵ This condition applies only to intentions to *bring about* X. If the intention is to *maintain* X, X had better obtain. But intending to maintain X would have its own equivalent of impossibility condition (i): If X is not the case, then X cannot be *maintained*.

⁶ This requirement of having positive beliefs may be too strong. Perhaps all that is needed is that I *do not believe* that the negations of (i)-(iii) obtain. If so, this would not affect the overarching argument, which depends neither on the concept of belief nor on the Davidson set.

⁷ See footnote 8 for the stronger view that the Davidson set serves to individuate the attitude *type*.

⁸ What if Davidson beliefs are not merely necessary for fixing the *content* of the intention, but for the state's being an *intention* at all? Then mere consistency will not be enough. Rather, we will need the beliefs in question to be true in order for the attitude to count as an intention in the first place. This suggestion would make superfluous my argument from curiosity. My reply is that, although on this suggestion Davidson beliefs get us a fairly good connection to the world, it is not of the right kind to give us acceptance of epistemic norms. That is, the *truth* of the first-order Davidson belief is enough to make my attitude an intention. It doesn't matter how the belief was formed or whether it is justified. For instance, I wake up confused from dreaming I had a bicycle, and form the intention to clean it. Unbeknownst to me, my friends have bought me one in the night. So, the important Davidson belief is true but epistemically as rogue as one might wish. Yet I can perfectly coherently have the intention to clean my bicycle (as long as I also truly believe it is dirty). So appeal to the Davidson set will not deliver the right interest in the justificatory status of one's beliefs, to explain our acceptance of epistemic norms. Curiosity, by contrast, comes automatically with acceptance of epistemic norms (see II.2).

⁹ Of course, forming an intention is committing to performing some action or other. And so if I believed that there was no action whatever that could be performed in order to finish the paper, then I could not form the intention to finish it. So, intention constitutively requires the belief that the instrumental set accompanying it is not empty. But this is still the *Davidson* belief that I *am capable* of finishing the paper. By contrast, no particular instrumental belief is a conceptual ingredient of the intention.

¹⁰ Thanks to Nick Tosh for a version of this experiment.

¹¹ Thanks to Ward Jones for this objection.

¹² Of course, without curiosity about outcomes, there are no intentions, on my account. But this is not a problem, for two reasons. Firstly, the objection turns on curiosity other than curiosity about outcomes (otherwise the objector and I would not be disagreeing). Secondly, the present point is about the comparative depths of the *concepts* of curiosity and intention, *relative to* the concept of agency.

¹³ Thanks to Hugh Mellor for the will counterexample, and to Ward Jones for the suicide counterexample.

¹⁴ These intentions involve the supposed *practical* impossibility of finding out. What of cases of *logical* impossibility? My intention to commit the perfect crime, say, necessarily involves that I take a potion which will erase my memories of having performed the crime. This makes it seem logically impossible for me to find out whether I have committed the perfect crime. So, as before, either I can't intend to commit the perfect crime or the Constitutive Claim is impugned. My strategy here is a lot simpler: I deny that it is impossible to find out that one has committed the perfect crime. Before the potion is taken I know that I have completed all but the last step of the perfect crime *and* that by taking the potion I will *absolutely guarantee* the full completion of the crime. So, as I take the potion, I know that I have completed it. Of course, I forget that immediately, but the point is that the intended outcome is possible to find out. So, such intentions present no problem for the Constitutive Claim.

REFERENCES

- ADLER, J. E. (2002), *Belief's Own Ethics*, Cambridge, Mass., MIT Press.
- BLACKBURN, S. (1996), *Oxford Dictionary of Philosophy*, Oxford, Oxford University Press.
- CLIFFORD, W. K. (1789), "The Ethics of Belief", reprinted in Burger, A. J. (ed.), *The Ethics of Belief* (2001), Roseville, CA, Dry Bones Press, pp. 9-40.
- HEAL, J. (1988) "The Disinterested Search for Truth", *Proceedings of the Aristotelian Society*, vol. 88, pp. 97-108.
- HEIL, J. (1992), "Believing Reasonably", *Noûs*, vol. 26: 1, pp. 47-72.
- KORNBLITH, H. (1993), "Epistemic Normativity", *Synthese*, vol. 94, pp. 357-78.
- MITOVA, V. (2008a), "Why pragmatic justifications of epistemic norms don't work", *South African Journal of Philosophy*, vol. 27:2, pp. 139-50.
- (2008b), "Why W. K. Clifford was a closet pragmatist", *Philosophical Papers*, vol. 37:3, pp. 471-88.
- SHAH, N. and VELLEMAN, D. (2005), "Doxastic Deliberation", *The Philosophical Review*, vol. 114: 4, pp. 497-534.
- STICH, S. P. (1990), *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*, Cambridge, MA, MIT Press.

- WEDGWOOD, R. (2002), "The aim of belief", *Philosophical Perspectives*, 16, *Language and Mind*, pp. 267-97.
- ZAGZEBSKI, L. T. (2004), "Epistemic Value and the Primacy of What we Care About", *Philosophical Papers*, vol. 33:3, pp. 353-77.