

DAVID JOHN OWENS

DOES BELIEF HAVE AN AIM?

ABSTRACT. The hypothesis that belief aims at the truth has been used to explain three features of belief: (1) the fact that correct beliefs are true beliefs, (2) the fact that rational beliefs are supported by the evidence and (3) the fact that we cannot form beliefs 'at will'. I argue that the truth-aim hypothesis cannot explain any of these facts. In this respect believing differs from guessing since the hypothesis that guessing aims at the truth can explain the three analogous features of guessing. I conclude that, unlike guessing, believing is not purposive in any interesting sense.

Why ascribe an aim to belief? Those who ascribe an aim to belief hope to explain at least three features of belief.¹

First, beliefs are right or wrong, correct or incorrect. Call the basis of this form of normative assessment *the standard of correctness* governing belief. Many writers think we can explain this standard of correctness by supposing that belief has a goal which it is striving towards, or a function which it is seeking to discharge. A correct belief is one which attains its goal and discharges its function.

Second, beliefs are rational or irrational, justified or unjustified. These too are normative assessments – rational belief is good belief, irrational belief is bad belief – but they are, to this extent, independent of judgements of correctness: a correct belief can be irrational and an incorrect belief can be rational. Call the basis for this second form of doxastic assessment *epistemic norms*. Many writers have thought that the standard of correctness for belief and the epistemic norms governing the acquisition of belief are connected as follows: the standard of correctness for belief sets a goal and epistemic norms are instructions about how to reach that goal.²

Making this connection simplifies our account of the normative assessment of belief. It also removes a worry which many people have about epistemic norms. We are familiar with instrumental or teleological norms from the case of action. The source of the



authority of these norms seems clear to many philosophers: they tell anyone who happens to have a goal how they should go about pursuing that goal. By contrast, epistemic norms appear *sui generis*. Why should believers be held to these norms? Where does their authority over us come from? Call this *the problem of authority*. If believing involves having an aim and these norms can be seen as instructions about how to achieve that aim, the solution is obvious.³

Postulating a goal for belief explains not only the authority of epistemic norms but also their motivational efficacy. And, in the eyes of many philosophers, this is essential for epistemic norms to have any authority. In calling a belief rational or irrational, justified or unjustified, we are judging not just the belief but also the believer. A person who believes irrationally is, to that extent, an irrational person: a normative assessment of someone's belief is a normative assessment of *them*. Now (many writers think) this can be so only if that person is capable of being moved by epistemic norms, only if these norms can guide their acquisition of beliefs. So we must explain what is it about believers which ensures they have this capacity, that they can indeed be moved by their recognition of the authority of these epistemic norms. Call this *the problem of doxastic control*.

Suppose all believers have a characteristic goal. Then all believers will be motivated to pursue that goal. Since epistemic norms tell us how to achieve that goal, it is no surprise that all believers have (at least some) motivation to conform their beliefs to these norms. Anyone who lacked all such motivation wouldn't count as a believer precisely because they would lack a goal that all believers necessarily have (Railton, 1997, pp. 54–59; Velleman, 2000, pp. 182–188). Problem solved. In the last section of the paper I shall discuss the nature and limits of doxastic control but first let us examine the account of the authority of our epistemic norms that is on offer.

TRUTH AS THE AIM OF BELIEF

What is the goal of belief? The most common suggestion is truth. On this hypothesis, when we intend to form a belief, we are intending to form a true belief. Believing is essentially purposive, so one is

not really forming a belief unless one is intending to form a belief which is true, unless one has that purpose in believing.

Since the truth of a proposition is, in general, quite independent of whether it is believed by me (or anybody else), truth is well suited to be the standard of correctness for belief. But not every action governed by a standard of correctness is performed with the purpose of meeting that standard. For example, false assertions are incorrect assertions but not every assertor has the goal of making an assertion only if it is true. Liars make genuine assertions with the aim of asserting something false, with a view to violating assertion's standard of correctness. If beliefs differ in this respect, if belief is necessarily adopted with the purpose of being true, that fact accounts for belief's standard of correctness and thereby discharges the first of our explanatory obligations.

What of the epistemic norms which govern the acquisition of belief? On the present model, believing propositions is somewhat like setting out to acquire stamps but only those which are Penny Blacks. We may try to explain the norms governing Penny Black collecting activity by seeing them as instructions about how to achieve that hobby's characteristic aim. Similarly, we might hope to explain the norms governing the acquisition of belief by referring to belief's characteristic aim (Railton, 1994, p. 75). A believer is intending to get at the truth and the way to do that is to base one's belief on evidence. That is why the epistemic norms, the norms which determine the rationality of belief, are norms of evidence.

Again, none of this applies to assertion. A liar may be perfectly reasonable in asserting that *p* regardless of the evidence. Just because a state or an activity is governed by a standard of correctness it does not follow that we can explain what makes it reasonable to be in that state or to engage in that activity by reference to this standard. This follows only if we further suppose that the phenomenon in question is purposive and that its purpose is to meet its own standard of correctness.

The truth-aim theory sets out to explain the fact that reasonable belief requires evidence by ascribing a purpose to belief. How exactly does this explanation work? The two key notions invoked are 'truth' and 'aim at'. The account will avoid circularity only if both notions can be defined without reference to the need for evidence.

Truth is relatively unproblematic: there are some who want to define truth in epistemic terms but, for the purposes of this paper, I shall assume that truth is a non-epistemic notion. What of 'aim at'?

Here we seem to encounter a problem. An agent with an aim has set themselves an end and is thereby obliged to take the means to that end. Call this *the means-ends principle*. The problem is that in order to apply this principle, the agent must already have some beliefs about what constitutes a means to that end. Where do these beliefs come from and what makes them reasonable? So long as we are concerned just with the rationality of agency, we can parry this question: agents conform to the means-ends principle by adopting the means they *believe* will secure their end and that policy makes them *practically* rational. But shouldn't a theory of epistemic normativity tell us what makes all beliefs (including beliefs about means) reasonable?⁴

Where truth is the aim in question, the means are evidence. Now principles of instrumental rationality alone are not going to explain why it is reasonable to believe that perception or memory or inductive inference are sources of evidence, are reliable guides to the truth. A person who has no idea how to go about seeking the truth will not be helped by our reminding him of the means-ends principle. But suppose we hold, on whatever grounds, that a sensory experience as of *p* is a reliable indicator of *p*'s truth. We may still ask: what has the fact that someone has a sensory experience as of *p* got to do with whether he ought to believe that *p*? The truth-aim theory of epistemic normativity sets out to answer *this* question and to answer it when it is raised about any belief (including our beliefs about evidence). The truth-aim theory tries to explain the normative force of evidential norms, not what constitutes evidence.⁵ How does it do this?

We are wondering whether a given subject ought to believe that *p*. He is aware of a certain amount of evidence for *p* but what has this got to do with whether he ought to believe *p*? (Perhaps he has been offered a huge prize to believe not-*p*: doesn't that make it more reasonable for him to believe not-*p*, regardless of the evidence?). The truth-aim theory offers the following reply: this subject won't count as believing *p* at all unless he adopts this belief with the aim of believing it only if it is true; given that he has this end, the means-

ends principle (whose authority we already acknowledge) tells us that his beliefs about the evidence for and against *p* are relevant to the rationality of belief in *p* (the prize notwithstanding).

It seems the hypothesis that belief aims at the truth can (without circularity) tell us why a believer should base his beliefs on (apparent) evidence rather than just on the prospect of financial reward. Has it thereby done all we should expect an instrumental theory of epistemic normativity to do? Not quite. There is a certain vagueness in the means-ends principle. In the practical case there is usually a choice of possible means each of which would be reasonably effective in attaining the end. Some of these means will make the attainment of the end more likely than others, but more effective means may also be more costly, or more unpleasant, or be more likely to conflict with the attainment of other ends. Clearly there is a balance to be struck here and it is the job of instrumental rationality to strike that balance. Should similar issues arise in the epistemic realm and given that believing is, as we are supposing, purposive, we shall expect the principles of instrumental rationality to resolve these issues for us.

Similar issues do indeed arise in the epistemic realm once we ask *how much* evidence is required for rational belief.⁶ Rational belief is rarely based on conclusive evidence and once inconclusive evidence is admitted, we must ask what level of inconclusive evidence suffices for rational belief. Requiring more evidence makes the attainment of the goal of truth more likely but it may also impose costs and burdens which a rational person would wish to avoid. We shouldn't expect precise answers here – norms of evidence are doubtless vague – but the truth-aim theory should be able to account for such facts as the following: (i) reasonable belief in *p* requires some evidence in *p*'s favour (ii) reasonable belief requires more than a little evidence in *p*'s favour (iii) reasonable belief in *p* requires that there be more evidence in favour of *p* than in favour of not-*p*.

To bring home the importance of specifying how much evidence is required for rational belief, consider how we should respond to the discovery of an apparent contradiction in our belief system. The answer might seem clear: since belief aims at truth and contradictory beliefs cannot all be true, we must remove the contradiction by abandoning (at least) one of the contradictory beliefs. But it is

not quite that simple. Firstly, not all contradictions are obvious. If the discovery of this contradiction involved a complex proof, or some contentious principles of reasoning, there will be a question of how much further confirmation of the soundness of the proof we should require before accepting that there is indeed an inconsistency. Secondly, even when the contradiction is obvious it need not be obvious what ought to be done about it. All we now know is that there is a falsehood among the inconsistent beliefs and there might be a substantial number of important beliefs in that inconsistent set. Should we abandon them all? Should we single out some member of the set to preserve consistency and if so, which one? Clearly the relative strength of the evidence in favour of individual beliefs, taken together with our need to have beliefs about a range of matters (a non-evidential consideration) should determine the answer. If there is no belief with conspicuously weak support, it may be reasonable to hang onto the inconsistent set for the time being rather than dropping one of its members (Foley, 1993, pp. 162–173; Harman, 1999, p. 19).

What is in question throughout this paper is the rationality of belief, *not* the rationality of the activities of evidence-gathering, or of evidence-storage and retrieval, or of thinking about the value of evidence before you. I shall call these activities *inquiry*. Inquiry often precedes belief formation and is motivated by all sorts of non-evidential considerations: only if I am interested in whether OJ is guilty will I go to the trouble of reading the newspapers, thinking about the evidence reported, trying to memorise it and so forth. This paper focuses on the role non-evidential considerations play *after* we have ensured (perhaps by means of inquiry) that a certain quantity of evidence on the matter is before us. Suppose the evidence favours OJ's guilt. Still the question remains: does it favour OJ's guilt enough to make it reasonable for me to make up my mind on the matter and form the belief that OJ is guilty? I maintain that this question cannot be answered without invoking non-evidential considerations (the importance of the issue etc.).⁷ And this question is not equivalent to the question: should I go on gathering evidence or thinking about OJ's guilt? There are all sorts of reasons why I might stop inquiring into OJ's guilt without forming a belief (I got bored) or continue my inquiry despite having formed a belief (I need to convince someone else). The same points apply where the issue

is whether I ought to abandon (rather than form) a belief, as in the example of discovering a contradiction.⁸

If we were out to explain the rationality of inquiry, we could simply observe that such inquiry is motivated by the desire for a true belief on a certain matter, where the satisfaction of this desire is constrained in the usual way by other demands. But, I take it, the truth-aim hypothesis is *not* put forward as part of a theory of the rationality of inquiry but rather as an account of the rationality of belief formation (and abandonment). That is a quite different task and the prospects of success are much more uncertain.

I have sketched how the truth-aim hypothesis *might* explain both belief's standard of correctness and what makes belief reasonable. In the next section, I shall assess the solution to the problem of authority we have just been offered, leaving the problem of doxastic control until later. My conclusion will be that the explanatory strategy behind it is perfectly sound: if believing were purposive, we could explain the above features of belief by reference to belief's purpose. We can see this by applying the strategy to something which does aim at the truth, which really is purposive in just the way we supposed believing to be: *guessing*. Yet the obvious differences between believing and guessing suggest that belief does not aim at the truth in the sense required by our explanatory project.

BELIEVING AND GUESSING: RATIONALITY

Before going any further, we must formulate the truth-aim hypothesis under discussion a little more precisely:

ϕ -ing that p aims at the truth if and only if someone who ϕ s that p does so with the purpose of ϕ -ing that p only if p is true.⁹

Three points are worth noting. First, the aim of truth is stated negatively as the purpose of avoiding error; it imposes no obligation on believers to seek out all, or even some truths. Since we are trying to find an aim which is common to all believers, negativity is a must. Second, the aim applies to beliefs taken individually; one can't pursue truth in this sense by adopting a particular belief, regardless of whether one has reason to think it true, provided its adoption is likely to lead to the adoption of many more true beliefs.¹⁰

Thirdly, an aim may be pursued more or less well. Someone who bases a belief on wholly inadequate evidence still counts as believing provided they are intending to get at the truth. It may be hard to draw the line between an irrational believer who is aiming at the truth, albeit ineptly, and someone who isn't really intending to get at the truth and so doesn't count as forming beliefs, whether rationally or irrationally. But there is a distinction here, as there is between a very poor shot and someone who isn't aiming at the target at all.

Guessing aims at the truth in the sense just defined. Clearly, truth is the standard of correctness for a guess and, I maintain, what explains this is the fact that a guesser intends to guess truly. The aim of a guess is to get it right: a successful guess is a true guess and a false guess is a failure as a guess. Someone who does not intend to guess truly is not really guessing.

In this respect guessing is unlike, for instance, imagining. To imagine that I have been elected President is to imagine that it is true that I have been elected President but imagining does not aim at truth as guessing does. The act of imagining may be a complete success in that it is extremely gratifying or deeply revealing or merely distracting: truth is not required for imaginative success. Similarly for supposition: I may suppose or hypothesise that an extension is added to my house as an aid to working out how desirable this building would be. It is no part of my purpose in supposing that *p* that I should suppose *p* only if it is true, though, of course, I suppose *p* to be true and do so in order to learn something (Railton, 1994, p. 74; Velleman, 2000, pp. 112–114, 183, 250–252). But we guess with the purpose of making a guess which is true.¹¹

Both believing and guessing satisfy the above definition of 'aiming at the truth'. A believer satisfies a further condition; in believing that *p* he actually believes that the aim of belief has been achieved, for he believes that *p* is true. By contrast, a guesser merely guesses that the aim of guessing has been achieved (and a Penny Back collector merely thinks it sufficiently likely that he has got a Penny Black).¹² Someone might suggest that we redefine 'aiming at truth' so as to exclude guessing etc. as follows:

ϕ -ing that *p* aims at the truth if and only if someone who ϕ s that *p* does so with the purpose of ϕ -ing that *p* only if (*p* is true and they believe this).

But such a restriction would be no better motivated than one which excluded belief from the set of states which aim at the truth by instead requiring that the subject merely guess *p* to be true. Better take the generic notion of aiming at the truth and see what work it can do for us in accounting for the normative character of both believing and guessing.

Guessing's standard of correctness is embodied in a purpose which any guesser must have. This should lead us to expect that guessing is governed by evidential norms. Since we guess with the purpose of guessing correctly and evidence is an indication that the guess will be correct (i.e., true), a reasonable guess should be based on the available evidence. Indeed it is so. I am asked to guess approximately how many people voted in the last US Presidential election and I plump for 98%. That is a silly guess, given what I know about turnout at previous elections. And it remains silly even if I happen to be right. I also guess that it will rain at some point during the next two weeks within the boundaries of Manchester city centre. That is a reasonable guess, one well grounded in my extensive experience of Manchester city centre. And it remains a good guess even if it turns out to be wrong.

Can we say anything more about the quality of evidence required to make a guess reasonable? If what I said in the last section is correct, the truth-aim hypothesis shouldn't just tell us that rational guesses are somehow based on evidence; it should also give us some way of determining how much evidence is required for rational guessing. Clearly this will be a function of two factors: first the aim of making only correct guesses which any guesser must have and second the subject's other desires and purposes. Not getting it wrong may be the only purpose constitutive of guessing but if that were the guesser's only purpose, he could achieve it simply by making no guess at all. We need to integrate error-avoidance with other aims which lead the guesser to guess and thus determine when the guesser has sufficient evidence to hazard a guess.

Let's take an example. Suppose the quiz master asks me whether the earth's population is greater than 7 billion. There is \$1 million at stake. I wouldn't claim to know the answer but no answer means no money. Obviously I ought to make a guess. So I plump for the answer 'yes'. My guess is a sensible one, even though I have no

evidence on the matter, because if I don't guess at all I am certain to lose the money. In fact, a guess would be sensible even if I had to get the population of the world to within ten million. Here I'm pretty unlikely to be right but, so far as my evidence goes, any guess I make is as good as any other and I have to make some guess to win.

Alternatively, suppose that, to win, I must answer a series of questions in a short time. After being asked a question, I am given a series of clues as to the answer. Here it might be sensible for me to wait until the evidence tips decisively in favour of a particular answer before making a guess (though at some point I will want to get this question out of the way and move onto the next). Clearly such variations in the quality of evidence required for a reasonable guess are due to variations in the purposes/situation of different guessers.

This example suggests that we treat the rationality of guessing just as we would the rationality of any other action: guessing that *p* is reasonable when aiming at the truth by means of a guess that *p* would maximise expected utility. That requirement integrates the purpose constitutive of guessing with the subject's wider purposes in a familiar and well understood fashion. On this view of guessing, not every expected-utility-maximising utterance of 'I guess' will count as a guess ('I guess' is not a performative).¹³ For example, a subject offered a large amount of money simply to utter the words 'I guess that I am seven foot tall' won't count as guessing that he is seven foot tall unless he utters these words with the purpose of uttering them only if he is seven foot tall. Nevertheless whether it is reasonable to do something with the truth aim and how it is reasonable to go about doing it will be determined by expected utility.

Must reasonable guesses satisfy any of the three evidential norms we formulated for belief? The quiz example establishes that neither (i) nor (ii) apply. What of (iii)? If the evidence favours *p* over not-*p*, would guessing not-*p* always be unreasonable? I'm not sure. Suppose I must choose between *p* and not-*p* and the evidence favours *p* but I am also told that I will get a much greater reward if I correctly guess not-*p* than I will if I correctly guess *p*. Given that, would it be reasonable for me to guess not-*p*?¹⁴ It might be reasonable for me to *say* 'I guess not-*p*'. But in saying this, am I

sincerely guessing that not-p? Am I really trying to win the prize by stating the truth? Or am I simply saying whatever will maximize my expected winnings?

One line would be that one is not really aiming at the truth unless one follows the (apparent) balance of evidence, so the above statement could not be an expression of a reasonable guess. But someone else might take the view that one can be aiming at the truth in guessing that not-p provided one does not have evidence sufficient to *establish* that p (i.e., sufficient to make it unreasonable to fail to believe p). If so, 'I guess that not-p' might well express a sincere attempt to get at the truth – the only way of winning – as 'I guess that I am seven foot tall' does not. Whatever view we take on this matter, the evidential norms governing guessing clearly differ from those governing belief.

In the quiz situation, with the \$1m on offer (and no clues forthcoming) I can make a perfectly reasonable guess as to the exact population of the world but I am in no position to form a reasonable belief about whether there are at least 7 billion people on the evidence before me. Even less can I reasonably believe not-p rather than p because the rewards for correctly believing not-p are greater than those for correctly believing p. However great an incentive I am given for forming a view about whether p, reasonable belief requires (at least) that I have more evidence in p's favour than against it. It looks as if our epistemic norms are much more demanding when it comes to determining the means which ought to be taken in pursuing the goal of truth. Can we explain this difference between believing and guessing by way of instrumental reasoning, on the hypothesis that belief aims at the truth?¹⁵

Some might think that believing is unlike guessing in that the *only* purpose with which the belief that p is formed (and not just the only purpose *essential* to believing) is the purpose of believing that p only if p is true. If so, there would be no question of having to integrate the goal of getting this particular matter right with other goals. Yet this can't be correct. Were avoiding error one's only objective in forming any particular belief, one could very easily achieve that objective simply by forming no beliefs at all. Reasonable people form beliefs and they do so on the basis of much less than conclusive evidence. This practice makes sense, on the purposive model of

belief, only if believers have some objective in forming beliefs other than the mere avoidance of false beliefs. Can we adapt our model of guessing rationality to the case of belief in an effort to explain how the rational believer integrates the aim of avoiding error with his other goals whilst also accommodating the obvious differences between believing and guessing?

One strategy would be to suppose that the believer attaches a higher value to avoiding error than the guesser. Since the believer's interest in the goal of truth is stronger than the guesser's, it will be more weighty relative to the agent's other goals and so more trouble will be taken to get things right. On this account, the \$1m prize must be enough to make the agent risk an erroneous guess but not a false belief. This is totally implausible. No sane person would want to avoid erroneous belief on this demographic issue more than they want a 50% (or even a 0.1%) chance of \$1m; yet the rational believer will still remain agnostic. Clearly we can't determine what beliefs it would be reasonable for someone to form by weighing the believer's interest in truth against the other interests which might be served by the formation of a belief.¹⁶

Another strategy would be to limit the sorts of goals which are relevant to the rationality of belief. So we might explain the quality of evidence required for reasonable belief by postulating a number of specifically epistemic goals – error avoidance, comprehensiveness and perhaps simplicity – and assigning weights to them. For example, most believers are prepared to run a substantial risk of error in order to avoid a widespread agnosticism. These epistemic goals could then be integrated with one another by seeking to maximise the expected *epistemic* utility of one's beliefs on given evidence, thus determining what it was reasonable to believe on that evidence.¹⁷

I doubt whether there is any way of isolating a purely epistemic set of goals. Think of all the factors which go to determine just how curious one is, which fix just how important it is to have a (correct) view on a certain matter. They will determine how much evidence one requires to make up one's mind. It is not the sheer comprehensiveness of a belief system (however that is to be measured) which makes it worthy of belief but rather the significance of the topics it is informing you about. Clearly practical concerns are critical here

and so we are confronted once more by the difficulty of using instrumental norms to integrate these ‘non-epistemic’ concerns with the goal of error-avoidance.¹⁸

But even if we could specify some set of distinctively epistemic goals which underlay our epistemic norms, this won’t help us with the broader explanatory project. By insisting that our other ‘non-epistemic’ goals are simply irrelevant to the rationality of belief we deepen the very crisis of normative authority which ascribing a goal to belief was meant to resolve. Why should we be held to norms which take no account of many of the things which matter most to us? And how, in any case, can we be moved by such norms?¹⁹

That completes my case against the truth-aim theorist’s account of the rationality of belief. Someone convinced by everything said so far may still wonder whether we are asking too much of the truth-aim theory. Might that theory provide a perfectly satisfying explanation of why evidence is relevant to the rationality of belief without also telling us how much evidence is required to make a given belief reasonable? Certainly the theory’s advocates have not felt any pressure to undertake the latter task. This may be because they overlooked the fact that belief isn’t the only thing which aims at the truth in their sense. But once this fact is acknowledged, can’t the truth-aim theorist offer to explain the similarities between believing and guessing – their shared correctness condition and the relevance of evidence – without pretending to account for their differences? I think not.

The truth-aim theory seeks to explain both the correctness condition of belief and the role of evidence in justifying belief by ascribing a *purpose* to every believer. If this talk of purpose is to be more than an empty metaphor, the notion of purpose invoked must be one that does explanatory work for us outside the domain of epistemic norms. And, surely, it is our ordinary notion of a purpose which the truth-aim theorist means to be employing. Now such purposes interact with each other in certain familiar ways, so if a subject really does form a belief with the purpose of forming it only if it is true, his pursuit of that goal should be constrained by his other goals and objectives in (something like) the usual fashion. That is just what happens in the case of guessing which is why we are fully justified in treating guessing as a purposive activity. But we failed in

our attempts to treat the ‘purpose’ allegedly shared by all believers as something which interacts with their other objectives in familiar ways. So what is left of the idea that they share a purpose? The truth aim theorist can’t evade this problem simply by declining to explain the differences between believing and guessing, since what these differences suggest is that our notion of a purpose does not apply to belief at all.

BELIEVING AND GUESSING: CONTROL

Until now I have concentrated my attention on attempts to explain the authority of our epistemic norms. But even if ascribing an aim to belief cannot help with that question, an aim might still be needed to explain certain facts about how belief is motivated, facts which can’t be accounted for simply by noting either belief’s standard of correctness or the content of our epistemic norms. Perhaps the truth-aim hypothesis can help to explain the nature and limits of doxastic control.

It is possible to violate, deliberately and self-consciously, many of the norms which govern our mental life. But philosophers who hold that belief aims at the truth have inferred from this that a believer can’t self-consciously disregard indications of the truth (i.e., evidence) in forming a belief. Williams writes that

If in full consciousness I could will to acquire a belief irrespective of its truth, it is unclear that before the event I could seriously think of it as a belief, i.e., as something purporting to represent the truth (1973, p. 148).

If a believer must form a belief with the purpose of believing something true, he cannot “in full consciousness” decide to believe that he is a foot taller than he is on the grounds that it would be very desirable for him to have this belief (as a boost to his confidence) (Velleman, 2000, p. 118).

For Williams’ line of reasoning to have any force ‘aim at truth’ must mean more than ‘have truth as your standard of correctness’. Assertion aims at truth in that sense and yet it would be superstitious to imagine that this somehow deprived people of the ability to make groundless assertions. What might deprive people of the ability to believe whatever they like, in full consciousness of what they are

doing, is the fact that belief's standard of correctness is embodied in the very purpose of the belief-former. That restricts the grounds on which he can form a belief in a perfectly non-magical way.

To see how this works, consider the action of shooting at a goal. No one can shoot at a goal unless they do so with the purpose of hitting the goal. True, they can perform many of the bodily movements which would be involved in shooting without having that purpose but these movements will not then constitute shooting at the goal. Now suppose someone comes along and offers me the usual \$1m if I perform the action of shooting at the goal but in a rather unusual way, by deliberately throwing my ball into the stands. I just can't comply with such a request, however much money is on offer. To comply with it I would have to cease intending to reach the goal and then I would no longer be shooting at it.²⁰

Something like this is true of guessing. I don't know my height though I know what sort of height I'm likely to be. Say that someone offers to give me \$1 million provided I guess that I am exactly 7 feet tall. What does it matter to me whether I make this guess or not? The money is more than adequate compensation. I tell myself to go ahead and guess but I can't do it. The explanation is obvious: to guess I must make some attempt to get at the truth and, given what I know, I could not make such an attempt by guessing that I was exactly 7 feet tall.²¹

Of course, I could insincerely hazard an opinion in an effort to give the impression that I had made the required guess. Alternatively, I could take steps to make myself guess, regardless of the evidence, by means of drugs or hypnosis etc. But I can't get myself to make this guess in the way I can get myself to make most guesses, simply by reflecting on the desirability of so doing. To make myself hazard such an absurd opinion, I need to go in for a bit of self-manipulation. I encounter no such obstacles to imagining or supposing that I am 7 foot tall: I can imagine this 'at will' simply because imagining it seems desirable. So the fact that guesses aim at the truth limits our ability to guess at will.

Are there similar limits on how belief can be motivated? And can we explain these limits only by supposing that belief aims at truth in just the sense in which guessing does? Let's grant that it is indeed impossible to acquire a belief simply by reflecting on how desirable

it would be to have that belief and regardless of the evidence for it. That gives us no grounds for thinking that belief is purposive unless we can also get ourselves to form beliefs, as we can get ourselves to make guesses, by reflecting on what the best way to achieve the aim of belief would be. The fact that belief aims at the truth will not explain why knowingly forming certain absurd beliefs is impossible unless our consciousness of the aim of belief can also explain how we get ourselves to form more normal beliefs.

The guesser aims to get it right. He plumps for an answer when he thinks the answer sufficiently likely to be correct, sufficiently likely given his other purposes in making the guess. If he is making the guess – aiming to get the answer right – with the further purpose of winning the quiz then all sorts of considerations, both evidential and non-evidential, will be relevant to whether his guess is reasonable. It is by reflecting on these factors before making a guess that he exercises control over his guessing. By contrast, a believer doesn't shoot at truth once he decides that it is sufficiently likely he'll get it, given his wider purposes in believing.

That is not because a believer has only one purpose: to avoid error. As already noted, this objective could be achieved easily enough by believing nothing at all, by declining to shoot at the target. Believers with limited cognitive resources wish to form a view on large a number of matters and so the formation of a particular belief will be reasonable only if it is part of a general policy of belief formation which is itself reasonable.²² Such policy must be sensitive to various non-evidential considerations. As James points out, believers must strike a balance between the agonies of agnosticism and the risks of error; evidence alone cannot strike this balance for them (James, 1956, pp. 17–19).

If believing were purposive, this sensitivity to non-evidential considerations would involve integrating the goal of having this belief only if it is true with the subject's wider purposes in forming beliefs. And this is something agents do all the time. But this sensitivity to non-evidential reasons for belief is not the sort of thing which is registered efficaciously in the subject's deliberations about what he ought to believe. Our subject can't get himself to believe by reflecting that he has sufficient evidence to form a belief, given his wider purposes in forming beliefs. When we bystanders consider

whether it is reasonable for our subject to form a belief on this matter, we will attend to the importance of the issue, the limited cognitive resources which he has to devote to it and so forth. But reflection on these matters will not move the believer himself to belief, as it moves him to guess (Owens, 2000, Chapter 2). Guessing can be controlled by reflection on how best to achieve an aim, believing cannot.²³

Chasing Penny Blacks is a goal directed activity. You're out to acquire Penny Blacks: you don't want to be taken in by fakes but nor do you want to miss out on the genuine Penny Blacks which cross your path. You must strike a balance between the risk of being duped and the risk of ending up with no Penny Blacks at all. Given that, when offered a stamp you'll run some checks and ask for some documentation but at some point you will be satisfied, you will have sufficient evidence that it is a Penny Black and you'll buy. What we do about an individual stamp will be constrained not just by the goal of buying that stamp only if it is a Penny Black but also by the more general goal of which this is a sub-goal – buying Penny Blacks as such.

From the outside, belief formation seems much the same. You are out to acquire true beliefs. You don't want to be taken in by falsehood but nor do you want to miss out on truths. You must strike a balance between the risk of error and the costs of agnosticism. When evidence suggests that a proposition is true, you may run some checks and wait for a bit more confirmation but at some point you are satisfied, you have sufficient evidence that the belief is true and you'll believe. Isn't that process under your control no less than stamp collecting?

From the inside, the comparison between believing and Penny Black collecting feels much less apt. I can get myself to purchase a stamp by telling myself that I have considered the matter quite long enough and, if I am ever to purchase another stamp, I must now make a decision about whether to buy on the basis of the evidence before me. In the rational Penny Black collector, the subject's pursuit of the goal of acquiring this stamp only if it is a Penny Black is regulated by the subject's reflections on how this goal fits with his wider objectives. But I can't get myself to form a *belief* by rehearsing the same train of thought, by reflecting on some broader

aims, epistemic or non-epistemic. Or, at least, my rationality alone does not guarantee that such reflections will have any influence on my beliefs.

In respect of control, guessing is much more like imagining or supposing than believing. Guessing is a mental action executed for a purpose in way that believing is not. A guess aims at the truth; the guesser has the truth as his goal and can guide his guess towards the truth, taking into consideration his other goals and purposes. This whole process is, in a very obvious sense, an exercise of control. If what I have said about belief formation is correct, nothing like that is true of belief. Belief has a standard of correctness but this standard of correctness does not set a goal towards which the subject can direct his beliefs by asking himself whether he is sufficiently likely to have hit it.

CONCLUSION

Truth may be belief's standard of correctness but, I have argued, there is little explanatory point in describing truth as the *aim* or the *goal* of belief. Those who employ these locutions mean them to do real explanatory work:

To say that belief aims at the truth is not simply to re-express the norm stipulating that a belief must be true in order to be correct; rather, it is to point out a fact about belief that generates this norm for its correctness (Velleman, 2000, pp. 16–17).

Belief's aim is meant to underwrite that norm, to explain not just belief's standard of correctness but also what makes our beliefs reasonable or unreasonable and how we go about controlling our beliefs. I have argued that the truth-aim hypothesis cannot fulfil its early promise: it accounts for neither the authority nor the motivational efficacy of our evidential norms. I suspect that ascribing the truth-aim to belief is indeed just "to re-express the norm stipulating that a belief must be true in order to be correct".

In considering whether belief has an aim, I have examined only the most common view: the truth-aim hypothesis. Are there alternatives? To have a hope of explaining the norms of correctness and rationality governing belief, the aim in question must entail truth and I can think of only one truth-entailing aim other than

truth itself, namely knowledge. Some authors (including myself in Owens, 2000, p. 35) have recently proposed that belief aims at knowledge. Since a guesser aims at truth and not knowledge, this might explain why reasonable belief requires more evidence than a reasonable guess: the goal of a belief is simply more demanding than that of a guess. Do the arguments of this paper work only against the truth-aim hypothesis?

I have no more objection to the idea that knowledge sets the standard of correctness for belief than I have to the view that truth sets this standard but neither hypothesis can do much explanatory work. Plausible accounts of what it is to know that *p* insist that the possession of evidence for *p*, or at least the absence of evidence against *p*, is a necessary condition for knowledge of *p*; this is because it is a necessary condition for reasonable belief in *p* and one can't know that *p* unless one reasonably believes *p*. This raises the question: what level of positive evidence is required for reasonable belief in *p*, or else what level of countervailing evidence is required to undermine reasonable belief in *p*? It is hard to see how this question can be answered, rather than begged, simply by ascribing a knowledge-aim to belief.²⁴

What of the two problems which we hoped to solve by attributing an aim to belief? Are these to fester unsolved? We might instead query the assumptions which generated those problems. Is the authority of instrumental norms really any more secure than that of non-instrumental norms? Can obedience to norms be understood only on the model of an agent pursuing some goal which he finds valuable? Why not allow that different mental states – beliefs, desires, emotions, intentions, etc. – are governed by different norms, only some of which take a teleological form? Perhaps the authority and motivational efficacy of these norms looks problematic only once we try to force them into the mould of the teleological norms which govern action.²⁵

NOTES

¹ (Velleman, 2000, pp. 244–246) gives us a list of three ‘ulterior reasons’ for being interested in the aim of belief which mention explananda corresponding to each of the three just described, though he does not set out to defend the claim that ascribing an aim to belief can explain all three.

² For example (Foley, 2001) recasts all epistemic norms as teleological norms.

³ Most of those who endorse this idea hold that belief has a constitutive aim: truth. But at least one such philosopher (Papineau, 1999) denies that belief has any constitutive aim. Papineau holds that our ‘most significant epistemic norms’ can be explained on the weaker assumption that much of the time we do happen to value true belief.

⁴ A similar line of thought is deployed against internalism about epistemic rationality in (Goldman, 1980). For a response from an internalist who wants to defend a really strong form of instrumentalism about epistemic norms, see (Foley, 1993, 2001).

⁵ Velleman allows that believers may have extremely perverse views about what counts as evidence and yet still be aiming at truth. See (Velleman, 2000, p. 113 n.34).

⁶ The phrase ‘how much’ refers to the probative force of the evidence. In what follows I shall speak indifferently of the *quality* and the *quantity* of available evidence; both phrases should be taken to denote probative force.

⁷ Several authors have noted that practical considerations help to determine how much evidence is required for rational belief (e.g., Nozick, 1993, Chapter 3) and (Wedgwood, 2002, pp. 273–278).

⁸ For more on the distinction between the rationality of inquiry and the rationality of belief, see (Owens, 2000, pp. 18–20, 24–34 and 85–87).

⁹ This is close to formulations offered by (Sosa, 2001, p. 52), (Humberstone, 1992, p. 73) and (Velleman, 2000, p. 251), although it differs markedly from Velleman’s in one respect: it attributes the aim in question to the subject of the belief. Velleman argues that believing is purposive even if the aim of belief is not an aim adopted by the subject of the belief (Velleman, 2000, pp. 184–188, 252–253). At least sometimes, he suggests, it is a belief-forming sub-system in the subject, rather than the subject himself, which pursues the aim of belief. If this were so, it would be hard to see why the normative standing of the belief-forming subsystem should have any implications for the normative standing of the subject. How can the norms of rational belief formation have any authority over me, how can violation of them by my beliefs make *me* irrational, unless I share the aims of the sub-system which give those norms their authority?

Velleman tells us that the subject might choose to endorse or associate himself with the doings of his belief-forming subsystem, in which case he will get egg on his face should the system misbehave (Velleman, 2000, pp. 252–253). But such endorsements are contingent and even voluntary. That’s not how it is between me and my beliefs. If I have an irrational belief, I am (to that extent) irrational regardless of the endorsements I have made or withheld. In so far as I have beliefs, I am always subject to the authority of epistemic norms; that is the fact we have to explain.

¹⁰ Railton observes that a belief does not become reasonable simply because adopting that belief would have the effect of making you acquire lots of other true beliefs and he concludes from this that “if belief is to be said to ‘aim at’ truth, then this would appear not to be ‘aiming at’ in the familiar, teleological sense in

which a goal is regulative of action” (1994, p. 74). But all that really follows from this observation alone is that each belief must be adopted with the aim of adopting that particular belief only if it is true.

¹¹ Railton says that “It is *distinctive* and constitutive of belief not only that it represents its content as true but that it takes itself to be correct only if that particular content really is true” (my italics) (Railton, 1994, p. 74). If a psychological state of the sort I take guessing to be is even possible then Railton is wrong about this. (Williamson, 2000, p. 244) makes an analogous observation about assertion and conjecture.

¹² (Williams, 1978, p. 39) notes this disanalogy between seeking true beliefs and collecting pre-historic flints.

¹³ (Cohen, 1974, p. 194) argues that the English word ‘guess’ is in fact ambiguous, having both a speech-act and a mental-act sense. He may well be right about this.

¹⁴ I am grateful to Mark Kaplan for this example.

¹⁵ (Nozick, 1993, pp. 85–100) considers a quasi-instrumental model of belief rationality which makes believing look rather like guessing and he notes some difficulties for it.

¹⁶ (Williamson, 2000, pp. 244–249) presents a different but related argument for the conclusion that we can’t explain the amount of evidence required for assertion by supposing that assertion aims at the truth.

¹⁷ For a good discussion of the formal aspects of epistemic utility theory, see (Kaplan, 1981).

¹⁸ Would things be any easier if we ascribed the aim to a belief-forming subsystem, rather than to the believer himself? (cf. note 9) I can’t see how. The system would need to be sensitive to a whole range of non-evidential considerations for it to produce reasonable beliefs and the same problem of integration would arise.

¹⁹ Where a subject’s wider purposes would be served by having an irrational belief, it might be rational for them to induce that belief. But what they are doing here is performing an action which, though it may be a rational action, has the effect of making them irrational. And an irrational subject is someone who is violating norms which have some rational authority over them. Our problem is to explain that authority.

²⁰ Compare Pink’s discussion of what he calls ‘special-purpose agency’ (Pink, 1996, pp. 248–252).

²¹ Might guessing what you know to be false be impossible not because guessing aims at the truth but rather because guessing requires ignorance of the truth? As Alexander Bird pointed out, it is no easier for me to *guess* that I am more than six inches tall than it is for me to guess that I am more than 7 ft. tall.

This is certainly the right thing to say about states like doubting or wondering. The correct explanation of why we can’t doubt that or wonder whether *p* where we know *p* is false is that doubting or wondering requires ignorance of *p*’s truth. Neither state aims at the truth in the sense I defined. But consider fantasizing that *p*. Here too we can’t fantasize that *p* when we know *p* to be true but the reason isn’t that we must be ignorant of whether *p* is true; on the contrary, we can fantasize

that p when we know p to be false. And this is because of the distinctive aim of fantasy: it aims to provide pleasure by entertaining thoughts which we don't know to be true, or something like that.

I would argue that we should look to guessing's aim in order to explain why we can't guess something we know to be true. I suggest that we can't guess something we know to be true because guessing aims at truths of which we are ignorant. And that also explains why we can't guess something we know to be false. So the orientation which guessing has towards truth in virtue of its aim is indeed essential to any account of the limits on what you can guess.

²² Velleman maintains that "one's acceptance of a proposition can amount to a belief without being part of any global epistemological project of accumulating true beliefs. What distinguishes believing a proposition from imagining or supposing it is a more narrow and immediate aim – the aim of getting the truth value of that particular proposition right, by regarding the proposition as true only if it really is" (Velleman, 2000, p. 252). One can agree with Velleman that this is the sole purpose essential to the acquisition of a particular belief, the only purpose one must have to qualify as a believer in p , whilst also insisting that all beliefs must be acquired with some other purposes in mind, purposes which help to determine whether the belief in question is reasonable.

²³ Of course, I don't deny that our actions can affect our beliefs but, as I have argued elsewhere, such indirect control over belief does not help us understand the motivational role of epistemic norms (Owens, 2000, pp. 82–84).

²⁴ For further discussion of the truth aim and the knowledge aim see (Williams, 1978, pp. 38–45; Williamson, 2000, p. 208; Wedgwood, 2002, pp. 289–291).

²⁵ Many thanks to Jonathan Adler, Alexander Bird, Paul Horwich, Barry Loewer, Peter Milne, Mark Kaplan and, especially, Paul Noordhof for comments on earlier drafts of this paper.

REFERENCES

- Cohen, J. (1974): 'Guessing', *Proceedings of the Aristotelian Society* 74, 189–210.
- Foley, R. (1993): *Working Without a Net*, Oxford: Oxford University Press.
- Foley, R. (2001): 'The Foundational Role of Epistemology in General Theory of Rationality', in Fairweather and Zagzebski (eds.), *Virtue Epistemology*, Oxford: Oxford University Press.
- Goldman, A. (1980): 'The Internalist Conception of Justification', in French, Uehling and Wettstein (eds.), *Midwest Studies in Philosophy V*, Minneapolis: University of Minnesota.
- Harman, G. (1999): *Reasoning, Meaning and Mind*, Oxford: Oxford University Press.
- Humberstone, L. (1992): 'Direction of Fit', *Mind* 101, 59–83.
- James, W. (1956): *The Will to Believe*, New York: Dover.

- Kaplan, M. (1981): 'A Bayesian Theory of Rational Acceptance', *Journal of Philosophy* 78, 305–330.
- Nozick, R. (1993): *The Nature of Rationality*, Princeton: Princeton University Press.
- Owens, D. (2000): *Reason Without Freedom*, London: Routledge.
- Papineau, D. (1999): 'Normativity and Judgement', in *Proceedings of the Aristotelian Society Supplementary Volume LXXIII* (pp. 17–43).
- Pink, T. (1996): *The Psychology of Freedom*, Cambridge: Cambridge University Press.
- Railton, P. (1994): 'Truth, Reason and the Regulation of Belief', *Philosophical Issues* 5, 71–93.
- Railton, P. (1997): 'On the Hypothetical and non-Hypothetical in Reasoning About Belief and Action', in Cullity and Gaut (eds.), *Ethics and Practical Reason*, Oxford: Oxford University Press.
- Sosa, E. (2001): 'For the Love of Truth?' in Fairweather and Zagzebski (eds.), *Virtue Epistemology*, Oxford: Oxford University Press.
- Velleman, D. (2000): *The Possibility of Practical Reason*, Oxford: Oxford University Press.
- Wedgwood, R. (2002): 'The Aim of Belief', *Philosophical Perspectives* 16, Language and Mind, 267–297.
- Williams, B. (1973): 'Deciding to Believe', in his *Problems of the Self*, Cambridge: Cambridge University Press.
- Williams, B. (1978): *Descartes: The Project of Pure Inquiry*, London: Penguin.
- Williamson, T. (2000): *Knowledge and its Limits*, Oxford: Oxford University Press.

Department of Philosophy
University of Sheffield
Western Bank
Sheffield, S10 2TN
UK
E-mail: d.owens@sheffield.ac.uk

