

Truth: the Aim and Norm of Belief*

Daniel Whiting

All of it was true. The truth and nothing but the truth. But not quite all the truth. What I left out was my business. (Raymond Chandler, *The Long Good-Bye*)

The final version of this paper will appear in a special issue of *Teorema* on 'The Aim of Belief', guest-edited by José Zalabardo. When citing, please refer to the published version.

1 Norms

It is correct to believe that Steinbeck wrote *East of Eden*, incorrect to believe that Updike wrote *Cannery Row*. It is right to believe that tomatoes are a fruit, wrong to believe that apples are a vegetable. More generally, true beliefs are correct or right, false beliefs are incorrect or wrong. This suggests that a norm of truth governs the state or attitude of belief. Since it is wrong or incorrect to do what one may not do, one might try to capture this suggestion with the following principle:¹

(TRUTH) One may believe that p if and only if it is true that p .^{2,3}

Other candidate norms for belief include:

(RATIONALITY) One may believe that p if and only if it is rational to believe that p .⁴

* Thanks to Tsung-Hsing Ho and Javier González de Prado Salas for comments on this paper, and to Conor McHugh for many discussions of the issues it concerns.

¹ Some deny that belief is governed by a general normative principle of any sort (see Papineau Forthcoming). Such views are not my present concern.

² I advance the claim that (TRUTH) governs belief in Whiting 2010. Though they differ in how they formulate the norm, Boghossian (2008), Gibbard (2005), Littlejohn (2012), Millar (2009), Shah and Velleman (2005), and Wedgwood (2002) defend similar claims.

³ Note that 'may' here and in all the norms which follow has narrow scope.

(KNOWLEDGE) One may believe that p if and only if one knows that p .⁵

It is an open question to what extent these principles are in competition. A common proposal is that a standard like (RATIONALITY) is in some way derived from or generated by the more fundamental (TRUTH) (cf. Boghossian 2008: 101). A similar though less common proposal is that (KNOWLEDGE) too is a consequence of (TRUTH) (cf. Wedgwood 2002). A less compromising position would be to reject one or both of (RATIONALITY) and (KNOWLEDGE), at least as stated.⁶ In this paper, I shall not pursue such matters. Instead, I shall focus on defending the formulation I have given of the norm of truth governing belief against various objections.

Other than to accommodate examples like those above, why accept (TRUTH)? There are a number of considerations which support doing so; for the purposes of this paper, I shall highlight the following. First, (TRUTH) promises to explain what is problematic about holding so-called Moorean⁷ beliefs, such as *that I believe that dogs bark but dogs do not bark*, a belief which is consistent and (so) might be true. Very roughly, the explanation proceeds as follows. If I were to have such a belief, I would take myself to hold an attitude governed by a certain norm while at the same time taking that attitude not to satisfy that norm (cf. Millar 2009).

Second, if (TRUTH) holds, that might go some way to explaining why subjects take only *evidence* to provide *reason* or *justification* for believing.⁸ Only evidence that p indicates that, were a subject to believe that p , she would satisfy the standard to which that attitude is subject. By the same token, if (TRUTH) holds, that might account for other general principles governing belief, such as that one ought not to believe something when the evidence indicates

⁴ Feldman (2000) defends a version of (RATIONALITY).

⁵ Williamson (2000) defends a version of (KNOWLEDGE).

⁶ For critical discussion of the view that knowledge is the norm for belief, see Whiting Forthcoming-a.

⁷ After Moore 1962: 277ff.

⁸ As I use the term here, a *justification* is a consideration which stands to what one may do as a *reason* stands to what one ought to do. So, while a reason for ϕ ing speaks in favour of ϕ ing, a justification might not speak in favour of ϕ ing, though it does not speak against it.

otherwise. One should not believe against one's evidence since, in doing so, one might violate (TRUTH).⁹

Third, by appeal to (TRUTH), one might explain why a subject is *motivated* (not) to believe a proposition when she takes herself to have evidence for (or against) it (cf. Shah and Velleman 2005). If a subject accepts (TRUTH), she is moved by considerations which seem to show that having a given belief would (not) accord with that norm, which considerations constitute evidence.

I shall return to each of these points in due course. (TRUTH) might be explanatorily significant in a further respect. Some suggest that it is *constitutive* of an attitude's being one of belief that it is subject to a norm of truth. I shall assume that, if the norm holds, it does so necessarily but I shall not take a stand on what this might tell us about the nature of belief. That is a topic for another occasion.

2 Aims

In addition to talk of a *norm* which governs belief, it is common to hear talk of an *aim* which governs belief. Some think that talk of belief's aim can only be metaphorical, more specifically, that it is best understood as a figurative expression of (TRUTH):

It is often claimed that belief aims at truth. [...] But this claim is obviously not literally true. Beliefs are not little archers armed with little bows and arrows: they do not literally 'aim' at anything. This claim must be interrupted as a metaphor.

I propose to interpret this claim as a *normative* claim—roughly, as the claim that a belief is correct if and only if the proposition believed is true. (Wedgwood 2002: 267)

⁹ This is an example of the attempt to derive a version of (RATIONALITY) from (TRUTH).

Others take the claim at face value (see Velleman 2000). Belief aims (only) at truth, they suggest, in the sense that it is part of the *function* of a state or attitude of belief to be true, where its having that function consists in its being causally regulated in ways designed to ensure that it is true.

It is common for those who understand belief's aim in this way to advance the (teleological) proposal that the norms governing belief, like (TRUTH), are *derived* from its aim. A false belief is wrong or incorrect, on this view, because it fails to fulfil its function. As Velleman says:

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. (2000: 16-17)

There is a way to take talk of aims governing belief seriously without taking it to concern the functional role of belief-states or the sub-personal mechanisms regulating them—one can take it to concern an aim which the *subject* has with respect to believing. It is legitimate to say that a newspaper article aims to shock but what that comes to is that the person writing the article does so with the aim of shocking. In a similar fashion, it is legitimate to say that a belief aims only at the truth but what that comes to is that a person aims to believe something only if it is true.¹⁰

Some are sceptical that beliefs in general are governed by an aim which the subject possesses:

Only some instances of belief are caused by the goal-directed activity of their subjects; many others are the product of processes such as perception, which don't involve any agential goals

¹⁰ For discussion and defence of this formulation of the aim, see Whiting 2012.

or intentions. If the metaphor that belief aims at the truth is not to rule out most cases of belief, it will have to draw on a wider notion of truth-directedness, encompassing non-agential mechanisms that track the truth. (Shah and Velleman 2005: 498-499)

This seems confused. In general, the fact that a subject has a certain aim has consequences for how she is inclined or motivated to act, and for what reason or justification she takes (what appear to her to be) the facts to provide (cf. Alvarez 2010: ch. 3). If I aim to lose weight, it does not follow that when I lose weight I do so intentionally or as a (direct or indirect) result of aiming to do so; the weight-loss might be due to an overactive thyroid. Nonetheless, given my aim to lose weight, I am inclined to take exercise, motivated not to consume protein and carbohydrates at the same time, and take the fact that the shake is high in calories to be a reason not to drink it. In a similar way, if I aim to believe only the truth, it does not follow that when I believe a proposition I do so intentionally or as a (direct or indirect) result of having that aim; rather, if given that aim, I am inclined not to pay attention to unreliable witnesses, motivated not to believe that Smith killed Jones having learned that Smith was on holiday at the time of the shooting, and take the fact that it is snowing to justify believing that it is cold outside. Even if a certain belief, say, that the kettle is boiling, results from perception, it remains the case that I would be moved to revise that belief were I to take it to clash with other things which I know, say, that the kettle is broken.

These remarks in turn point to something which speaks in favour of attributing the aim which governs belief to a *subject*. It is a subject, not a *state* of the subject, who takes (what appears to her to be) a certain fact to provide reason or justification for believing something, and who is motivated not to believe certain things when presented with counterevidence. The suggestion that a subject aims to believe only the truth promises to provide a straightforward explanation of this in a way which the suggestion that belief-states

are causally regulated by certain sub-personal mechanisms seems not to (which, of course, is not to deny that belief-states are so regulated).

As the preceding discussion shows, I do not think that talk of aims with respect to believing is merely a dark but suggestive way of talking about the norms of belief. But nor do I accept the teleological proposal that from the aim governing belief one can *derive* the norms to which that attitude is subject. The fact that someone aims to do something does not generally entail that she has reason to do that thing (unless there is reason to have that aim).¹¹ That Augustus aims to gorge himself on chocolate does not entail that Augustus has reason to gorge himself on chocolate, that he should do so, or that his doing so would be correct. However, if Augustus takes himself to have reason to gorge himself on chocolate, or if he accepts a norm according to which he should do so, he might as a result aim to gorge himself on chocolate. In a similar fashion, I suggest, the fact that a subject aims to believe only the truth is a *consequence* of her (perhaps tacit) acceptance of (TRUTH).

3 Ought and may

According to (TRUTH), if a proposition is false, one ought *not* to believe it but, if a proposition is true, it is only the case that one *may* believe it, not that one *ought* to do so.

Why, one might ask, formulate the norm in this way? Why not think, with Gibbard (2005) and Horwich (1998: 187), among others, that one *should* believe a proposition when it is true? Indeed, Olinder (2012: §3) objects to my formulation of the norm governing belief precisely on the grounds that it does not do justice to some version of this thought.

The claim that, if a proposition is true, one ought to believe it clashes with the principle that *ought implies can* (cf. Boghossian 2008: 100; Bykvist and Hattiangadi 2007).

¹¹ The standard response to this point is to appeal to the suggestion that the relevant aim is *constitutive* of belief, in the sense that an attitude or state is one of belief only if governed by that aim. Without denying the constitutive claim, I doubt this answers the objection. It might be true that having a certain aim is constitutive of ϕ ing and also true that that aim is not worth having (in which case, it is not worth ϕ ing). For trenchant criticism of the idea that *practical* norms might be grounded in what is constitutive of *agency*, see Enoch 2006.

There are truths which are so complex that it is impossible to believe them, for example, the truth consisting of the conjunction of all the truths about the dimensions and mass of each grain of sand on each beach in Cornwall. Since one ought to ϕ only if it is possible for one to ϕ , it follows that it is not the case that one ought to believe that truth.¹²

One might think that (TRUTH) faces a similar problem, since, according to it, I may believe the (unbelievable) complex truth about grains of sand. However, as I have argued elsewhere (Whiting 2010: 216-217), there is no principle that *may implies can*. To say that one may ϕ is to say that it is not the case that one ought not to ϕ . Clearly, it does not have to be possible for one to ϕ for it not to be the case that one ought not to ϕ .

As a variation on the theme, suppose that Mulder is simply unable *not* to believe that there are aliens—he cannot but believe this. Suppose also that there are no aliens. According to (TRUTH), Mulder ought not to believe that there are aliens. Surely, the critic of (TRUTH) might say, this conflicts with the principle that *ought (not) implies can (not)*.

The case of Mulder is one in which it is *psychologically* impossible for a subject not to believe something. In contrast, it is *humanly* impossible for a subject to believe the conjunction of all the truths about the grains of sand (or, for that matter, all of those truths individually). Standardly, proponents of the principle that *ought implies can* take the relevant possibility to be weaker than psychological possibility. If Scrooge is psychologically unable to raise staff salaries, this hardly releases him from the obligation to do so. In contrast, since it is not humanly possible for Scrooge to travel to Christmas past, it cannot be the case that he should do so.

A critic of (TRUTH) might persevere. Perhaps there are falsehoods which it is not humanly possible not to believe. If so, it cannot be the case that one ought not to believe them.

¹² There are various ways in which the proponent of the idea that truth is governed by a norm according to which subjects *ought* to believe what is true might try to reformulate the norm so as to avoid this problem. Bykvist and Hattiangadi (2007) anticipate and criticise (convincingly, in my view) many such reformulations.

I do not think we need to take this line of thought seriously. Regarding truths which it is humanly impossible to believe, we know that they *actually* exist. But the critic of (TRUTH) has only appealed to the *mere possibility* of falsehoods which it is humanly impossible not to believe. She has not shown that such cases are genuine, or even that they are genuinely possible. Indeed, if there were an attitude toward a proposition which a human could not but have, it is not clear that it would count as a *belief*, properly so-called, insensitive as it is to evidence, isolated as it is from our practices of belief revision, insulated as it is from the processes through which we rationally update our belief-system, and so on.

Pursuing these matters would take us far afield—the principle that *ought implies can* is controversial and even those who accept it disagree over how to interpret it. Fortunately, there are considerations which have nothing to do with that principle which count against building into (TRUTH) the idea that one ought to believe what is true.

It is a platitude that it is wrong not to do what one ought to do.¹³ If Elliot ought to apologise to Stanley, it is wrong of him not to do so. It is also a platitude that it is wrong to believe a falsehood. If Holly believes that Spielberg directed *Days of Heaven*, she is wrong. In contrast, it is *not* wrong *not* to believe a truth. It is not wrong for Holly not to believe that Ozu directed *Late Spring*. Since it is not wrong not to believe a truth, and since it is wrong not to do what one ought to do, then it is not the case that one ought to believe the truth.

Perhaps there are situations in which it *is* wrong not to believe a truth, and so in which one ought to do so. It might be wrong for a driving instructor in the UK not to believe that vehicles there drive on the left—she really should know this. But in this case it is not the mere truth of the relevant proposition that makes it the case that the instructor ought to

¹³ This claim requires qualification. Some distinguish what one ought to do relative to the facts, relative to the evidence, relative to one's beliefs, and so on. One might then distinguish corresponding respects in which something can be wrong. The claim would then be that it is wrong in such-and-such respect if one does not do what one ought relative-to-such-and-such to do.

believe it but additional considerations relating in a fairly obvious fashion to her role and responsibilities.

This asymmetry between truth and falsity—the fact that there is reason not to believe falsehoods in a way that there is not reason to believe truths—is interesting and, perhaps, surprising. I do not pretend to have a satisfactory explanation for it but here is a thought. If one has false beliefs, it follows that to that extent one is out of touch with reality, one does not have a grip on it. If one has true beliefs, it does *not* follow that to that extent one is in touch with reality or that one has a grip on it; after all, one's beliefs might be irrational, or unjustified, or accidentally true, or.... Consider a subject who believes that there is a barn in the field but does not *know* that there is a barn in the field. Even if what she believes is true, she does not believe what she does in light of the fact that there is a barn in the field, and so in turn she cannot act or hold other attitudes in light of that fact.¹⁴ It seems, then, that false beliefs (as such) are detrimental in a way that true beliefs (as such) are not beneficial. This might go some way to helping us to make sense of the asymmetry.¹⁵

Interestingly, one finds a similar asymmetry with respect to things which are analogous to beliefs: maps. If a London Transport map represents Waterloo as being on the Victoria Line, it is a bad map. A cartographer ought to revise the map, and a tourist or commuter has reason not to use it. In contrast, the actual London Transport map is not a bad map because it does not represent certain details of the area it covers, such as the geographic locations of the stations, or because there are areas it does not cover, such as Chicago; on the contrary, it is widely acknowledged to be, for these very reasons, a good map. It is not the case that a cartographer ought to revise it, and a tourist or commuter has no reason not to use it.

¹⁴ For more on the relationship between knowing that *p* and being able to act in light of the fact that *p*, see Hyman 1999.

¹⁵ For more on this theme, see Whiting Forthcoming-b.

So, generally-speaking, a map should not be inaccurate but it need not be comprehensive; everything that a map covers it ought not to cover unfaithfully while it is not the case that a map ought to cover everything faithfully. Presumably, what details a map should include, or what area it should encompass, depends in appreciable but difficult to spell out ways on the needs and interests it is to serve.

Returning to the issue at hand, if, as Ramsey suggests (1990: 146), a belief or belief-system ‘is a map of neighbouring space by which we steer’, it is to be expected that those beliefs one has about that space ought not to be false but not that one ought to have true beliefs about everything in that space, let alone of everything outside of it.¹⁶

4 Moorean propositions

I suggested above that appealing to (TRUTH) might explain what is problematic about believing certain Moorean propositions.¹⁷ One might think that, on the contrary, such propositions create difficulties for (TRUTH). Consider:

(MOORE) Dogs bark but I do not believe that dogs bark.

Suppose that (MOORE) is true. According to (TRUTH), I may believe it. However, were I to do so, (MOORE) would be false.¹⁸ So, I would then have a belief which, according to (TRUTH), I should not have.¹⁹

As I admit in an earlier paper, ‘there is [...] something rather fishy about the notion of a permission which, when acted upon, results in one’s doing something one is not permitted to do’ (2010: 218). In the same paper, I propose amending (TRUTH) so as to exclude cases

¹⁶ This point depends only on the thought that beliefs are analogous to maps, not that they *are* maps. For discussion and defence of a ‘map theory’ of belief, see Braddon-Mitchell and Jackson 2007: ch. 11.

¹⁷ One might think that the appeal to (TRUTH) alone cannot explain what is problematic about believing a Moorean proposition such as *that it is raining but I do not know that it is raining*. For discussion, see Whiting Forthcoming-a.

¹⁸ The second conjunct would be false on the assumption that, if I believe the conjunction, I believe each of its conjuncts. There is a longer way of telling the story which does not depend on that assumption.

¹⁹ Bykvist and Hattiangadi (2007) are responsible for this line of thought, though they do not discuss the formulation of the norm of truth I provide here.

like (MOORE) by restricting it to those propositions one can truly believe. But I now recognise that this looks *ad hoc* (cf. McHugh 2012: n34).

Sorensen (1988) calls a proposition like (MOORE) a ‘blindspot’ for belief. It is a blindspot in the sense that the proposition, though consistent, is unavailable to belief; it is not a proposition toward which one can adopt that attitude. Very roughly, Sorensen’s explanation for why Moorean propositions are blindspots for belief proceeds as follows. Subjects have the ‘doxastic goals of getting truth and avoiding error’ (1988: 37). These goals impose ‘constraints’ on believing, constraints which could not be satisfied if (*per impossibile*) the object of belief were to be a proposition like (MOORE). Were a subject to believe (MOORE), she would be ‘committed’ (in an intuitive sense) to both believing that dogs bark and to not believing that dogs bark. Were the subject to follow through on her commitments, she would have inconsistent beliefs, and hence would fail to satisfy the relevant doxastic goals. Thus, ‘Given the constraints imposed by certain desiderata of belief, I cannot believe [(MOORE)] even though it is a consistent proposition’ (1988: 53).

Insofar as this story appeals to goals—akin to the aim to believe (only) what is true—and resultant constraints—or norms—which govern believing, it is congenial to the proponent of (TRUTH). By appeal to it, she might respond to the above objection. Since a proposition like (MOORE) is unbelievable, there will be no situation in which a subject believes what, according to (TRUTH), she may believe, only to find herself believing what, according to (TRUTH), she should not.

One might worry about a norm which states that one may do something one cannot do (cf. McHugh 2012: n33). But, as discussed earlier, there is no principle that *may implies can*.

I have a lot of sympathy for the view that Moorean propositions are blindspots, so conceived. But I do not know how to defend or bolster the view and, in my experience, it is not widely shared. Fortunately, I do not think it is necessary to pursue the issue for present

purposes. Even if it is possible to believe the relevant Moorean proposition, the thought that this poses a problem for (TRUTH) involves a misunderstanding. According to (TRUTH), a situation in which (MOORE) is true is a situation in which I may believe it. But *any* such situation is one in which I precisely do *not* believe (MOORE), since any situation in which I *do* believe it is *not* a situation in which it is true. Hence, the situation which supposedly presents a counterexample to the principle is not genuinely possible.

One might respond that the relevant situation is supposed to be one in which (MOORE) is true *at a certain time*. Obviously, in that situation, I do not in fact believe (MOORE) at that time but, according to (TRUTH), I may do so. If I were to proceed in light of this verdict to believe (MOORE) at some later time, I would then believe a falsehood, which, according to (TRUTH), I should not believe.²⁰

Bringing in temporal indices in this way only makes the problem with the objection from Moorean propositions more apparent. According to (TRUTH), as originally formulated, if a proposition is true, one may believe it. If one adds a temporal index to the antecedent of the relevant conditional, then surely one must also add the *same* temporal index to the consequent, as follows: if a proposition is true at *t*, one may believe it at *t*. Once again, then, the situation which is supposed to present a problem for (TRUTH) is not a genuine possibility. In a situation in which (MOORE) is true at a time, then, according to (TRUTH), I may believe that proposition at that time. But any such situation is one in which I do *not* believe (MOORE) at that time.

So, there is no possible situation in which a subject believes what, according to (TRUTH), she may believe only to find herself in a situation in which she believes what, according to (TRUTH), she may not believe. Again, one might worry about a norm which

²⁰ One might worry about the idea that a proposition has a truth-value at a certain time, at least insofar as that invites the idea that it might have a different truth-value at a different time. I share this worry but set it aside for present purposes.

states that one may do something one cannot do but, again, there is no principle that *may implies can*.

5 Evidential norms

McHugh complains about the version of the norm governing belief I advance on the grounds that it ‘was supposed to explain other epistemic norms as being derivative from it’:²¹

[(TRUTH)] will not explain why you *should* tend to believe [...] propositions for which you *have* sufficient evidence, even utterly compelling evidence (if you form any attitude to them). At most, they will entail that this is not something you shouldn’t do. [...] So, on this account, the epistemic normative landscape becomes rather sparse. (2012: 18)

Relatedly, McHugh complains that (TRUTH) cannot explain why evidence that *p* is a reason to believe that *p*, that is, a consideration which favours of doing so, only why evidence against *p* is reason not to believe that *p* (2012: 18-19).

I agree that, if (TRUTH) holds, one would expect there to be a derived norm according to which one may believe that *p* if and only if one has sufficient evidence that *p*, but not one according to which one should believe that *p* if one has sufficient evidence that *p* (cf. Whiting 2010: §6).²² I also agree that, if (TRUTH) holds, evidence that *p* is at most justification for believing that *p*, not a reason for so believing.²³ Unlike McHugh, I take all of this to be a virtue of the formulation I recommend, not a vice. The epistemic normative landscape is indeed a sparse one.²⁴

²¹ McHugh (2012: 17-18) also criticises (TRUTH) by appeal to considerations relating to *withholding belief* or *suspending judgement*. The nature of this attitude and of the norms governing it are large issues, which I shall have to address on another occasion.

²² I shall not here tackle the issue of what it takes for evidence to count as sufficient, or of how to reconcile or adjudicate the verdicts the evidential norm and (TRUTH) deliver where they appear to diverge.

²³ Cf. n8 above.

²⁴ I defend this claim in Whiting 2012: §5. See also Littlejohn 2012: 46-48; Nelson 2010.

Suppose that I have evidence that the cakes are burning. Whatever provides this evidence provides evidence for an infinite number of other beliefs, such as that there are cakes, that the cakes are burning or that Tolstoy wrote *Great Expectations*, that if the cakes are burning then the cakes are burning, that there is more carbon dioxide in the atmosphere than there was before I made the cakes, that there is no dancing monkey singing the national anthem in the space the cakes occupy, and so on without end. Given the principle that *ought implies can*, it is not the case that I should believe each and every one of these propositions.

As noted earlier, the principle that *ought implies can* is controversial. Fortunately, I do not need to rely on it. It is just implausible to suggest that, for each and every proposition which the evidence supports, one ought to believe it. Surely, a norm which asks this of a subject is absurdly demanding.

If someone were to continue to insist that there is an evidential norm according to which one ought to follow one's evidence, there is a danger of stalemate. To avoid this, consider a point McHugh makes to further his case:

It might be suggested that the epistemic normative landscape is sparse in this way. But that would not tally with our ordinary epistemic assessments. (2012: 18)

On the contrary. If I were to believe that the Liberal Democrats will win the next election in the face of overwhelming evidence that they will lose, or simply in the absence of any evidence that they will win, I might be stupid, daft, irrational, stubborn, or rash. It would not be out of the ordinary to criticise me in one of these ways. But it would be out of the ordinary to criticise me were I simply not to believe all and any propositions which the evidence supports. After all, right now there are many such propositions which I do not believe. If I am stupid, daft, etc., it is not for that reason.

McHugh might reply that I have overlooked the parenthetical remark: 'if you form any attitude to them'. The idea is not that, for each proposition one's evidence supports, one

ought to believe it but that, if one adopts some attitude toward a proposition one's evidence supports, it ought to be an attitude of belief.

If evidence that a proposition is true on its own provides no reason for believing that proposition, and so could not make it the case that one should do so, as the parenthetical remark seems to grant, it is hard to see how the mere fact that one *is* adopting an attitude toward the proposition could supply the outstanding reason so to believe or make it the case that one should. That looks like bootstrapping.²⁵

Of course, if one has a reason to adopt some attitude toward the relevant proposition, then perhaps one should adopt belief, given that the evidence supports it. A teacher might have reason to have a view about the aptitude of her pupil and so, if the evidence shows that the pupil lacks talent, she ought to believe this. But, of course, one should not expect the norm which governs believing on its own to explain why, in such a case, the subject has reason or ought to believe what the evidence supports, since that turns on non-epistemic or non-alethic considerations.

6 Motivation

McHugh makes another complaint concerning the explanatory adequacy of (TRUTH). The principle (or, I would add, our acceptance of it) 'will explain why we are motivated not to believe propositions for which we lack evidence', but not 'why we are ever motivated to believe anything', at least when we have no existing interest in or concern for the subject matter (2012: 19). As discussed earlier, if one accepts (TRUTH), one aims to believe a proposition *only if* it is true, but one does not aim to believe a proposition *if* it is true.

Where a subject is motivated to ϕ , she is typically motivated by certain considerations, that is, by (what she takes) to be reasons for ϕ ing; and, if a subject takes there to be a reason

²⁵ In something like Bratman's sense (1987: 24-27).

for her to ϕ , she is typically motivated to ϕ (cf. Alvarez 2010: 56-59). Since I deny that evidence provides a reason for believing, it is no surprise that I deny that subjects are motivated to believe by evidence alone. However, it is entirely consistent with the view I recommend that subjects often are motivated to believe, and motivated by their evidence, since it is consistent with that view that subjects often (take themselves to) have independent reasons for having beliefs about certain subject matters, and hence aim to do so. It is also consistent with the view that subjects are often motivated to believe out of natural curiosity and the like. Still, it is worth asking what I might say about those cases where such an aim or motivational state is lacking.

Suppose George opens his eyes, sees that he has two hands, and as a result forms the belief that he has two hands. No doubt his perceptual experience is in some way responsible for his belief but I am not sure what the grounds are for insisting that what George saw *motivated* his belief. Perhaps the worry is that to deny this is to deny that George believes what he does *for* a reason. But one can still say that George believes *with* or *in light* of the justification for doing so, namely, the fact that he has two hands, and that he does so thanks to the successful exercise of a perceptual capacity.²⁶ So, to deny that George's belief is motivated is not to say that it is formed blindly.

Suppose, however, that one asks George *why* he believes that he has two hands. It would be entirely unsurprising if he were to answer: Because I saw them. If George were here citing the reason for which he believes, a reason provided by the evidence made available by his perceptual experience, and if taking oneself to have such a reason involves being motivated so to believe, then it might seem I am committed to claiming that George's answer is false (if not insincere).

²⁶ It is telling that it is odd to speak of doing something *for* a justification.

In giving his answer, I suggest, George is citing an *explanatory* reason for his belief—that he saw his hands is a reason *why* George believes that he has two hands, but it need not be a reason *for which* he believes.²⁷ Moreover, in giving this explanation, George might be indicating that he has justification for his belief, insofar as seeing that one has two hands is a very good way of securing justification for believing that one does. So, (TRUTH) notwithstanding, it is not difficult to make sense of both the question George is asked and his answer to it.

Of course, one should not put too much weight on one example. The point is only that it is not a given that subjects are motivated to believe propositions (simply) by the evidence they take themselves to have, and so it is something one can appeal to, at least not without further ado, in criticising (TRUTH).

7 Conclusion

In this paper, I have defended my formulation of the norm of truth which governs belief. On the assumption that subjects in some sense recognise and accept that norm, I have thereby defended the idea that subjects aim to believe only what is true. A recurring theme in the discussion has been whether a norm which does not incorporate some version of the thought that one ought to believe what is true will do the business, that is, whether it will explain the ways in which we are moved in forming and revising beliefs and the various dimensions along which those beliefs are evaluated. I have suggested that it does, in large part by suggesting that the number of those movements and dimensions is less than it might seem.

References

Alvarez, M. 2010: *Kinds of Reasons*. Oxford: Oxford University Press.

²⁷ For discussion of the difference between explanatory reasons and motivating reasons, see Alvarez 2010: ch. 2.

- Boghossian, P. 2008: *Content and Justification*. Oxford: Oxford University Press.
- Braddon-Mitchell, D. and Jackson, F. 2007: *Philosophy of Mind and Cognition*, 2nd ed.
Oxford: Blackwell.
- Bratman, M. 1987: *Intention, Plans and Practical Reason*. Cambridge, Massachusetts:
Harvard University Press.
- Bykvist, K. and Hattiangadi, A. 2007: Does Thought Imply Ought? *Analysis* 67: 277-85.
- Enoch, D. 2006: Agency, Shmagency: Why Normativity Won't Come from What Is
Constitutive of Action. *Philosophical Review* 115: 169-198.
- Feldman, R. 2000: The Ethics of Belief. *Philosophy and Phenomenological Research* 60:
667-695.
- Gibbard, A. 2005: Truth and Correct Belief. *Philosophical Issues* 15: 338-350.
- Hyman, J. 1999: How Knowledge Works. *Philosophical Quarterly* 49: 433-451.
- Littlejohn, C. 2012: *Justification and the Truth-Connection*. Cambridge: Cambridge
University Press.
- McHugh, C. 2012: The Truth Norm of Belief. *Pacific Philosophical Quarterly* 93: 8-30.
- Millar, A. 2009: How Reasons for Action Differ from Reasons for Belief. *Spheres of Reason*,
ed. S. Robertson. Oxford: Oxford University Press.
- Moore, G. E. 1962: *Commonplace Book: 1919-1953*. London: Allen & Unwin.
- Nelson, M. 2010: We Have No Positive Epistemic Duties. *Mind* 119: 83-102.
- Olinder, R. F. 2012: Rescuing Doxastic Normativism. *Theoria* 4: 293-308.
- Papineau, D. Forthcoming: There are No Norms of Belief. *The Aim of Belief*, ed. T. Chan.
Oxford: Oxford University Press.
- Ramsey, F. P. 1990: *Philosophical Papers*, ed. D. H. Mellor. Cambridge: Cambridge
University Press.

- Shah, N. and Velleman, J. D. 2005: Doxastic Deliberation. *Philosophical Review* 114: 497-534.
- Sorensen, R. 1988: *Blindspots*. 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: 267-297.
- Whiting, D. 2010: Should I Believe the Truth? *dialectica* 64: 213-224.
- . 2012: Does Belief Aim (Only) at the Truth? *Pacific Philosophical Quarterly* 93: 279-300.
- . Forthcoming-a: Nothing but the Truth: on the Norms and Aims of Belief. *The Aim of Belief*, ed. T. Chan. Oxford: Oxford University Press.
- . Forthcoming-b: The Good and the True (or the Bad and the False). *Philosophy*.
- Williamson, T. 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.