

XII*-BELIEVE WHAT YOU WANT

by Paul Noordhof

ABSTRACT The Uncontrollability Thesis is that it is metaphysically impossible consciously to believe that p at will. I review the standard ways in which this might be explained. They focus on the aim or purpose of belief being truth. I argue that these don't work. They either explain the aim in a way which makes it implausible that the Uncontrollability Thesis is true, or they fail to justify their claim that beliefs should be understood as aimed at the truth. I further argue that the explanations don't cut deep enough. Making the aim of truth internal to a state does not explain why we can't produce at will states without this feature but sharing these states' motivational role. I put forward a different explanation. I argue that consciousness makes manifest the attraction of the norm of truth. If we are consciously attending to the question of whether p , we cannot help but make a judgement in line with what the evidence gives us grounds for believing true.

It seems a striking fact that our beliefs are not under the control of our will. It is harder to state what is actually the case and why it is so. For instance, we can want to believe that p and, as causal consequence of this, in fact, believe that p . We can also manipulate our beliefs if wishing it were so is not enough. If we say 'Every day I'm getting better and better' as we rise each morning, we may end up believing it. The last example involves mediated belief production. We have to utter the phrase in order to manipulate our belief. Other cases of mediation would include: motivated treatment of the evidence, and deliberately putting ourselves in circumstances where we are likely to change our way of assessing evidence because of environmental influences such as peer group pressure (Cook, 1987: 441–446). There are also cases in which we can believe at will because our willing the belief makes the proposition believed very likely to be true and we recognise this. If I truly believe you will give me a million pounds if I believe I am a millionaire, I think I can get the money.

Mediation does not appear essential. There seems nothing to rule out a situation in which a desire gives rise to an intention which directly, but unbeknownst to us, produces the belief. What

* Meeting of the Aristotelian Society, held in Senate House, University of London, on Monday, 14th May, 2001 at 4.15 p.m.

does not seem possible is unmediated conscious belief-production:

- (i) I form the conscious intention to believe that p now;
- (ii) as a result of having this intention, I believe that p now without there being any mediating act which helps to produce this belief or the support of evidence for p or a belief that by believing that p I make p true;
- (iii) the intention and the belief stand in the same relation to each other as intention and action do in the case of intentional actions.¹

By contrast, I can form the conscious intention to raise my arm now and, just as a result of having this intention, raise my arm. The characterisation just given explains what I mean when I use the phrase ‘cannot believe/judge at will’. The last clause is needed to distinguish what certainly seems possible, namely that a deviant causal chain connects my intention with my belief. I could intend to believe that I am morally worthless and be so struck by the pretentiousness or self-serving character of this intention I end up believing that I am morally worthless. In such cases, my belief would not be an action for the standard reason. I will call the claim that unmediated conscious belief-production is impossible the *Uncontrollability Thesis*.

My arm does not rise just because I intended to raise it. My body must oblige. I took that as read. Bringing it to the fore suggests a boring explanation of why the Uncontrollability Thesis is true. My body is just not set up to cooperate. In which case, the thought would run, we might be able to control our beliefs with a bit of practice, rather like learning to wiggle our ears.

I leave this position as the default option. It would take the Uncontrollability Thesis to be a contingent *psychological* truth. It seems to me that there is plausibility in the claim that the Uncontrollability Thesis should be interpreted as claiming that

1. Some disagree. Jack Meiland argues that there might be sufficient reason to believe that not- p without the evidence being so compelling that not- p is forced on us. In this case, we can believe that p as a matter of unmediated will. I doubt that. If I am wrong, that doesn’t mean that the considerations I offer later have no application. Meiland’s point will just qualify my claim about the force of the norm of truth in consciousness. Meiland still needs an explanation of why some beliefs are forced on us against our will. To that, I offer an answer. (See Meiland, 1980: 16–17.)

(i) to (iii) are jointly metaphysically impossible (I shall so understand it hereafter). But if this is right, why is it? My particular interest in the question lies in its relationship to two other issues, first, the claim that beliefs are, in some sense, mental states which aim at the truth and, second, the claim that beliefs are governed by the norm of truth. Consideration of these issues will introduce a third, the role of consciousness in the appreciation of norms.²

I

The Aim of Belief. I agree with David Velleman that the right way to characterise what it means for beliefs to be aimed at the truth is that beliefs are regardings of propositions as true for the purpose of getting their truth value right (Velleman, 2000: 251). There are a number of ways in which this may be understood. One option would be to think of beliefs as the products of a system which has been designed to record truths by God or, more likely, by natural selection or by people (in the case of computers of the future). I don't deny that beliefs may have purposes in one of these ways. It just seems to me unlikely that this will provide an explanation of why we can't believe at will. A designer might set up a system so that beliefs are produced in response to evidence or facts. However, any system can malfunction and any designer can knock something off whose purpose is to believe the truth but which is wholly unsuited for that purpose, for example, by having beliefs at will. Yet we have no explanation of why there are no cases of malfunctioning in which somebody produces a belief in themselves just as a result of a conscious intention, or why there are no creatures which have defective belief systems of this general kind. If beliefs derive their purposes from their designers, then there is no reason why being conscious

2. Let me briefly mention one alternative theory. Scott-Kakures argues that the direct production of a judgement by the intention to judge that p requires the intention to monitor the progress of the action. This monitoring, he alleges, requires a rational connection between the cognitive perspective at the time of the intention and the cognitive perspective at the time of the judgement. By hypothesis, this will be absent in the case of intentional judgement unsupported by evidence. The main problem with his suggestion is this. A rational connection is not required in the case of monitoring arm-raising. To claim that it is required in the case of the direct production of judgements seems to be a restatement of the claim that we cannot judge at will (unsupported by reason) rather than an explanation of it. (For further discussion of Scott-Kakures position, see Radcliffe, 1997: 145–151.)

of intending to produce a state with that aim should cause us any difficulty. Their aim need not be ours. So, while beliefs may have truth as a biological aim (say), that is not an explanation of the Uncontrollability Thesis.

There is another possibility. Actions have aims too. But now we hit a snag. There seems a clear sense in which beliefs cannot have aims in virtue of being actions. Beliefs are not things we do. So it does not seem open to us to explain the Uncontrollability Thesis in this way. Indeed, the very fact we seek to explain seems to render impossible the application of this idea to resolve the problem.

We may get round this by focusing on judgement. A judgement that p is an action, and yet it seems that a version of the Uncontrollability Thesis holds just as much for judgement as it does for belief. The claim that judgement is an action is contentious. Nevertheless, it seems to capture the phenomenology of judgement. In making a judgement, I can be aware that I could have made no attempt to make the judgement but chose otherwise. I may also be aware of the effort which went into gathering the material necessary for making the judgement, for example, evidence. I may also be aware of my careful treatment of the evidence and assessment of what it supports. All of these things bespeak *activity*.

The feeling that judgement is not an action seems rooted in various inconclusive observations. For instance, it might be thought that I cannot do A if it is not possible for me to fail to do A. But it is clear that this is not true even if it were true that I would not do A of my own freewill. A brain surgeon could make sure that my arm moves at time t even if I don't will it. That doesn't mean that my raising my arm at t is not an action (Frankfurt, 1969: 5–8). For this reason, the Uncontrollability Thesis is *prima facie* compatible with judgement being an action. Often it is remarked that judgement involves an element of passivity. We put ourselves in the right position evidentially speaking and then it just happens that we judge. On the face of it, this remark would impugn more than judgements. Kicking you in the shins is something I did. But according to this view, that's not so clear. I just put myself in the right position: drawing my foot back, bringing it forward smartly. The momentum forward of my foot made it impossible for me to stop it just before it made contact with your shin. An element of passivity is not enough to

turn something from action to non-action. Another observation is that actions are something we set out intentionally to do. I do not set out intentionally to judge *that p*. Nevertheless, I do set out intentionally to judge and, indeed, to make true judgements. Actions must be intended under some description, not every description. With phenomenology in favour and no good reasons against, I think we may conclude that judgements are actions.

If judgements are actions, then it is plausible they have purposes. My judgement that *p* invariably involves the formation of a belief that *p*. Beliefs may have, in an extended sense, the aims of the judgements that involve their formation. So the key question is whether we can explain the Uncontrollability Thesis by noting that the judgements which involve their formation are actions aimed at the truth.

There seem to be broadly two ways in which judgement may be aimed at the truth in the specified sense. The first is that truth somehow enters into our action of judging rather like moving our legs enters into our action of running. Something may be an aim of an action by being part of the constitutive means by which the action takes place. If I am running, I aim to move my legs. If I am judging, I aim to endorse the truth. The second way in which truth may be the aim of judgement is that truth characterizes the distinctive motive for our action of judging.

Is truth a constitutive means to judgement? One way in which truth might seem to enter into our act of judging is that when we judge that *p*, we, in effect, judge that *p* is true (so long as we have the concept of truth). However, this suggestion is not going to provide us with something which explains the Uncontrollability Thesis. When I desire that *p* or imagine that *p*, I desire or imagine that *p* is true. Yet it is not the case that I can only desire what I take to be true, and I can imagine that *p* and that not-*p* at will (Velleman, 2000: 247–248).

Desiring and imagining aren't assessable as true or false, whereas judgement is. Judgement is not the only propositional attitude that is assessable as true and false though. There is belief (of course) and there is also thinking or entertaining a proposition. You can have the thoughts you consciously choose to have. Why doesn't an equivalent of the Uncontrollability Thesis hold for thought? An initially plausible answer is this. When subjects think and judge, they wield concepts. The concepts that constitute a particular thought or judgement have conditions of

application specified in terms of the conditions under which the propositions they form would be likely to be true. If a subject attempts to use concepts according to these conditions of application, the subject judges; otherwise he or she is having a thought. The reason why a subject can't choose whether to judge that p or not- p is that, if a subject were not to attempt to adhere to the concepts' application conditions, the subject would not be making a judgement.

The problem with this proposal is that the particular means by which an act of judgement takes place is not the only candidate for being a distinctive feature of judgement. Judgements, along with beliefs, also have typical effects. Upon them we base our behaviour to satisfy our desires. Judgement and belief have a distinctive motivational role which thought does not. This opens up the possibility that, even if *we* cannot consciously judge that p without considering the grounds of the truth of p , there are other creatures who can. This might not immediately be clear because our use of the term 'judgement' seems to display a process-product ambiguity. According to this view, the act of judging that p (the process) culminates in the judgement that p (the product). It could be argued that the act of judging must be characterised in terms of a certain kind of concept use, even if the resultant state can be characterised in terms of motivational role. Unfortunately, this doesn't help. For one thing, it's not clear that we can effect the separation. We could characterise the act of judging in terms of its successful production of a state with an allegedly distinctive motivational role. For another, we would still have no explanation of why the resultant product of the process of judging cannot be produced at will. Even if I cannot consciously produce a judgement that p in this sense without engaging in a certain kind of concept use, it is not ruled out in other creatures. We are back with wiggling the ears. So the substantial issues remain in place. My own preference is to use 'judgement that p ' to capture the same kind of idea as the description 'scores a goal'. The action of scoring a goal involves a process which requires a certain outcome in order to be realised. In the case of judgement, the outcome is belief. The issue is how one should characterise the success: should this be in terms of a certain motivational role being realised or in terms of a certain type of truth-aiming activity?

Velleman has argued that the motivational role of judgement and belief is, in fact, not distinctive of it. It is shared by imaginings that *p*. If this were right, then there would be a way to resist the point just made. He argues that, in the context of children's play, *imagining* that one is an elephant (say) disposes a child to behave as if he or she were an elephant (i.e. as if he or she believed that he or she were an elephant) (Velleman, 2000: 255–263). There are two natural anxieties to have about this observation. The first is that, even if there are some limited behavioural similarities, in the round the child would behave differently if he or she actually *judged* that he or she was an elephant. Would the child seek to live outdoors, consume elephant edibles and so on? Would he or she be worried at the transformation? I think not. Imaginings have their motivational role limited in clear ways that judgements do not. Second, even if imagining that *p* in a certain context has the same motivational role as judging that *p*, there will be other contexts in which these things come apart. These different contexts may be enumerated to enable us to distinguish belief or judgement from imagination. I think versions of these points extend to the other cases Velleman considers. Thus, it still seems possible that some creatures might be able to will a state which plays the motivational role associated with judgements in all contexts, i.e. a judgement.

This point can be reinforced by noting that we can make unconscious judgements without being attentive to the grounds of the judgement. Or, if it is preferred to limit the notion of judgement to consciousness, then acquisitions of belief may occur unconsciously without attention to the grounds for the belief. Whichever way we put it underlines the thought that, even if a certain kind of concept use invariably occurs when we make conscious judgements, it does not follow that it will for other creatures (Bennett, 1990: 94–95). Of course it is possible to insist that judgement must involve the appropriate kind of concept use, but this needs motivation if the existence of the alternative characterisation of judgement is borne in mind. It also fails to promise a sufficiently deep explanation of the phenomenon in which we are interested. It seems just as true that I cannot consciously intend to produce a state upon which I am prepared to act in all the circumstances I would be prepared to act upon a judgement.

We could deal with both concerns if we could establish that the truth-aiming feature and the motivational role go together. It might be argued that we could not view a judgement as something upon which we may act in serious pursuit of what we desire unless we also view it as something whose existence is grounded in the truth of what we judge (Velleman, 2000: 272–274). Unfortunately, whatever its general merits, this thought is ill-suited as a response to objections to truth being a constitutive means to judgement. Even if we arrive at a judgement by dodgy means, it does not follow that the way in which we view the judgement now is infected by it. Once we have judged that p , we might engage in rationalisations about how p is well-supported after all. This would still be a case of consciously judging at will. The rationalisations would provide an explanation of why we were under the illusion that the Uncontrollability Thesis is true.

Let me now turn to the second way in which we might seek to understand truth being the aim of judgement. The suggestion is that, since judgement is an action, and actions are generally the result of means-end beliefs and desires, judgements derive their purposes or aims from the beliefs and desires that give rise to them. In which case, the thought must run, whenever we make a judgement we desire the truth (Walker, 1996: 104–107; Walker, forthcoming).

An explanation of the Uncontrollability Thesis based on this idea faces considerable problems. First, the explanation offered seems to be trivial. It is plausible that we cannot desire the truth as regards to whether p is the case and, as a result, intend to judge that p regardless of evidence. So *if* judgements are aimed at the truth, then we have an explanation of why we cannot judge at will. However, no reason has been given for supposing that judgements must be aimed at truth. Instead, that just seems to be stipulated. But why can't judgements occur as a result of other combinations of belief and desire, for example, the desire for something reassuring or the desire to behave in a certain way? We still need an explanation of why this should be ruled out.

Second, it does not seem plausible that we always desire and thereby intend to judge the truth. There are plenty of cases where it appears we would rather not know things or do not have the truth in mind. Questions of appropriateness sometimes come into play. This is particularly dramatic with regard to situations in

which we view endorsement of a certain proposition as favoured by bigots and discriminators. An opposing judgement then seems more appropriate.

Third, the explanation offered appears to be too general. I have suggested that we have no grounds for rejecting the idea that unconscious judgements may occur at will. This might be incorrect but it has some degree of plausibility. If the present explanation were right, it would be just as clear that unconscious judgements could not be produced at will.

Finally, even if we must always desire the truth when making judgements, why couldn't there exist creatures who sometimes elect to make smudgements rather than judgements? Smudgements are states which have the full motivational role of judgements and, indeed, may be aimed at the truth by design but which are not truth-aimed in the sense specified by the theory under consideration. The smudgers on occasions think it better to have control over what leads up to behaviour. Perhaps they are committed to a certain doctrine about what is required for free will. Allowing that such creatures are possible is as bad as allowing for the possibility that one might make judgements at will.

II

Consciousness, constitutively normative actions and the norm of truth. The discussion of the previous section has given us some idea of what we need for a successful explanation of the Uncontrollability Thesis. We must make sense of the thought that, in describing the aim of a judgement as truth, we are describing one of our aims in judging. Yet, we must not do this in a way which fails to accommodate the special role of consciousness. Perhaps part of the problem is that the options hitherto considered assimilate the aim of judgements to some previously well understood application of the talk of aims and purposes instead of taking it as sui generis. One way of cashing this out is to argue that judgements are constitutively normative actions. It is part of their nature that they are always criticisable for being false (Williams, 1970: 136–137). It is in this sense that they are aimed at the truth. In making a judgement, we are opening ourselves up to this

dimension of criticism.³ Another way of cashing it out is to suppose that our consciousness makes manifest the attractiveness of the norm of truth to us and we seek to judge in accordance with it.

It is pretty clear that merely taking judgements as constitutively normative actions criticisable for being false is not sufficient to explain the Uncontrollability Thesis. Assertions are also constitutively normative in this sense. Yet we can intend to make a false assertion at will. So what we need is an identification of the extra factor. It is in this context, that Williams' discussion is salient. The key elements of the discussion occur in the following passage.

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 reality. At the very least, there must be a restriction on what is the case after the event; since I could not then, in full consciousness, regard this as a belief of mine, i.e. something I take to be true, and also know that I acquired it at will (Williams, 1970: 148).⁴

There are two ideas in play in this passage. Each provides a potential explanation for why our consciousness of judgement as having a certain aim is incompatible with successfully intending to judge that *p* consciously at will. The first is that

- (i) It is not possible for me to judge at will in full consciousness because I could not think of it beforehand as a judgement.

3. Williams mentions two other features of being aimed at the truth that are best set aside. The first is that, if a subject recognises that his or her belief is false, the subject abandons the belief. This appears to be the consequence of the subject's recognition of the normative character of the state—rather as the Uncontrollability Thesis is alleged to be such a consequence—instead of being part of this character. The second is that it is, in some sense, paradoxical to believe that *not-p* but *I believe that p*. The observation is due to G. E. Moore. Various explanations have been offered of the paradoxicality of believing this which do not appeal to the truth-aiming character of judgement. Even if these explanations are wrong, it is unlikely that the paradoxicality serves to characterise the way in which judgements and beliefs are normatively constituted states rather than being a consequence of this character. In any event, Williams does not appeal to these two features in his discussion of the reason for the Uncontrollability Thesis being true (Williams, 1970: 136–137, 148–149).

4. Williams' argument for the impossibility of believing at will has generated many objections. I do not propose to rehearse all of these and shall just focus on the kernel of his proposal and how it displays a particular approach to the issue (see Winters, 1979; Bennett, 1990).

It is not clear that this suggestion provides what we need. I can intend to do something which I will later take to purport to represent reality. I can think of what I am trying to do now as a judgement because of my knowledge of what I will think about it then. It is hard to see why we should insist that it is essential to thinking of something as a future judgement of mine that I should now think of it as purporting to represent reality. This line of thought may be behind Williams' own shift to the second idea.

- (ii) It is not possible to judge that p if we judge afterwards that the basis for our candidate judgement was not sufficient evidence that p is true but rather intending to judge that p .

It is not clear that this suggestion is any better. First, even if we judge that the basis for our candidate judgement was a conscious intention to judge that p , that does not undermine the judgement if, now, there is evidence in favour of it (Winters, 1979: 252–253). However, let us suppose that there is no supporting evidence. Then I concede that, if we judge afterwards that the judgement was produced at will, our judgement that p may not form a *sustained* belief that p . But why does this rule out the judgement at the time?

The considerations offered against the two features Williams identifies are underlined by the comparison with assertion. When I make an assertion, I might open myself up to the dimension of criticism that concerns whether or not the assertion is true but I can quite consciously break the rules. Similarly, I can think of my past judgement, or the future judgement I intend to make, as aimed at the truth while at the same time thinking of it as something in whose production I broke or will break the rules. This applies even if I am the rule maker. For instance, I. L. Humberstone suggests that beliefs are states governed by the intention to alter so that, if not p were the case, then the subject would believe that not- p (Humberstone, 1992: 75–81).⁵ There is no reason why this governing intention could not be overridden on certain

5. I. L. Humberstone has a rather subtle discussion of the possible semantics of the governing conditional intention which is not reflected in my crude characterisation of his position. The reason why my simplification is acceptable is that my subsequent discussion does not depend on the subtlety and my characterisation retains the element he emphasises: failure of contraposition.

occasions. So long as I give an explanation of why the state produced is still governed by the rules, there isn't a problem with thinking of it as a judgement or belief. Towards such an explanation, I can note that when the judgement is produced, I will take it to purport to represent reality. Likewise, after the judgement is produced and in the absence of evidence to support it now, sustained belief will be undermined by recognition of its ancestry or by counter-evidence.

The case of false assertion also brings into focus the point about the aim of truth having to be *our* aim. When we make a false assertion we have other objects in view. It might seem that the same is possible in the case of judgement. Taking assertion and judgement as constitutively normative acts does not imply that, when we perform such acts, the norms which constitute them are our aims. So it does not much matter what the constitutive aim of judgement is. What matters is the aim which, in fact, governed its production. Hammers may be essentially nail bangers but their shaft could, in a particular case, point in the direction of hidden treasure.

Intentionally false assertion also involves a certain dissociation. We do not think that, if what was falsely asserted were relevant to our most important desires, we should now act upon it across all, i.e. not just imaginary, contexts. In the case of judgements we do. We would not, in full consciousness, act upon a judgement unless we took it as representing reality at the time at which we acted. This is not a matter of judgement's constitutive aim but rather of its actual apparent achievement. It is closer to the truth to observe that we cannot judge at will and at the same time take the judgement to represent reality because the actual purpose of that particular judgement would not be the representation of reality.

These observations lead me to favour the second approach I identified at the beginning of this section. According to this approach, the aim of truth is not internal to judgement but applied from outside *via* consciousness. Consciousness makes manifest the attractiveness of being disposed to act upon the truth. To set out my proposal, I need to introduce the notion of a *j-state*. J-states include smudgements, judgements (if the latter are different from smudgements by being constitutively aimed at the truth), perceptions, beliefs and their smudgement-produced

equivalents. A j-state is not undermined so long as the subject's remaining relevant j-states do not collectively suggest that there is no reason to suppose that it is true. The explanation of the Uncontrollability Thesis has three components. The first is a characterisation of the (or at least *a*) norm of truth.

- (A) You prima facie ought to have a j-state that *p* iff
- (i) the question of whether *p* arises for you,
 - (ii) *either* you have j-states to the effect that the evidence is sufficiently favourable to *p* *or* a j-state that *p* will be true if you have a j-state that *p*, and
 - (iii) you do not have a non-undermined j-state that the j-states mentioned in (ii) have origins that question their likelihood of truth.

Second there is a thesis about the nature of consciousness.

- (B) Metaphysically necessarily, you have a conscious j-state that *p* iff
- (i) the question of whether *p* arises for you,
 - (ii) you have conscious j-states to the effect that the evidence is sufficiently favourable to *p* *or* a conscious j-state that *p* will be true if you have a j-state that *p*,
 - (iii) you do not have a conscious non-undermined j-state that the j-states mentioned in (ii) have origins that question their likelihood of truth, and
 - (iv) you consciously attend to producing a j-state about whether *p*.⁶

Third there is a categorisation of the norm of truth.

- (C) The Norm of Truth is a practical norm.

The basic idea is that it is part of the nature of conscious attention that it gives determinative weight to the norm of truth. It

6. In fact, I believe that this claim must be moderated slightly. There are circumstances in which a very pressing practical or moral norm closely tied to what an agent most values can override the norm of truth. This does not affect the explanation of the Uncontrollability Thesis. Since I cannot establish the qualification here I leave the proposal in unqualified form. I suspect that other adjustments will also be necessary; so to an extent my proposal is programmatic.

makes truth-likelihood the deciding factor in the formation of a certain kind of motivational state.

It is easiest to understand the details of the proposal by considering how it deals with the alleged possibility that there might be creatures who sometimes elect not to make judgements but rather smudgements. Suppose one of our smudgers perceives that its cakes are burnt. This is one of the ways in which the question of whether p may arise. If a subject has an interest in whether p or is going to make a judgement one way or another as a result of perceiving something, then the question of whether p has arisen. Otherwise not. A subject is not required to believe every truth. In the case of the smudger and the cakes, its interests are also involved. It would be much happier if it believed that they were perfect when handing them out to friends. According to the hypothesis considered at the end of the last section, one can't make a judgement unless one desires to judge the truth. According to the first type of theory considered in this section, the constitutive norm of judgement makes it impossible for the subject to be conscious of making the judgement and not attempt to judge what is likely to be true. Both prescriptions seem avoidable by the smudger. However, the perception that the cakes are burnt is a j-state. The smudger does not smudge that the perception was brought about by will or, if it did, the smudgement is undermined. Hence, according to (A), the smudger ought to smudge that the cakes are not perfect but burnt. According to (B), if the smudger consciously attends to producing a j-state about whether or not the cakes are perfect, the smudger will have a j-state that the cakes are not perfect. The manifest attractiveness of the norm of truth to consciousness makes truth the actual aim of the smudgement.

A j-state is one upon which we would act as relevant if we were conscious of it. I suggested that this was the crucial difference between assertion and judgement. This is borne out by our consideration of the smudgers. We don't think that they can smudge that the cakes are perfect as a result of a conscious intention to make this smudgement. The diagnosis of the difference between assertion and judgement and the intuition about what is possible for smudgers coincide. Even if truth is the constitutive norm of judgement, there is no role for this idea in the explanation of the Uncontrollability Thesis. Instead, what is important is the way consciousness makes manifest the attractiveness

of being disposed to act upon what is true, and the unattractiveness of being disposed to act regardless of the truth. Talk of the smudgers is just a way of dramatising this fact.

Scepticism about my proposal is apt to fall under three headings. First, it might be questioned whether we have any reason to believe that conscious attention gives determinative weight to the norm of truth. Second, it might be questioned whether I even need to appeal to the norm of truth in the story. Why isn't (B) enough? Third, it might be questioned why we should suppose that the norm of truth is a practical norm.

I think there are the following non-conclusive reasons for supposing that conscious attention gives determinative weight to the norm of truth. First, when we are presented with strong evidence for p , or when we attend to the logic of a particular valid argument, we feel that we ought not to judge otherwise than in favour of p , or that we cannot maintain the truth of the premises and deny the conclusion. This feeling about how we ought to think displays our apprehension of the norm of truth in consciousness and its effect upon our reasoning.

Second, our attentive perceptual experience of objects and properties in the world or, for that matter, our conscious experience of features of our own mental life seems to have a normative force. We feel that we ought to believe that the facts presented hold. This feeling can only be overcome by powerful considerations from elsewhere in our system of beliefs. We should ponder, for a moment, why conscious perceptual experience has this feature of appearing to compel belief rationally. I suggest that it is because conscious experience makes manifest the attractiveness of the norm of truth. This is no accidental feature of conscious experience. In both cases, it is absolutely central to the nature of phenomenal consciousness that it manifests the attractiveness of the norm of truth in this way. It does not seem possible that we could conceive of conscious experience anything like ours which did not rationally compel belief according to what strikes us as evidence for the truth of a proposition that p , or perception of objects and properties in the world, or facts about our mental life.

Third, judgements in line with evidence which we would rather not make are still an exercise of our freedom. It is true that we sometimes say that we are not free to judge otherwise given the

evidence. But that does not mean that, in making the judgement, we are not exercising free will. It merely means we are unable to do otherwise. The connection between these two notions is rightly contentious. A plausible explanation of why such judgements are an exercise of freedom is that they are responses to the norm of truth. If a mental state is a response due to the recognition of what a norm requires of one, the upshot is an exercise of freedom. Whether it is the only way in which freedom is exercised, as some suggest, is another matter (Pettit and Smith, 1996: 444–449). So while it is quite possible to discern a way in which we feel constrained, there is a crucial difference between this kind of the case and the case of an addict. Addicts are neither responding to a norm nor able to do otherwise when they act so as to feed their habit.

Fourth, the appeal to norms and the nature of conscious attention explains why (B) holds. Of course, we can try consciously to produce a judgement against the evidence by a bit of brainwashing, i.e. indirectly. This only works by dulling the influence of the norm of truth by focusing on distracting issues and the like. It does not work when we actually attend to the evidence. The evidence both reveals the application of the norm of truth and, because we are consciously attending to the evidence, makes manifest its attractiveness.

Fifth, the appeal to the role of the norm of truth explains why we are inclined to move from reasons for judgement that p to reasons for ascribing to ourselves the judgement that p . The weight of the norm of truth in conscious attention means that, if we perceive that the reasons favour p , then we will have made a judgement that p . So it is no surprise that the self-ascription of a judgement is based upon our perception of reasons for the judgement.

Of course, I have not established that it is impossible for there to be some very different type of phenomenal experience which did not make the norm of truth so manifestly attractive. If there is, then I should qualify the thesis stated in (B) to apply only to consciousnesses like ours. This does not strike me as undermining the central claim. I think that, if we could imagine such a consciousness, we would think that it is an open question whether the Uncontrollability Thesis would be true for the subject of that consciousness. The actual reference to consciousness

in (B) would need to be more specific. All I want to urge is that something like (B) is correct.

Let me turn to the next source of scepticism, namely the question of why we need the claim about norms at all. Why can't we just assert that (B) is the case? The first point to make is that, although the combination of (A) and (B) asserts more than merely (B) on its own, the justification for believing that (B) is true is our phenomenal awareness of the influence of the norm of truth. Appeal to the norm of truth in no way precludes a reductive understanding of our mental life and the norms which govern it. For instance, the characterisation of the norm I have provided might be an expression of what we desire to desire. So there is no reason to reject the approach recommended here because of an antecedent commitment to some form of naturalism. Other considerations in favour of the assertion of both (A) and (B) rest upon the third claim, the classification of the norm of truth as a practical norm. So I will turn to that matter.

One reason for thinking that the norm of truth is a practical norm is that both intending that the norm of truth is p and judging that p are actions. The norm of truth provides considerations for acting in these ways. Broadly conceived, practical norms are precisely those which provide considerations for action. A second reason is that agents act so as to satisfy their desires. An agent's desires are only satisfied as a result of the agent's action if the beliefs and judgements upon which the agent acts are true. Therefore, it is part of practical reason that our beliefs should be true.

The classification of the norm of truth as a practical norm enables us to defend the claim that judgements have the aim of truth due to the beliefs, desires and intentions behind them. Judgements which are the product of conscious attention cannot fail to be the result of a desire for the truth since this desire is a response to the determinative weight of the practical norm of truth in conscious attention. This answers the first objection raised. We also have a justification of the allegedly implausible ascription of desires and intentions. The reason why the objection seemed plausible was because we were looking for a certain kind of evidence for the desire or intention. Roughly speaking, we wanted it to be explicit in our mind that we had that intention and we wanted to feel the force of the desire for the truth, in the way that we might feel the force of yearning. But now we see

that the nature of phenomenal consciousness bears witness to the presence of the desire by making manifest the attractiveness of the norm of truth.

III

Concluding Remarks. The picture we are left with is this. When I consciously attend to the evidence for whether p or not- p and focus on making a judgement, I cannot help but form a judgement in line with what I believe to be sufficient evidence for the truth. I cannot will to judge something else because of the force of the norm of truth upon conscious attention. Since this norm is a practical norm, I cannot will to judge something different because I cannot want anything more than the truth. Equally, when I consciously attend to my belief that the evidence fails to support p , the unattractiveness of being disposed to act regardless of the truth is made manifest. It determines that I do not judge that p . In one sense then, the Uncontrollability Thesis is not happily glossed as the claim that we cannot believe at will. I have tried to argue that my proposal is the only way in which we might seek to explain the Uncontrollability Thesis. It is, of course, open to you to reject my proposal. You could claim that the Uncontrollability Thesis is merely a contingent psychological truth. Now that I have drawn these options to your attention, though, I think you should believe what you want.⁷

*Department of Philosophy,
University of Nottingham,
University Park,
Nottingham NG7 2RD
email: paul.noordhof@nottingham.ac.uk*

7. I am very grateful for the stimulation, encouragement and trenchant criticism I have received from the following people: Michael Clark, Greg Currie, Eros Corazza, Bill Fish, Jennifer Hornsby, Jeff Ketland, Bob Kirk, Mike Martin, Stephen Mumford, David Papineau, Tom Pink, Charles Travis, Joss Walker, and Terence Wilkerson. Without their help, I probably wouldn't have had a glimmering of understanding of the complexity of the issues I discuss. I am also grateful to the Editor for helping me to be clearer. I would like to thank the AHRB Matching Research Leave scheme for their support.

REFERENCES

- Bennett, J., 1990, 'Why is Belief Involuntary?', *Analysis*, 50, 87–107.
- Cook, J. T., 1987, 'Deciding to Believe without Self-Deception', *The Journal of Philosophy*, 84, 441–446.
- Frankfurt, H. G., 1969, 'Alternate Possibilities and Moral Responsibility', *Journal of Philosophy*, 66(23) and in his (1988), *The Importance of What We Care About* (Cambridge: Cambridge University Press), pp. 1–10.
- Humberstone, I. L., 1992, 'Direction of Fit', *Mind*, 101(1), 59–83.
- Meiland, J., 1980, 'What Ought We to Believe? The Ethics of Belief Revisited', *American Philosophical Quarterly*, 17, 15–24.
- Pettit, P. and Smith, M., 1996, 'Freedom in Belief and Desire', *The Journal of Philosophy*, 93(9), 429–449.
- Radcliffe, D., 1997, 'Scott-Kakures on Believing at Will', *Philosophy and Phenomenological Research*, 57(1), 145–151.
- Scott-Kakures, D., 1994, 'On Belief and the Captivity of the Will', *Philosophy and Phenomenological Research*, 53(4), 77–103.
- Velleman, J. D., 2000, 'On the Aim of Belief', in his (2000), *The Possibility of Practical Reason* (Oxford: Oxford University Press).
- Walker, M. T., 1996, 'The Voluntariness of Judgement', *Inquiry*, 39, 97–119.
- Walker, M. T. (forthcoming), 'Williams, Truth-Aimedness and the Voluntariness of Judgement', *Ratio*.
- Williams, B., 1970, 'Deciding to Believe', in his (1973), *Problems of the Self* (Cambridge: Cambridge University Press), pp. 136–151.
- Winters, B., 1979, 'Believing at Will', *The Journal of Philosophy*, 76, 243–256.