

Leaps of Knowledge¹

0 Introduction

This paper considers two unpopular philosophical claims. One is that *doxastic voluntarism*, the view that agents can, at least sometimes, choose what they believe, is true.² The other is that *strict evidentialism*, the view that all reasons for belief are evidential reasons, is false. I do not aim to defend either view, but rather to show that, if defensible, they are consistent with two commonly held views in epistemology and the psychology of belief: that the aim of belief is either truth or knowledge and that this aim plays a central role in setting the norms of belief.

A difficulty in discussing doxastic voluntarism is that its falsehood is commonly taken as a datum.³ There is an influential body of work suggesting that it is a conceptual truth that beliefs are involuntary. My own view is that whether or not beliefs can be formed voluntarily is an empirical rather than a conceptual question, but that is not a matter that I care to take up here. This paper will assume, for the sake of argument, that doxastic voluntarism is in all relevant senses possible.

There is a parallel difficulty in discussing strict evidentialism (henceforth just ‘evidentialism’). As with the falsehood of doxastic voluntarism, the truth of evidentialism is also commonly accepted as a datum, indeed as a conceptual truth.⁴ I have argued in some detail elsewhere that

¹ This paper has been greatly improved by comments from Timothy Chan, Fritz-Anton Fritzon, Olav Gjelsvik, Mikkel Gerkin, Jessica Pepp, Nikolaj Jang Lee Linding Pedersen, Wlodek Rabinowicz, and audiences at the the University of Copenhagen and at the Lund University Philosophical Society. I also wish to thank especially the Centre for the Study of Mind in Nature at the University of Oslo, owing to whose generous support in the form of a visiting fellowship I was able to develop this paper.

² The classic modern discussion of doxastic voluntarism is Williams (1973).

³ See Hookway (2000), Kelly (2002), and Shah (2006).

⁴ See Parfit (2001), Skorupski (forthcoming).

evidentialism is not a conceptual truth and may not just be assumed as a datum.⁵ I shall not revisit the matter here. For the purposes of this paper, it will suffice to treat the truth of evidentialism as an open question.

Preliminaries now set aside, in this paper I shall argue that both a limited doxastic voluntarism and anti-evidentialism are consistent with the claims mentioned above: that the aim of belief is truth or knowledge and that this aim plays an important role in norm-setting for beliefs. More cautiously, I shall argue that limited doxastic voluntarism is (or would be) a useful capacity for agents concerned with truth tracking to possess, and that having it would confer some straightforward benefits of both an epistemic and non-epistemic variety to an agent concerned with truth tracking.

1 Unstable beliefs with single and multiple fixed points

There is considerable debate about how to interpret claims of the form: *the aim of belief is truth* or *the aim of belief is knowledge*.⁶ I shall understand these aims as pertaining to particular beliefs rather than to an agent's complete belief set. A sketchy, but intuitive, way of understanding belief's having the aim of truth or knowledge is this: belief's regulatory mechanisms are primarily geared towards truth tracking (or towards truth tracking in the right way for the knowledge aim), and a belief's role in our cognitive architecture is to provide an accurate representation of its objects. Even in this sketchy account of belief's having the aim of truth or knowledge, there is much to quibble with and still more left out, but I shall not attempt further elaboration here.

One might propose a variety of cognitive mechanisms for belief acquisition and change. The default assumption in writings about theoretical reason is that agents (or perhaps well-functioning agents) modify their doxastic states in response to changes in perceived evidence.

⁵ See Reisner (2007), Reisner (2008), and Reisner (2009).

⁶ See Velleman (2000), Wedgwood (2002).

This process is involuntary, although voluntary psychological activity (e.g. theoretical deliberation), may trigger the involuntary changes.

In ordinary cases, in which we as a matter of course trust our senses, our beliefs change as we encounter new stimuli. The tree outside my window that appeared to be an aspen in the gloaming now appears to be a sugar maple in the clear sun. These changes in appearance automatically trigger changes in my belief about what kind of tree I have been looking at.

For beliefs that we arrive at as the result of reasoning, belief change is triggered when we take ourselves to have sufficient evidence for the proposition about which we are reasoning. I may initially have no view about who the greatest quarterback of all time is, but on researching the statistics and news archives, I take myself to have sufficient evidence that it is John Elway. Taking oneself to have evidence need not be a conscious process nor entail any particular belief about evidence. Taking oneself to have sufficient evidence for a proposition that one did not already believe normally triggers the formation of a new belief.

This is an attractive model of cognitive regulation, but it is inapplicable to some circumstances. It is on these that I shall focus. The first circumstances of interest are those in which an agent has an unstable belief with either a single fixed point or multiple fixed points. Let us consider an example of each:

1.1 The Numbers game with a single fixed point⁷

Alice has volunteered to participate in a psychology experiment at her university. The experiment uses a new and absolutely flawless mind-reading machine. The parameters of the experiment are known and understood fully by Alice.

Alice is attached to the mind-reading machine, which is in turn attached to a large visual display that is within Alice's view. When the machine is turned on, Alice will be asked to

⁷ This example is adapted from Reisner (2007).

predict what number will appear in 20 seconds time on the display. That number will be determined by a function of what Alice believes the number will in fact be. In particular, the function is $1/2n + 1$, where n is the number that Alice believes will appear on the screen. If Alice changes her mind, the machine resets and will display the new number (according to the function) in 20 second time from when Alice changes her mind. Finally, if Alice has no belief about which number will appear on the display, the number 16 will appear, unless Alice develops a belief about what number will appear prior to the number 16's appearing. In that case, the clock will reset and the $1/2n + 1$ formula will become operative.

Alice appears to be in a pickle if she has any belief but that the number will be 2. Other beliefs will lead her to a new belief and are thus unstable. Suppose that she believes that the number will be 100. She knows that belief will cause the display to display 51, leading her to believe that the number will be 51. She knows that her new belief will cause the number to be 26.5, leading her to believe that it will be 26.5, and so on. If she believes that the number on the display is 2, however, her belief will be stable, true, and she will have perfect evidence that the number will be 2, given her other knowledge.

1.2 The numbers game with multiple fixed points

This experiment is the same as the numbers game with a single fixed point, save that the function that determines what number appears on the screen is more complicated. If n is greater than or equal to 0, then the formula remains $1/2n + 1$. But, if n is less than 0, it is $1/2n - 1$. In the modified experiment, Alice knows and understands the modification.

1.3 The single fixed-point numbers games and evidence

In the first version of the numbers game, there is a single fixed point: 2. The only belief that Alice can have that will be true (and also not self-undermining given her other knowledge) is that the number on the screen will be 2. This raises three interesting questions. The first is whether, and if so why, Alice ought to believe that the number on the screen will be 2. The

second is whether she can believe that the number on the screen will be 2. The third is if she can believe it, how can she believe it.

In answering the first question, the first thing to note is that if Alice ought to believe that the number will be 2, it is not on account of evidence in any very direct way. To see why, some stipulations are required. Let us suppose that Alice lacks a brute disposition to form stable beliefs and that she also lacks a brute disposition to believe that she has a disposition towards stable beliefs. Let us also suppose that Alice is not self-deluding such that she can get herself to believe there is evidence when it is plainly lacking.

To have evidence that the number on the screen will be 2, given her other knowledge, Alice would need to have evidence that she will believe that the number will be 2. Absent brute dispositions of the kind ruled out by the above stipulations, it is difficult to see where this evidence would come from. Alice arrives at the experiment lacking the belief that the number will be 2. She can reason her way to a conditional claim, namely that if she believes the number that will appear on the display will be 2, it will be 2. Conditionals are not evidence for their own antecedents, so this new piece of information will not provide evidence that Alice will believe that the number on the screen will be 2.

To the extent that Alice's cognitive regulatory mechanism for belief is sensitive to her evidence, Alice will not come to believe that the number will be 2. Indeed, naïve reasoning leads Alice down a very difficult path. As seems likely, she will have no view about what number will appear on the display, when she starts the experiment. She knows that this will cause the number on the display to be 16. If she comes to believe on that basis that the number will be 16, that new belief will cause the number to be 9. Her arithmetic reasoning will lead her to the conclusion that the number will, as a result of her new belief, actually be 5.5, and so on. Eventually Alice's mathematical reasoning capacity will be exhausted if she keeps on, and presumably she will cease to have a belief, leading, as she will know, to the number on the display's being 16. This will lead to her believing that the number will be 16, and the cycle will repeat.

Alice will not have evidence on a simple probabilistic view about evidence: p is evidence for q just in case the probability of q given p is higher than the probability of q given not p .⁸ Perhaps some more sophisticated account of evidence would allow for the following: p is evidence for q if p entails that any belief but q will be false.⁹ On this view, p would be the facts (or Alice's knowledge) about the numbers game; q would be that the number will be 2. This more sophisticated view about evidence provides an answer to the first question. Alice does have a reason to believe that the number will be 2, and that is because she has evidence for it. It also provides an answer to the second and third questions. Alice can believe that the number will be 2, and she can do so because she is sensitive to her evidence.

1.4 The multiple fixed-point numbers game and evidence

The same considerations, *mutatis mutandis*, apply to the multiple fixed-point numbers game as apply to the single fixed-point numbers game, when working with the simple view of evidence. This is not the case when working with the sophisticated view of evidence.

In the multiple-fixed point numbers game, there are two fixed points: 2 and -2. Thus, it is not true of 2 that one will have a false belief unless one believes it, and the same goes for -2. The sophisticated account of evidence does not even yield that one has evidence that the number will be either 2 or -2. Alice's beliefs will be false unless she believes that it will be 2 or -2 (the scope of the belief is $B(2 \vee -2)$). Holding in place the supposition that she has no brute dispositions to believe one way or the other, or brute dispositions to believe that she will believe one way or the other, Alice cannot have evidence for any belief in the multiple fixed-point numbers game.

2 Believing absent evidence

⁸ See Harman (1999).

⁹ This has been suggested to me in response to the original discussion of the single fixed-point numbers game in Reisner (2007). I do not endorse this view, but I have no interesting argument against it.

If the sole regulatory mechanism of belief formation is sensitivity to evidence, then Alice could not form the belief that the number will be 2 in the single fixed-point numbers game, if what she is in fact sensitive to is evidence of the simple kind. If she is sensitive to the complex kind of evidence, then she can form her belief based on evidence. Neither the simple nor the more sophisticated accounts of evidence would trigger a belief that the number will be 2 or the belief that the number will be -2 in the multiple fixed-point version of the numbers game. In these cases, Alice cannot believe that the number will be 2 (or believe that it will be -2), if the only cognitive regulatory mechanism for belief formation is sensitivity to evidence.

Someone wishing to defend the view that Alice can believe that the number will be 2 in the single fixed-point numbers game has various non-evidential regulatory options at his disposal for explaining how. I shall briefly canvas two of them, but more of the same style may be possible. The first of these is what I shall call the ‘stability proposal’:

C1. *Stability*: Absent any other regulatory considerations, an agent’s beliefs are sensitive to stability.

In cases for which there is a single fixed point, stability regulation may be sufficient to form a belief. In the single fixed-point numbers game, that belief would be 2.¹⁰ Here is a second proposal:

C2. *Knowledge*: An agent can, and sometimes does, believe something when she knows (or has a very high confidence) that it will be true, if she believes it.¹¹

I do not intend to discuss the plausibility of either the stability proposal or the knowledge proposal as cognitive regulatory mechanisms, but I do wish to highlight an important feature of them both. In the single fixed-point numbers game, the stability proposal and knowledge proposal would both yield the fixed-point solution. If, when confronted with epistemic

¹⁰ I discuss the stability proposal in more detail in Reisner (2007).

¹¹ The knowledge proposal is also discussed in more detail in Reisner (2007).

environments in which (simple) evidence was lacking, one's beliefs were formed by *stability* or by *knowledge*, one could still form all of one's beliefs without choosing. To put things another way: *stability* and *knowledge* would lead to the stable belief in an agent for whom all belief formation was involuntary.

These mechanisms are not sufficient for forming a stable belief, absent the capacity to voluntarily choose beliefs, for the multiple fixed-point numbers game. It is easy to see why for both *stability* and for *knowledge*. In the case of *stability*, there is *ex hypothesi* more than one stable belief. Thus, a cognitive regulatory mechanism that selected for stability would have its choice underdetermined. Exactly the same is true for *knowledge*. In the multiple fixed-point numbers game, Alice knows that her belief will be true, if believed, when the contents of that belief are that the number on the display will be 2, and also when they are that the number on the screen will be -2. Underdetermination is the death of involuntary belief forming mechanisms.¹²

Within the range of cases to which *stability* or *knowledge* apply, an agent who possessed a limited capacity to choose her beliefs would be able to arrive at one of the fixed points. She could choose either 2 or -2. If *stability* or *knowledge* restricted the membership of her belief choice set, her cognitive regulatory mechanism would still in an important sense be truth governed. Only true beliefs would (and could) be acquired and held in the multiple fixed-point numbers game and other scenarios with the same basic structure. I shall return to this point in §4.

3 Non-evidential reasons for belief

We can distinguish between two broad classes of normative reasons for belief: evidential and non-evidential.¹³ It will be useful to have working definitions of each:

¹² Strictly speaking, one could start rigging up more elaborate cognitive models. When encountering multiple fixed-point scenarios, an agent's cognitive regulatory mechanism could have a randomiser that arbitrarily selects one of the options. This addition strikes me as *ad hoc*.

¹³ See Reisner (2009) for an extended discussion of the difference and Reisner (2008) for a discussion of how evidential and non-evidential reasons for belief can be weighed against each other.

R1. *Evidential reasons for belief*: Fact f is an evidential reason for agent a to believe b just in case f is evidence¹⁴ for the contents of b .

and

R2. *Non-evidential reasons for belief*: Fact f is a non-evidential reason for agent a to believe b just in case f is a reason for a to believe b and f is not evidence for the contents of b .

R1 is designed to ensure that reasons for belief are present or absent at the same time and in the same ways as evidence. I am hoping to capture the views of strict evidentialists¹⁵ in doing so, because ‘normative reason for belief’ is a term of art. I doubt that there is a properly intuitive notion to capture. R2 is designed to capture any reason for belief that is not an evidential reason. This leaves open the question of whether there are non-evidential reasons for belief. I have argued that there are elsewhere,¹⁶ but here I am assuming that there are non-evidential reasons for belief for the sake of argument.

If R1 is correct, then whether there are evidential reasons for belief in the single fixed-point numbers game depends wholly on whether one accepts the simple or the complex account of evidence. Neither the simple nor the complex account of evidence will deliver an evidential reason for belief in the multiple fixed-point numbers game. This is because the presence (or absence) of an evidential reason for belief is entirely parasitic on the presence (or absence) of evidence. As discussed in §1.4, no evidence is available to Alice in the multiple fixed-point numbers game.

A normative version of both *stability* and *knowledge* delivers reasons for belief in both the single and multiple fixed-point numbers games. I shall only discuss the multiple fixed-point numbers game here for the sake of economy, but the same considerations apply *mutatis mutandis* in the

¹⁴ If one thinks that evidence is relative to an agent, then this can be expressed as ‘evidence for a ’. For a thorough discussion of evidence and theoretical reasons, see Skorupski (forthcoming).

¹⁵ See Kelly (2002), Parfit (2001), Shah (2006) and Skrupski (forthcoming) as examples.

¹⁶ See Reisner (2007), Reisner (2008), and Reisner (2009).

single fixed-point version, save that the reasons are unique in the single fixed-point case (because there is a single fixed point). A first attempt at a normative version of *stability* and *knowledge* is:

R3. *Normative stability*: Fact f is a reason for agent a to believe b if f makes it the case that b is a stable belief for a .

and

R4. *Normative knowledge*: Fact f is a reason for agent a to believe b if f makes it the case that a knows that b will be true, if a believes b .

As they stand, neither R3 nor R4 is very plausible as an account of a type of normative reason for belief. A modified version of R4 can salvage *normative knowledge*, but for reasons that I shall presently explain, *normative stability* is more difficult to preserve.

The benefits of having stable beliefs depend on the context. One could generate evil demon cases in which one is either rewarded or punished for having stable beliefs. They are of no special interest here. Outside of evil demon cases, and more earthly cousins, most of the benefits derived from having stable beliefs are in the conservation of cognitive resources. Provided that the topic of the belief is of no great moment (as in the example in this paper), an unending cycle of reasoning is a waste of mental energy, not to mention time. Stability is no guarantor of truth, and in cases in which the stable beliefs are the true ones, one suspects that the reason giving force of stability is merely derivative from the reason giving force of truth.¹⁷ Thus, I tentatively conclude that the fact that a belief would be stable is unlikely to have non-derivative reason-giving force. Cases in which being a stable belief provides a reason are cases in which there is already some other more basic reason-giving consideration.

¹⁷ For an interesting discussion of the way that truth crowds out other epistemic considerations, see Pritchard (forthcoming).

Normative knowledge, on the other hand, has inherent plausibility. Its inherent plausibility comes from an assumption set out in §0, namely that the aim of belief plays a role in setting the norms of belief. This requires elaboration.

In this paper, our working understanding of what it is for belief to aim at truth is that the cognitive regulatory mechanisms for the formation of individual beliefs are geared towards tracking truth. We may modify our understanding of the aim of belief if we take the aim to be knowledge, including appropriate riders concerning an agent's being appropriately epistemically situated.¹⁸

If one allows that belief's constitutive aims play a role in setting the norms for belief, then one expects at least some of the normative reasons for belief to be concerned with arriving at the truth. This is vague. A precise way of understanding it is offered by strict evidentialism: all reasons for belief are evidential reasons.¹⁹ A more ecumenical evidentialist can offer the weaker position: insofar as the aim of belief is truth, some normative reasons for belief are evidential reasons.²⁰ Evidence, however, does not exhaust the ways in which the aim of belief might set truth-tracking belief norms. To see why, consider the numbers game example.

In the multiple fixed-point numbers game, Alice knows that if she believes that the number that will appear on the display is 2, then that number will be 2. Thus, Alice knows that she will acquire a true belief if she comes to believe that the number on the display will be 2. For the reasons discussed in §1.4, Alice lacks evidence that the number on the display will be 2, but she is guaranteed that if she forms the belief that it will be, then she will have acquired a new true belief.

¹⁸ See Williamson (2002) for an extended discussion of the idea that belief aims at knowledge. The notion of being *well situated* is much more difficult to spell out than one might expect and is problematic. For a concise discussion, see Blöme-Tillmann (2007).

¹⁹ See Adler (2002)

²⁰ See Steglich-Petersen (2009) for more discussion on the relationship between the aim of belief and our reasons.

Acquiring this new belief is consistent with the aims of her cognitive regulatory mechanisms for believing. She aims at having her individual beliefs come out true, and it is only by believing that the number on the display will be 2, or by believing that the number on the display will be -2, that she will have beliefs that she knows will be true. Indeed, she knows that all other beliefs will be false. Her beliefs are guaranteed to fail to track the truth, unless she forms the belief that the number on the display will be 2, or unless she forms the belief that the number on the display will be -2. If we take the aim of belief (and the concomitant goals of her cognitive regulatory mechanism) seriously as setting norms for belief, there is a reason to believe that the number will be 2, or to believe that the number will be -2.

At first blush, this appears to commend R4, the principle that I earlier said was flawed:

R4. *Normative knowledge*: Fact f is a reason for agent a to believe b if f makes it the case that a knows that b will be true, if a believes b .

It does not. R4 gets things wrong by being too permissive. Here is a counter-example. Suppose that Alex is perfectly reliable about predicting his own behaviour. Alex knows that if he believes that he will drop an anvil on his foot, he will in fact drop an anvil on his foot. Alex knows this because his beliefs about his own behaviour are perfectly reliable indicators. Nonetheless, it is clearly wrong to say that these facts about Alex jointly give Alex a reason to believe that he will in fact drop an anvil on his foot. What is the difference between Alex's case and Alice's earlier on? Alice's belief causes its own truth, whereas Alex's belief is perfect evidence of its own truth. Alex's coming to believe that he will drop an anvil on his foot will not cause him to drop an anvil on his foot. Rather, some disposition to perform a particular action, let us say, causes Alex to believe that he will perform that particular action.

This points the way to a less permissive version of *normative knowledge*:

R4*. *Normative knowledge*: Fact f is a reason for agent a to believe b if f makes it the case that a knows that b will be true, if a 's believing b causes b to be true.

The difference between R4 and R4* is in the last clause. An agent's belief must cause itself to be true ('cause' should be interpreted loosely here) in order for there to be a reason for the agent to have that belief.

One might still worry about R4*, owing to a view in the philosophy of action, *cognitivism*,²¹ which holds that intentions are special cases of beliefs about what one will do. In combination with R4*, cognitivism appears to lead to the conclusion that we have a great many reasons to believe that we will act in certain ways, since having the intention to ϕ under favourable circumstances causes one to ϕ . Nonetheless, cognitivism should not be treated as a source of concern for R4*, as any account of reasons for intending according to the cognitivist understanding must exempt intention-beliefs from having the aim of truth as understood here.

Normative knowledge as spelled out in R4* is consistent with the pair of views that the aim of belief is truth or knowledge and that the aim of belief plays a role in setting the norms for belief. In cases like the multiple fixed-point numbers game, where evidential norms do not help an agent track the truth, *normative knowledge* is not only consistent with this pair of views, but suggested by them. *Normative knowledge* entails the falsehood of strict evidentialism. Strict evidentialism may yet be true, but it is not because it is the sole norm of belief that is consistent with the aim of belief's being truth or knowledge.

4 Knowledge, normative knowledge, and doxastic voluntarism

Normative knowledge, if true, is sufficient in the single fixed-point numbers game to determine what Alice ought to believe, because there is only one belief that satisfies the conditions as set out in R4*. The non-normative cognitive regulatory mechanism, *knowledge*, is sufficient, *ceteris paribus*, in the single fixed-point version of the numbers game to determine what Alice will believe. This can be seen by spelling out *knowledge* more precisely:

²¹ See Setiya (2007) and Velleman (2000a) for discussions of cognitivism. See Langton (2004) for worries about cognitivism.

C2* *Knowledge*: Absent evidence, an agent can and sometimes does form a belief when she knows that having that belief will cause it to be true.

If we accept the simple view about evidence, then in the single fixed-point numbers game, Alice lacks evidence. *Ex hypothesi*, the belief regulatory mechanism that is available to her is *knowledge*, which selects the only available belief that satisfied C2*, the belief that the number on the display will be -2.

Neither *knowledge* nor *normative knowledge* are sufficient in these ways in scenarios that have the form of the multiple fixed-point numbers game. This is because there is no unique solution to the multiple fixed-point numbers game. On the cognitive regulatory side, *knowledge* is left with more than one belief to choose from. On the normative side, *normative knowledge* provides reasons of equal strength for more than one belief.

That *knowledge* and *normative knowledge* suffice and fail to suffice for their respective purposes in the same scenarios is unsurprising. One way of understanding *knowledge* is as being a special case of a general cognitive regulatory mechanism:

C3 *Reasons*: Under normal circumstances, agents involuntarily form beliefs in response to what they take to be their epistemic reasons.

C3 is phrased subjectively in order to avoid mysteries about how external epistemic reasons could play a direct causal role in belief formation. It is also designed to be neutral about what an agent can take to be an epistemic reason. Two constraints on how to understand C3 must be introduced.

The first is that *taking* is not meant itself to be a reflective or conscious process. Agents need not possess the concept of a reason or be able to articulate the considerations on which a belief is formed. This restriction runs contrary to much of the current thought about what it is to be a reason, which holds that for something to be taken as a reason, an agent must be able to articulate it in some way, at least retrospectively, if asked. Current thought may be correct,

although I suspect it is not. My weaker condition does not preclude instances of *taking to be a reason* that meet the stronger and more popular condition. I offer the weaker condition to focus on an idea that I believe enjoys broad consensus. That idea is that cognitive regulatory mechanisms, or at least rational ones, are sensitive to what an agent takes to be her reasons. What counts as being taken as a reason need not be settled here.

The second constraint is that ‘epistemic’ must be understood in some reasonable way. Epistemic reasons cannot be pragmatic reasons (although the same fact could in some circumstances be both an epistemic and pragmatic reason) for belief. I suggest that we understand epistemic reasons as non-evidential reasons that are nonetheless truth-concerned. It strikes me as implausible to suppose that most agents are working with a sophisticated theory of what an epistemic reason is, but barring some interesting further beliefs,²² they do not take the fact that it would be a great thing if p were true to suggest that p is in fact true.

With these constraints on board, *reasons* covers the more common circumstance in which an agent forms beliefs in response to her (apparent) evidence. *Knowledge* is just a special case of *reasons*. Because evidence-based belief formation is the normal circumstance, accepting *reasons* makes doxastic involuntarism an especially attractive view about belief formation. This is because given an evidential state, e , and a proposition, p , an agent will have one of three clear and mutually exclusive kinds of reasons: a reason to believe p , a reason to disbelieve p , or a reason to suspend judgement about p . Because these three kinds of reasons are mutually exclusive, according to *reasons* a well-functioning cognitive regulatory mechanism will always be able to determine a single doxastic state for an agent based on what she takes to be her evidence. Being able to form beliefs voluntarily would therefore be either superfluous or would go against the norms set by the aim of belief. It would be superfluous if the agent chose the evidentially recommended belief state, because *reasons* already would ensure that she forms that belief state. It would go against the norms set by the aim of belief if the agent chose a belief state other than that recommended by the evidence, as doing so would violate the norm set by the aim of belief that belief be truth tracking.

²² E.g. perhaps a Panglossian believes that because this is the best of all possible worlds, the fact that something would be the best state-of-affairs is a reason to believe that it is the actual state-of-affairs.

The situation is different when we turn to the multiple fixed-point numbers game. Here, neither *knowledge* nor *normative knowledge* provides a unique solution. There is, by *normative knowledge*, equally good reason to believe that the number on the display will be 2 and to believe that the number on the display will be -2.²³ *Knowledge*, lacking a unique solution, also does not have a determinate answer as to which belief to form.

In essence, this is a case in which there are no epistemic reasons that favour believing 2 over -2 or *vice-versa*. And, it is also a case in which there is no epistemic reason to suspend judgement, because the rules of the numbers game are set up so as to ensure that suspending judgement is unreasonable. Despite this being the state of affairs, the only way to track truth is to believe one of 2 or -2. Because the rules and operation of the numbers game are transparent to its players, they cannot (or ought not, at any rate) to take themselves to have a reason to believe 2 over -2 or *vice-versa*. Tracking truth dictates that they believe one or the other, but not either one in particular.

This case is a kind of doxastic Buridan's Ass. To avoid starvation, a hungry donkey must choose one of the two bales of hay despite not being able to form a rational prejudice for one over the other. In cases of action, we assume this is possible. The donkey, or rather his human counterparts, has a psychological capacity to choose when the reasons for favouring one choice over another are exhausted. This capacity for brute choice is sometimes regarded as extra-rational. If it is, it is nonetheless a precondition of one's ability to satisfy rational requirements of the form: choose *a* or *b*.

In the multiple fixed-point numbers game, a capacity to choose a belief would be required for a non-self-deluding agent to meet her requirement to track the truth. If we suppose that the right way to understand *normative knowledge* is that an agent has a reason, or (more strongly) is required either to believe 2 or to believe -2,²⁴ then an agent must arrive at one of those two

²³ Note that *normative knowledge* does not generate a reason to believe (2 v -2). Having this disjunctive belief will not cause it to be true that the number on the display will be 2 v -2.

²⁴ To avoid confusion about scope: Reason (B2 v B-2).

beliefs despite having no epistemic prejudice for one over the other. No sound theoretical reasoning will lead to one belief over the other. Some capacity for spontaneous belief formation is required; most plausibly this would be a capacity to choose what to believe.

As I am not seeking to defend the possibility of doxastic voluntarism, let me make it clear that ‘most plausibly’ is not intended to convey unrestricted plausibility. If doxastic voluntarism is impossible or highly improbable, then the issues raised by the multiple fixed-point number game do constitute, as far as I can see, an independent argument for doxastic voluntarism. Rather, if doxastic voluntarism is possible in the relevant senses and not wholly implausible, it looks like the right cognitive capacity for this job.

The upshot is that doxastic voluntarism, if possible, is useful for a truth-tracking agent, at least in these limited circumstances. Indeed, it is more than useful. It looks like the best way for an agent to securely track the truth in cases where *knowledge* and *normative knowledge* apply. This very limited kind of doxastic voluntarism is consistent with the view that the aim of belief is truth (or knowledge) and that the aim of belief plays a role in setting the norms of belief.

5 Leaps of knowledge: from instability to a fixed point and from one fixed point to another

It is possible now to give a more precise account of the sort of doxastic voluntarism that is consistent with, and possibly commended by, the view that the aim of belief is truth and that this aim sets some of the norms of belief. This is a limited kind of doxastic voluntarism, one that not only tracks truth, but ensures knowledge:

D1 *Voluntarism*: An agent can choose her belief just in case three conditions are met. *A*. Evidence does not issue a relevant requirement (either for a belief, disbelief, or suspension of judgement). *B*. The agent knows that her having the belief will cause the belief to be true. *C*. *Normative knowledge* does not issue a reason for just a single belief.

Voluntarism as spelled out in D1 allows for doxastic choice in cases in which the agent is in a position to make a leap of knowledge. It does not permit an agent to make a leap of faith. An agent can only choose a belief that she knows will be true, if believed, not one about which the evidence is merely indecisive.

One may therefore think of the kind of doxastic voluntarism that is consistent with the aim of belief's being truth or knowledge as 'leap-of-knowledge' voluntarism, as opposed to the much stronger 'leap-of-faith' voluntarism that is often the target of doxastic involuntarists.

Whether or not we have leap-of-knowledge voluntaristic capacities seems to me to be an empirical question. If our cognitive regulatory systems have evolved under pressures to track truth, it at least seems possible, given the epistemic utility of being able to choose beliefs in some circumstances, that human cognitive regulatory mechanisms have evolved with a limited capacity to choose beliefs.

One of the limits on our capacity to choose beliefs according to *voluntarism* as spelled out in D1 is an implied trumping rule concerning the operation of our cognitive regulatory mechanisms. This appears in clause A) of D1. An agent can only choose her beliefs when what to believe is not settled by evidence (or apparent evidence). On the normative side, *normative knowledge* is only relevant when evidential reasons are lacking. On the psychological side, *knowledge* is only operative when evidential regulatory mechanisms are not relevant. In the multiple fixed-point numbers game, evidence has no role to play on either the normative or the psychological side. There are other cases in which this is not so. It is interesting to consider them, suspending clause A) for the moment.

5.1 *The power of positive thinking*²⁵

Consider the following situation. Robert visits an internist for his annual physical examination. He receives a phone call from his physician later in the week with some unsettling news.

²⁵ This section's name is borrowed from Harman (1999)

Robert has a rare and dangerous illness that is connected to his brain states. Improbably, it is connected in particular to brain states that encode beliefs about this specific illness. Patients who believe that they will recover from the illness always do. Patients who believe that they will not recover from the illness never recover. And, patients who have no view about the matter never recover. These outcomes have been shown by extensive research, Robert is truthfully advised, to be caused by the brain states that encode beliefs about the illness.²⁶

Having been told this, Robert will realise that he is in a bad way. Not having felt ill or ever having heard of this illness prior to his visit to the doctor, Robert had no belief about whether he would recover from this illness. He now has very strong evidence that his lack of belief will cause him not to recover. This leads him to form the belief that he will not recover, which is itself further and self-sustaining evidence that he will not recover. Of course, if Robert were to come to believe that he will recover, he in fact would. That belief would become evidence for itself and would cause itself to be true.

5.2 Pragmatic encroachment and leaps of knowledge

Voluntarism as spelled out under D1 rules out the possibility of Robert's choosing to believe that he will get better. This is only because of clause A). In cases like the power of positive belief, a direct change in one's belief would also be a leap of knowledge. There are two fixed points for Robert: believing he will recover and believing that he will not recover. Suspending judgement, once he has been informed of his diagnosis and the nature of the disease, is unstable.

If agents can at least choose their beliefs in the limited way allowed by *voluntarism*, are there any grounds for thinking that perhaps a slightly less restrictive version of doxastic voluntarism that would allow Robert to change his belief directly might be the case? This depends on one's reaction to two questions.

²⁶ One can also produce evil demon variants of this example.

The first question is whether leaps of knowledge from fixed point to fixed point are possible under any circumstances. If not, then Robert will need to seek out a hypnotist or find some other indirect way to effect a change in his belief. The second question is under what conditions a fixed point to fixed point leap can occur, if it is possible at all.

I suggest, tentatively, that if one can make a leap of knowledge from one fixed point to another, it is because of a weak kind of pragmatic encroachment. Let us return briefly to the multiple fixed-point numbers game to see the case for this weak kind of pragmatic encroachment. In a slight modification to the original example, the rules as set out previously still apply, but now Alice will win a prize if she believes that the number on the display will be -2 , but she will not win the prize if she believes that it will be 2 .

From an epistemic point of view, nothing has changed in this modified version over the original. If Alice's cognitive regulatory mechanisms are governed solely by truth-tracking considerations, and reasons are only truth-concerned, changing the pragmatic inducements will affect neither the psychological nor the normative balance sheet. Nonetheless, it does not seem unreasonable to think that Alice, with her capacity to choose which number she believes will appear on the display, would choose -2 in order to win the prize. If we grant her even this limited doxastic voluntarism, she will have the ability to make some choices on a pragmatic basis. Her ability to do so will be consistent with the view that epistemic considerations are completely dominant, both psychologically and normatively, in belief formation.

This modified version of the multiple fixed-point numbers game and the power of positive thinking are both cases in which there are multiple fixed points that differ with respect to pragmatic inducements. In the modified multiple fixed-point numbers game, it seems likely that if Alice has any ability to choose her beliefs at all, she will have the ability to do so taking pragmatic inducements into account. This particular instance of weak pragmatic encroachment does not come at an alethic cost.

It is in alethic cost that the power of positive thinking example differs importantly. In order to make a direct change in his beliefs, Robert's choice would have to override the evidential

considerations. This would undermine the lexical dominance of (apparent) evidence in determining belief states implied by clause A).

How high a cost would it be? There is no straightforward answer. It is not clear to me that anything about the aim of belief requires that evidence take priority in cases in which *knowledge* (suitably modified to remove the stipulation that evidence is dominant) could apply. *Knowledge* is motivated by an interest in tracking-truth and gaining knowledge, just as evidential belief regulation is. At the same time, there is a difference in the degree of pragmatic encroachment between the power of positive thinking and the multiple fixed-point numbers game. The former has a larger pragmatic encroachment, because there is an epistemic consideration, evidence, to which one can appeal to determine which fixed point to settle on. There is none in the latter. Thus, pragmatic considerations can do more work in the former case than in the latter, moving from mere, and optional, tiebreakers to considerations that can override properly epistemic cognitive regulatory mechanisms.

6 Conclusion

I have argued in this paper that a limited version of doxastic voluntarism is consistent with and perhaps suggested by the suppositions that the aim of belief is truth or knowledge and that the aim of belief plays a role in setting the norms for belief. That is the most solid conclusion of this paper.

A less solid conclusion is that a limited pragmatic encroachment on cognitive regulation remains consistent with the aim of belief's being truth or knowledge in the modified multiple fixed-point numbers game and cases with a similar structure. That is to say, in those cases in which one is moving from unstable beliefs to a stable one. I have also considered cases in which an agent is at a fixed point, but has other fixed points available. The example discussed in this paper is the power of positive thinking. I have argued that any plausibility of accepting a degree of pragmatic encroachment in the multiple fixed-point numbers game case does not automatically carry over to these fixed point to fixed point cases. Further, accepting that an

agent has a limited capacity to choose her beliefs in the multiple fixed-point numbers game does not imply that she has that capacity when she already has a belief that is a fixed point.

The reason why the two kinds of cases are different both with respect to pragmatic encroachment and doxastic voluntarism is that an agent who starts at a fixed point already possesses evidence that her belief is true. Thus, while consistent with the aim of having individual beliefs track truth (or having them become knowledge), doxastic voluntarism is not required for truth-tracking or knowledge acquisition.

I have also argued that strict evidentialism is not required by the aim of belief's being truth or knowledge and the aim of belief's playing a role in setting the norms of belief. I have argued in the past²⁷ that philosophers who try to connect strict evidentialism with doxastic involuntarism are mistaken. That claim was too strong. Doxastic involuntarism is unproblematically consistent with the aim of belief's being truth or knowledge, if strict evidentialism is true. If strict evidentialism is false, doxastic involuntarism may leave us short of living up to the aim of belief.

²⁷ Reisner (2009)

Works Cited

- Adler, J. (2002) *Belief's Own Ethics*. Cambridge: Bradford Books.
- Blome-Tillmann, M. (2007). The folly of trying to define knowledge. *Analysis* 67, 214-219.
- Harman, G. (1999) Rationality. In *Reasoning, Meaning, and Mind* (pp. 9-45). Oxford: Clarendon Press.
- Hookway, C. (2000). Epistemic norms and theoretical deliberation. In J. Dancy (Ed.) *Normativity* (pp. 60-77). Oxford: Blackwell Publishers Ltd.
- Kelly, T. (2002). The rationality of belief and some other propositional attitudes. *Journal of Philosophical Research*, 110, 163-196.
- Langton, R. (2004) Intention as faith. In H. Steward and J. Hyman (Eds.) *Action and Agency* (pp. 243-258). Cambridge: Cambridge University Press.
- Parfit, D. (2001) Reasons and rationality. In D. Egonsson, J. Josefsson, B. Petersson, and T. Rønnow-Rasmussen (Eds.) *Exploring Practical Rationality* (pp. 17-39). Aldershot: Ashgate.
- Pritchard, D. (forthcoming). The swamping problem. In A. Reisner and A. Steglich-Petersen (Eds.) *Reasons for Belief*. Cambridge: Cambridge University Press.
- Reisner, A. (2007). Evidentialism and the numbers game. *Theoria*, 73, 304-316.
- Reisner, A. (2008). Weighing pragmatic and evidential reasons for belief. *Philosophical Studies*, 138, 17-27.
- Reisner, A. (2009). The possibility of pragmatic reasons for belief and the wrong kind of reason problem. *Philosophical Studies*, 145, 257-272.
- Setiya, K. (2007). Cognitivism about instrumental reason. *Ethics* 117, 649-673.
- Steglich-Petersen, A. (2009). Weighing the aim of belief. *Philosophical Studies*, 145, 395-405.

Shah, N. (2006). A new argument for evidentialism. *The Philosophical Quarterly*, 56, 481-498.

Skorupski, J. (forthcoming). *The domain of reasons*. Oxford: Oxford University Press.

Velleman, D. (2000). On the aim of belief. In *The Possibility of Practical Reason* (pp. 244-282). Oxford: Clarendon Press.

Velleman, D. (2000a). The guise of the good. In *The Possibility of Practical Reason* (pp. 99-122). Oxford: Clarendon Press.

Wedgwood, R. (2002). The aim of belief. *Philosophical perspectives* 16, 267-297.

Williams, B. (1973). Deciding to believe. In *Problems of the Self* (pp. 136-151). Cambridge: Cambridge University Press.

Williamson, T. (2002). *Knowledge and Its Limits*. Oxford: Oxford University Press.

