

The courage of conviction

Sarah K. Paul

Department of Philosophy, University of Wisconsin-Madison, Madison, WI, USA

ABSTRACT

Is there a sense in which we exercise direct volitional control over our beliefs? Most agree that there is not, but discussions tend to focus on control in forming a belief. The focus here is on sustaining a belief over time in the face of 'epistemic temptation' to abandon it. It is argued that we do have a capacity for 'doxastic self-control' over time that is partly volitional in nature, and that its exercise is rationally permissible.

ARTICLE HISTORY Received 6 November 2015; Accepted 17 November 2015

KEYWORDS Belief; volition; temptation; self-control; diachronic rationality

As agents, we sometimes determine for ourselves what we will do. Some of our intentions and actions are attributable to *us*, rather than to some external influence or some automatic, subpersonal mechanism within us. And we have the capacity to be more or less self-governing over certain aspects of our practical lives, in that the extent to which our actions are the product of external forces is sometimes under our control. Can anything similar be said about our relationship to our beliefs?

According to a widespread and appealing view of agency, an action is most attributable to an agent when it expresses that agent's values, higher-order desires, or practical commitments. The idea is that the structure of a person's will is constituted by a network of such attitudes, such that when this network guides thought and behavior, the agent governs herself. It is when a person non-accidentally does what she takes to be most worth doing, or what she wholeheartedly wants to do, or what is most consistent with her plans and self-governing policies, that her agency is fully manifested. When her behavior deviates from what is called for by these lights, in contrast, her agency is diminished. This is sometimes a matter of failing to be an agent altogether, as when she is passively acted upon by psychological or environmental forces that are

CONTACT Sarah K. Paul  skpaul@wisc.edu

© 2015 Canadian Journal of Philosophy

completely external to her will. In other cases, she is akratic; she actively does something other than what is called for by her values and other commitments.

On this view, there is a substantial role for something worth calling 'volition' in constituting oneself as an agent. For one thing, it denies that autonomous agency consists solely in responding cognitively to one's reasons for action by forming normative beliefs and acting accordingly. This is because there is – at least apparently – widespread normative underdetermination in our practical lives. There are countless incompatible things that are all worth doing, and countless ways of doing them that are equally good. Even the most angelic agent who is determined to do what is best must commit to pursuing only some of the valuable things, and to doing them in only some particular ways. There is therefore ample latitude for discretion in one's intentional actions that is consistent with one's normative reasons for action. Where this latitude exists, we face a genuine choice: there are multiple actions or intentions available, each of which would be an expression of the agent's will if that course of action were adopted.

Moreover, the wide middle ground between 'full-blooded' agency and total passivity that is afforded by the phenomenon of akrasia allows for the exercise of self-control. In cases where an agent is assailed by temptation to do something other than what she by her own lights ought to do, it is in part up to her whether she succumbs or succeeds in resisting. When she gives in and takes the second dessert, she is not simply overcome by appetite; there is a real sense in which she chooses to give in. Likewise, when she refrains, this is not merely a hydraulic process of good judgment overcoming the rogue force of desire. Though it is notoriously difficult to flesh out exactly what this capacity is, it is not merely a conative matter of one's strongest desire winning out, nor a cognitive matter of having the correct beliefs. Like determining for ourselves what to do in the face of normative underdetermination, it is a volitional matter.

In contrast, it is almost universally denied that volition plays any direct role in forming or sustaining our beliefs and other doxastic commitments. Some have argued that this is a conceptual constraint on belief, while others see it as a merely contingent psychological claim, but most agree on its truth. As William Alston puts it in his seminal discussion, 'volitions, decisions, or choosings don't hook up with anything in the way of propositional attitude [belief and withholding] inauguration, just as they don't hook up with the secretion of gastric juices or cell metabolism' (1988, 263). The explanation of why this is so is controversial, but a natural thought is that it derives from the absence of the kind of normative latitude in our epistemic lives that is present in our practical lives. There is only one kind of consideration that properly bears on the question of what to believe: considerations that bear on what is true. And because there is a univocal measure of success with respect to any particular belief – namely, truth – there is no sense in which we can actively elect to believe on the basis of epistemic reasons we know to be outweighed, in the way we can elect to

pursue one genuine good while knowing that we ought to pursue a different good. As Joseph Raz writes, 'the weaker reasons are just less reliable guides to one and the same end ... because there is no possibility that the lesser reason for belief serves a concern that is not served better by the better reason there is no possibility of preferring to follow what one takes to be the lesser reason rather than the better one' (2011, 42). It has thus seemed to many that we must be 'passive' in response to (Feldman 2001, 83) or 'at the mercy' (Hedden 2015, 25) of the evidence.

To be sure, it is widely acknowledged that we can voluntarily shape what we believe in many *indirect* ways, and that non-evidential influences frequently affect what we believe. We can curate the evidence we expose ourselves to, or even change our evidence; we can choose how thoroughly to deliberate; we can develop good or bad habits with respect to the possibilities we consider. And if we are at the mercy of our evidence when we are explicitly considering what to believe, we are equally at the mercy of implicit bias, wishful thinking, fatigue, intoxication, and all sorts of other forces that influence how we see the evidence. The denial that volition plays any role in forming or sustaining a belief is specifically the claim that volition plays no *direct* role, unmediated by some distinct action.

My constructive purpose in this paper is to argue that there is more latitude for volition in constructing our doxastic lives than this traditional picture admits. I think there are important cases in which what we believe is directly up to us. This in turn grounds a dimension in which we can be more or less self-governing in our doxastic lives, just as we can in our practical lives. My destructive purpose is to argue against an increasingly influential conception of epistemic rationality, one that emphasizes 'transparency,' on which this kind of doxastic self-governance is condemned as problematically alienated.¹

1. Transparency

The first task is to lay out more precisely the philosophical theses that I wish to challenge. My chief target will be a way of thinking about the nature of belief and epistemic rationality that I will refer to under the umbrella label 'the Transparency View,' although I do not mean to suggest that any particular philosopher has advocated for the view quite as I will present it here. I focus on this view because I find much in it that I agree with, and yet it purports to rule out what I think is the genuine possibility of believing at will. I also find the explanation it offers for why we cannot believe at will to be the most compelling of those I am aware of.² What I ultimately hope to show is that we can preserve much of the insight of the Transparency View without accepting its implications for the relationship between belief and volition.

The starting point is the insight that belief is an attitude that embodies a thinker's take on what is true. According to the Transparency View, this fact about belief is centrally manifested by a conceptual constraint on doxastic

deliberation. The claim is that as a matter of conceptual truth, the deliberative question of whether to believe that P is transparent to the question whether P, in that the former question can only be settled by considerations bearing on the factual question of whether P is true. As Nishi Shah characterizes the thesis, it is manifest in the phenomenology of doxastic deliberation: '... as long as one is considering the deliberative question of what to believe, these two questions must be considered to be answered by, and answerable to, the same set of considerations. The seamless shift from belief to truth is not a quirky feature of human psychology, but something that is demanded by the nature of first-personal doxastic deliberation' (2003, 447). According to Shah, transparency is explained by the fact that belief is a normative concept, such that possessing the concept of belief requires that one accept the norm that it is correct to believe that P if and only if P is true. Because engaging in doxastic deliberation requires that the concept of belief be at least implicitly deployed, the disposition constitutive of accepting this norm of correctness will be activated, compelling the thinker to exclude any considerations that she does not take to bear on whether P is true.

This conceptual constraint on doxastic deliberation purports to rule out the possibility of believing at will. To believe at will, it is not enough that one voluntarily do something in order to bring it about that one has a particular belief. Rather, the belief itself would need to be the conclusion of a process explicitly aimed at issuing in a belief, but that did not appeal only to evidence. For if the thinker does appeal only to questions bearing on the truth of P, then she is simply oriented toward her belief in the ordinary epistemic way; this would not count as believing at will. On the other hand, if she appeals to non-evidential considerations in her effort to form the belief, such as that it would be worthwhile to believe that P, then she has violated a conceptual requirement for being engaged in a process aimed at issuing in a belief. By her own lights, the considerations she has mustered do not settle the question that must be positively answered if she is to believe that P – namely, whether P. Whatever issues from this process would explicitly fail to meet the correctness norm on belief, and thus would not count as that state.³

A second context in which transparency is often said to be manifested is in a thinker's capacity to know what she believes. The idea here is that a rational thinker need not rack her brain or theorize about herself in order to know whether she believes that P. Rather, because a rational thinker's beliefs are settled by what she takes to be true, she can come to know that she believes that P simply by reflecting on her reasons for thinking P is true and avowing her conclusion as her belief (Moran 2001; Boyle 2009). That is, the question 'do I believe that P?' is held to be transparent to the question 'is P true?' On the version of the view I am interested in, there is no inferential step between concluding that P is true and avowing it as one's belief. Rather, the claim is that when all goes well,

the conclusions are not distinct; to judge that P is both to make up one's mind that P and to know that one believes that P.⁴

Perhaps surprisingly, it is meant to be consistent with both of these appeals to transparency (if not exactly highlighted) that most of us are in fact woefully irrational thinkers. We all believe many false things, and our beliefs frequently do not comport with our evidence. We are subject to all sorts of illicit influences like implicit biases, wishful thinking, neuroticism, priming effects, confirmation effects, stereotype threat, and so forth. Knowing this, how can I be justified in taking my beliefs to be settled by my evidence, let alone true? The answer, I take it, is that transparency only holds from the first-person perspective. If I adopt a 'theoretical,' third-personal point of view on myself, conceiving of myself as any other person who is subject to cognitive error and frequently mistaken, the possibility comes into view that my beliefs are false or unsupported by my evidence. But when I adopt the first-personal perspective, my potentially fallible mind is rendered invisible to me; it does not allow me to entertain the possibility that my beliefs might deviate from what fact and evidence require. As Bernard Williams famously argued, reflection about what to believe makes no essential use of the concept of the first person: 'When I think about the world, and try to decide the truth about it, I think *about the world*, and I make statements, or ask questions, which are about it and not about me' (Williams 2006, 67).

This is not to deny that it is possible, and sometimes necessary, to adopt a third-personal perspective on oneself. If my belief that P is recalcitrantly resistant to my take on whether P, I will not be in a position to know of my belief by reflecting on my evidence; I will have to theorize about myself, or pay an analyst to do it for me. And I can certainly reflect in general on my cognitive failings and devise strategies for improving my epistemic success. But the key claim here is that with respect to any particular belief, to adopt a theoretical perspective on that belief is to become alienated and estranged from it – the belief no longer seems to embody *my* take on how things are. We are sometimes so alienated, but the advocates of transparency argue that this is necessarily the exception and not the rule. It is a symptom of a rational breakdown, in that one's judgment has failed to determine one's belief in the way it ought to. Transparency is thus understood as a normative ideal for belief as well as a conceptual constraint on doxastic deliberation.

The transparency of doxastic deliberation and the transparency condition on self-knowledge might seem to be only loosely related. After all, one need not deliberate anew every time one wishes to know what one believes. Indeed, critics of the transparency condition on self-knowledge have objected that deliberating about whether P is no method at all for knowing what one already believes (e.g. Shah and Velleman 2005, 506–508). It can lead one to make up one's mind anew, but it cannot reliably tell one whether one believed that P before deliberation commenced. While true, I think this objection fails to appreciate the scope of the rational ideal that is advocated by proponents of transparency. In

roughly the same way that believing against one's take on the evidence results in estrangement, some have held that alienation also arises for a thinker who takes her past beliefs to have settled for her what she now believes.

The argument for this claim appeals to the kind of answerability involved in believing. Belief is such that a thinker can always be asked for her reasons for believing that P, and even if she has no reason to offer on a particular occasion, she must acknowledge the appropriateness of the question. Matthew Boyle and Pamela Hieronymi have independently pointed out that within this practice of holding a thinker answerable for her belief, it is inappropriate to give an answer in the past tense – 'because I believed that P' or even 'because I consulted the evidence and judged that P' (Boyle 2011, 2013; Hieronymi *manuscript*).⁵ As Boyle puts it, 'The relevant why-question does not inquire into the explanation of [one's] coming, at some past time, to hold the belief in question, except insofar as the subject's knowledge of how he came to hold the belief speaks to the reasonableness of his continuing to hold it now. Our interest is not in his psychological history, but in the present basis of his conviction' (2011, 11). The objection is that if I take my beliefs to be passive states that I occupy now as a result of some past intervention on myself – even a past event of forming a judgment – this is more akin to programming myself than to the exercise of reason that the why-question aims to elicit.

These claims about the structure of answerability have led both Boyle and Hieronymi to conclude that beliefs must be in the metaphysical category of ongoing activities rather than passive states. The idea is that believing that P is a continuous exercise of the capacity to settle for oneself what one believes, generally (though not necessarily) by deeming one's reasons for so believing to be sufficient. At each point in time, the thinker is unalienated from her belief that P only if she *now* actively holds P to be true. And if this is right, then the two faces of transparency are connected after all. On pain of alienation, higher-order questions about what one believes should if possible be treated as deliberative questions that elicit one's current assessment of P. Questions about past beliefs might be of psychiatric or biographical interest, but something has gone wrong if they play any distinct rationalizing role in determining what one now believes.

Admittedly, the notion of 'alienation' invoked in these arguments is somewhat impressionistic, and it is at first difficult to see how the same threat of alienation should arise both from the inability to know of one's beliefs transparently and in the reliance on past judgments. The former signifies a divergence between what the thinker believes and what she takes to be true, but the latter need not. As I understand it, however, the problem with conceiving of belief as the causal product of some event of making a judgment, and thus as standing states rather than activities, also traces back to the infelicity of taking a third-personal perspective on oneself. If beliefs are states that are formed at a particular time by an event of making a judgment, the worry is that judgment becomes a matter of intervening on oneself to bring about a state of believing. And intervening

on oneself is in turn understood as requiring that one adopt a third-personal perspective on oneself, treating one's own mind as a transitive object. It is indeed plausible that any belief I instill or discover from this perspective is one I am alienated from, in that it could only be self-ascribed by way of a referential stipulation that it belongs to the person who happens to be me. If this is correct, then the connection between the two manifestations of alienation is that they each involve a problematic abandonment of the first-personal perspective.

In sum, I have tried in this section to explicate several related philosophical theses (omitting many nuances, I am afraid). These are: (1) believing is necessarily involuntary; (2) reflection about what to believe makes no essential use of the concept of the first person; and (3) beliefs that are opaque to one's current take on the evidence, or that are held on the basis of a past judgment, are a rational failure and exhibit a problematic form of alienation. These claims are each meant to follow from the premise that transparency is both a conceptual constraint on doxastic deliberation and a normative ideal for belief in general. I do not claim that any particular philosopher endorses all of these theses, or the reasoning I have suggested is behind them, but I do think the spirit of this view has wide appeal. I aim in the rest of the paper to deny all three of them.

2. Epistemic temptation

In this section, I will highlight what I take to be a major drawback of belief transparency. Briefly, if the line of thought developed in Section 1 is correct, it turns out that there is no way for a thinker to weather fluctuations in her judgment in a rational manner. If the loss of transparency between belief and evidence is necessarily a rational breakdown, and especially if transparency is a conceptual constraint on treating the question of whether one believes that P deliberately, then a rational thinker must be bound to her present perspective on her evidence even if that perspective is corrupted. The only way to avoid having one's beliefs fluctuate along with such corruptions of judgment will be to engage in estranged self-manipulation.

The kind of fluctuation I have in mind might be called 'epistemic temptation,' in that it is structurally similar to the kind of temptation that afflicts our desires and evaluative judgment. As the ancient Greeks perceptively emphasized, practical temptation often works on us in a way that is akin to suffering an illusion (though I do not mean to be endorsing the view that it is always like this). It distorts an agent's capacity to assess what she genuinely has most reason to do by distorting the perceived value or desirability of some actions or states of affairs. Whereas she would ordinarily take herself to have much better reason to get a good night's sleep than to stay up late watching 'Broad City,' the temptation of watching just one more episode corrupts her judgment and makes the value of watching the episode appear temporarily outsize. We can identify this as a distortion rather than a genuine change of mind by reference to its temporal

profile, which is bookended by the reverse preferences, as well as by the fact that it tends to be caused by ease of opportunity and followed by regret.

Although it is far less frequently discussed,⁶ I think the corresponding phenomenon of epistemic temptation is equally commonplace. As noted earlier, most of us are regrettably vulnerable to non-evidential influences like emotion, peer pressure, priming effects, stereotype threat, and myriad other factors. The very same evidence can strike one as supporting different conclusions depending on how it is presented, the company one is in, or the kinds of emotions one is currently experiencing. When these factors vary over time, a thinker's ability accurately to assess what her evidence supports varies with them, and can lead to a rise or fall in her confidence in a proposition without a change in the evidence she possesses. Of course, not every change in confidence without a change in evidence is a mistake; sometimes, a thinker rightly recognizes that she has incorrectly deliberated and redeliberates to a better conclusion. But some such changes can be properly classified as temptation because they are by nature temporary, caused by emotions that will pass or environmental stimuli that will eventually be escaped.⁷

Although I will ultimately argue against the claim that transparency to evidence is in any way a conceptual constraint on belief, a significant kernel of truth in that view is that it is psychologically very difficult to be clear-eyed about the effects of epistemic temptation. Even more than in the practical case, tempting influences tend to work by directly pervading and distorting the thinker's perception of the truth. They cause some subset of the evidence to become much more vivid while the rest pales, or lead the thinker groundlessly to suspect that she has deliberated unsoundly. They can also affect us at the level of the evidential standards we hold ourselves to, inclining us to alter the threshold we employ in favor of assenting more readily or withholding more strictly. Either way, rather than driving a wedge between an accurate judgment as to what one's reasons support and the akratic belief one actually forms, epistemic temptation directly pervades the thinker's perspective on the reasons themselves. It is almost irresistible to resort to visual metaphors here, as Michael Smith and Philip Pettit do: beliefs that run counter to what fact and evidence require '... will not allow those requirements to remain visible because the offending beliefs themselves give you your sense of what is and your sense of what appears to be' (1996, 448).

Because I am interested in the framework of all-out belief rather than in credences here, I will specifically focus on cases in which a thinker has formed the belief that P at some previous point in time, and later loses confidence in her previous judgment without acquiring new evidence or having specific reason to think her deliberation was flawed. Let us suppose that she is aware of a conflict between how things seemed to her before and how they seem to her now, but that she has no strong evidence one way or the other as to which perspective might be a result of some malfunction. In at least some such cases, the fact of

the matter is that her previous deliberation was sound and her current loss of confidence is a result of epistemic temptation.

Consider an example: suppose that at some point in the past, I deliberated about a philosophical question, considering all the major arguments for and against the possible views. Eventually, I formed the belief that View X is the correct one, thereby coming to believe in the truth of X. But when I arrive at the conference to present on X, my confidence in my previous deliberation plummets (though I gain no specific information concerning a flaw in that deliberation). The arguments in favor of X now strike me as much less forceful than they previously did. Although my time and psychic energy could be better used by concentrating on the next session, I instead spend it by re-opening the question and deliberating anew with the same evidence I previously had, with my insecurity-infused judgment now leading me to abandon my belief in X. Finally, although I previously held that the prestige of a philosopher's home institution is no evidence at all that his or her views are correct, I now perceive the arguments of those with prestigious positions as much more compelling and form the new belief that Y is the correct view.⁸

The outcome of this episode is clearly suboptimal; I have lost a belief that I was entitled to on the basis of my evidence, ended up with a belief my evidence does not support, and wasted time and energy in the process. The interesting question is whether I could have been more autonomous or self-governed than I was: could I have maintained my previous belief throughout the conference, even though it no longer seemed during that time to be true or adequately supported by the evidence? Now, it is uncontroversial that we can exercise what I will call 'self-manipulation' to weather bouts of temptation and continue in the belief we previously had. I could have gone to sleep, drugged myself, or otherwise undermined my capacity to redeliberate. But this is not especially philosophically interesting (nor good conference etiquette); the more pressing issue is whether it is possible to exercise 'doxastic self-control' in a way that is epistemically rational, without strategic self-manipulation.

We are now in a position to see the dark side of transparency. If we are rationally or conceptually bound to treat our beliefs as transparent to our current take on what is true, then there is no way (or no unalienated way) to insulate our beliefs from temporary corruptions of judgment. For this would be a matter of actively maintaining a belief through points in time at which it does not seem true or sufficiently supported by the evidence. It might be that one has most *practical* reason to maintain one's belief, even at the cost of feeling alienated, but this is not something that the standards of epistemic rationality could directly condone.

I think this is a mistake. Before I proceed to argue against this conclusion, though, let me set aside two tempting but inadequate responses to this kind of case. First, one might attempt to assimilate instances of epistemic temptation to the examples typically discussed in the literature on higher-order evidence,

concerning what a thinker should do when she has evidence that she is likely experiencing a cognitive malfunction – she is severely sleep-deprived, or has taken a drug that interferes with normal belief-forming mechanisms. The idea would be that in cases of epistemic temptation, the thinker will frequently have higher-order evidence about her own cognitive functioning that will dictate a rational response – to withhold belief, perhaps, or bracket the deliverances of any potentially compromised thought (e.g. Christensen 2010).

In my view, this response will not do on its own, since cases of epistemic temptation are often relevantly different from cases in which one is drugged or severely sleep-deprived. There will generally have been no obviously compromising event in the former kind of case, or even anything much out of the ordinary, and therefore no clear higher-order evidence about the cause of the conflict between past and present. Whereas drugs and all-nighters are not inescapable facts of everyday life (for those of us who are not rock stars or college students), having emotions and being in social situations are. I am not denying that it is possible in some cases of epistemic temptation to have conclusive higher-order evidence that one's current perspective on the truth is clouded and one's past judgment is therefore more likely to be sound. In such cases, there may be a rational response that is dictated by the combination of one's first-order and higher-order evidence. But I think there are many cases in which the thinker lacks conclusive evidence one way or the other as to which of her perspectives is corrupted. She may have some ambiguous or complex evidence, such as that she is feeling intimidated or buoyant, but this will not always suffice to tell her what to do; we cannot be required to suspend trust in ourselves whenever we are experiencing emotion.⁹

Second, it might seem that there is a simple way to incorporate one's past perspective in such cases while still maintaining transparency: treat one's past judgment that P as testimonial evidence as if from any other person. But while this may in general be admissible, it is not in fact an option in the relevant cases, in which the thinker has called the soundness of her previous deliberation into question. If she were nevertheless simply to accept the deliverances of past deliberation wholesale, treated as the testimony of another person with the same expertise in possession of precisely the same evidence, this would be a matter of continuing to believe that P against her own best judgment. This is clearly epistemically irrational. On the other hand, if she were to treat her past judgment merely as some evidence that P and factor it in to new deliberation as to whether P, then the game is all but lost – her corrupted perspective on the evidence will incline her to draw the incorrect conclusion, especially since the fact that she believed that P will not generally carry much evidential weight. Treating past judgments merely as more evidence will not solve the problem.¹⁰ The trouble is that in cases of epistemic temptation, the thinker should not be making a current evidence-based judgment at all, even one in which higher-order evidence and the testimony of past perspectives is factored in.

3. The diachronic first-person perspective

My view is that it is possible to weather epistemic temptation in a way that is not self-manipulative, rationally impermissible, or alienated. I think that the arguments we have considered to the contrary rest on a mistaken conception of what the first-personal perspective must be.

The mistake is to presuppose that the first-personal perspective is necessarily synchronic, encompassing only the present moment. Although it is never made explicit, the ideal of transparency in fact demands this presupposition. If it were allowed that multiple, conflicting answers to the question 'Is P true?' are first-personally available, the relation between truth and belief could not be conceptually or metaphysically unmediated. And yet, memory does give us access to appearances of the truth that can differ from what now appears to be true. It may even be possible to have psychological access to a future take on what is the case. If these past and future stances were not implicitly ruled out as part of a thinker's first-personal perspective, transparency would fail; the capacity to avow the belief that P on the basis of reflection on evidence would require the mediation of a further answer to the question 'When?'

It may sound odd for me to claim even of the view that believing is an ongoing activity that it is committed to a synchronic first-personal perspective – surely activities are diachronic. But I think it is so committed: as I understand it, it is a time-slice view of believing in which the slices are maximally thin, like stop-animation. At each point in time, the activity view entails that whether the thinker believes that P in an unalienated way is entirely dependent on the synchronic fact as to whether she *now* takes some set of considerations Q to answer the question of whether P (even if in some cases, the consideration is simply 'P'). It does not matter whether she took Q to be a sufficient reason to believe P a moment ago, or whether she will do so in the next moment. Only the present moment matters, or could rationally matter, on pain of becoming estranged from that belief. As Boyle writes, 'I do not recall what I believe about whether P unless I recall what now looks to me to be the truth as to whether P. What I call to mind must be not merely my past assessment of it, but my present assessment of it – the assessment that currently strikes me as correct' (Boyle 2011, 10).

I think we should reject the idea that we are rationally limited to occupying a synchronic, present-directed perspective – that the first person is essentially indexed to 'now'. It is open to me to conceive of myself as occupying a genuinely diachronic first-personal perspective that encompasses past, present, and even future assessments of the truth as potentially my own.¹¹ I do not think that this diachronic self-conception is required in order to be a believer, or to possess the concept of belief, but I do not think it is ruled out either. We saw in Section 1 that according to the Transparency View, possessing the concept of belief involves accepting a norm of correctness to the effect that a belief is correct

only if it is true. Let us grant that this is so. The mistake is to interpret this norm as compelling us to believe whatever the evidence *now* seems to us to require. As a reflective creature, I am in a position to recognize that my capacity to evaluate what is true vacillates over time. I can therefore see that the best way of satisfying the norm of believing P only if it is true may not be always to let my present perspective determine what I believe. A legitimately epistemic concern for the truth does not compel me to conceive of past and future judgments as akin to the testimony of another person, something that could never settle for me what to believe unless I now take the further step of deeming it to be sufficient first-order evidence. Rather, I can consider all of them mine, and therefore candidates for constituting my considered stance on what is true.

If this is right, then we should reject the thesis examined in Section 1: that as a matter of conceptual truth, the deliberative question of whether to believe that P must be settled by reasoning solely about the truth of P. Or rather, we should make a further distinction between the deliberative question of whether to *form* the belief that P and the deliberative question of whether to *continue* to believe that P. The transparency claim may well be true with respect to the first question, but I think it is false with respect to the second. For the diachronic believer, the present moment is not uniquely authoritative in determining what to believe; there is always an implicit question concerning *which* moment in time settles what the diachronic thinker believes. Normally, that question is settled simply by when the thinker takes herself to be in possession of the best evidence (whether or not she is right). But in the cases I am interested in, when P sometimes seems to me to be true and sometimes not without a change in evidence, the question is not settled in this way. It is up to me whether I continue to take my past perspective from which I judged that P as constituting my best judgment, or whether to abandon my belief and deliberate anew, thereby taking my current perspective to have primary authority.¹²

Of course, the truth contained in the doctrine of presentism about the first-personal perspective is that the question of whether to continue believing can only ever arise for the thinker *now*, and must be answered on the basis of considerations that are accessible to her now. Richard Foley dismisses the view I am advocating on this basis:

... insofar as one is deliberating about what to do and think, it is one's current self that is doing the deliberating. This means that conflicts between past opinions and current opinions cannot be treated by me as conflicts between a past self and a current self, where the latter is merely one more part of my temporally extended self. To view the matter in this way is to overlook the banal truth that at the current moment, if I am to have opinions at all, they will be current opinions. Correspondingly, if I am to arbitrate between my current opinion and past opinions, it will be my current self that does the arbitration. (2001, 149)

Foley concludes that 'the solution ... is not to pretend that I can somehow escape from my current perspective. It is to burrow deeper in my current perspective ...'

(151–152). But this simply does not follow. The banal truth cited is merely a formal condition of thought; the fact that the arbitration must *happen* now does not force us to suppose that the *conclusion* of that arbitration must therefore prioritize the present in cases of conflict.

What this shows is that contra Williams, the concept of the first person does make an essential contribution to rational reflection about what to believe. When I think about the world, and try to decide the truth about it, I am at the same time taking a stance on my own ability to evaluate what is true. I am at least implicitly, and occasionally explicitly, answering the question ‘Which of the temporal perspectives that are accessible to me constitutes my take on what is true?’. In cases of epistemic temptation, I can elect to continue to treat my past perspective rather than my current one to speak for me and thereby refrain from changing my mind. It is the fact that I conceive of both the past and the present moment as belonging to my own diachronic first-personal perspective that allows me to avoid believing against my own best judgment.¹³ Instead, I take my best judgment to consist in my earlier assessment rather than in the way things seem to me now. This is the phenomenon I have previously called ‘doxastic self-control.’¹⁴

The possibility of retaining my grasp on my prior assessment of the truth in the face of epistemic temptation may seem ruled out by my earlier observation that epistemic temptation tends to work by directly pervading and distorting the appearance of the truth. I do not mean now to be retreating from this claim; I think epistemic temptation is exceedingly difficult to identify and overcome. To see how it is even possible, it is helpful to think of the phenomenon as a kind of cognitive illusion. Illusions generally do not dissipate merely because one starts to suspect that one is undergoing an illusion. Still, it is something one can in principle recognize and thereby refuse to believe accordingly.¹⁵ In order to do this, the thinker need not actually retain a memory of precisely how the evidence struck her in the past. She merely needs to remember *that* she previously took herself to have sufficient reason to think P is true, and this is precisely the fact that is represented by her belief that P.¹⁶

Still, it is difficult to give a philosophical argument for the psychological possibility of something. At this point, I will resort to aping Alston’s tactic for demonstrating that voluntary control over belief is psychologically impossible: ‘simply ... asking you to consider whether you have any such powers’ (1988, 122). Unlike Alston, I am hoping that you respond with a positive answer. For my part, it seems to me that I do have the power to hold fast to a previously-formed belief even when it currently does not strike me as true, by determining that my present perspective does not speak for me on this matter.

In my view, then, transparency to the present is neither a conceptual nor a psychological constraint. Exercising doxastic self-control is consistent with belief being a state that is regulated for truth, and with accepting the norm that a belief is correct only if true. The capacity for non-manipulative self-control

over one's beliefs is enabled by conceiving of oneself as a thinker who takes a diachronic stance on what is true, rather than occupying only the 'ongoing present.' This self-conception allows one to entertain the possibility that the truth is independent of what seems at any given point to be true, without shifting to a spectatorial third-personal perspective on oneself. It is possible to weather epistemic temptation by taking my past perspective on the evidence, rather than my current one, to be *my* take on what is true.¹⁷

4. Rationality and alienation

Even if I am right that we can choose to believe against our current take on the evidence, the question remains of whether this can be epistemically rational, and whether it need involve a problematic form of alienation.

There are various substantively different notions of rationality we might be interested in here. One such notion concerns an impersonal standard by which we evaluate epistemic states, where this standard is not meant to provide guidance to the thinker herself or obey any kind of 'ought implies can' principle. However important this notion might be, I wish to set it aside here. My concern is not with ideal thinkers, but with the limited epistemic agents that we actually are. To what norms should thinkers like us be disposed to conform, and are they norms that we could reflectively endorse and be guided by? On this personal understanding of rationality, the question is whether we should accept a norm that permits the exercise of doxastic self-control.

An initially tempting thought is that once we reject a time-slice conception of the first-person perspective, with the incumbent commitment to the priority of the present, then the answer is obviously 'yes' – of course it is rational to continue in one's beliefs. To suppose that there is a philosophical puzzle about how one can rationally or autonomously take a past judgment or perspective to be authoritative for oneself might seem to just assume that different temporal stages of a person are rationally equivalent to the viewpoints of different thinkers. On the time-slice view, it is indeed puzzling how a past judgment could influence one's current beliefs without it being a matter of the past self 'programming' the future self. This is what leads us to think that past judgments must be treated as mere testimony. But once we reject this conception of what it is to be a believer, it might seem to follow that there is no philosophical puzzle here – of course a thinker can rationally and autonomously continue to believe what she judged in the past to be true, for no other reason than that the past judgment was *hers*.¹⁸

There is a sense in which I quite agree with this response, but I do not think it suffices in this context merely to point out that we are not time-slice believers. True, we are thinkers that occupy a point of view that is more or less unified over time, but this is not a given or a happy accident – it is an achievement that requires explanation. This is particularly the case in the circumstances under

discussion, in which diachronic unity is threatened by a fluctuating perspective on the evidence. It is easy to slide into time-slice language to describe the problem, couching it as a matter of disagreement between a past and future self, but I think this is a dispensable crutch rather than an illicit assumption necessary to generate the problem in the first place. The philosophical puzzle concerns how a thinker ought to conceptualize and respond to the fact that her reasons seem to her to support different conclusions at different times. Must she always conceive of her past view as a mistake when this happens, or not? Simply pointing out that the past perspective is hers does not solve the problem, since the present perspective is also hers, and yet they conflict. More needs to be said in support of the claim that in cases like these, the present need not always trump.

And this is the relatively modest claim I wish to defend – that as long as the thinker does not take herself to have acquired significant new evidence since forming her belief, she is rationally permitted to respond to a dive in confidence by electing to accept her previous view of the evidence in place of her current view. I do not think it is rationally required, or that it is rationally forbidden to respond by withholding belief or even redeliberating and changing one's mind. There might of course be features of any given case that make one response required; the claim is merely that there is no general requirement prohibiting the exercise of doxastic self-control.

The first thing to note is that the successful exercise of doxastic self-control does not require that one believe against one's epistemic reasons, as I argue in more detail in Paul (2015). A major difficulty for explaining how self-control in the practical case could be rational is that it often does require that an agent follow through on a prior resolution even when she lacks sufficient reason in the moment to do so. This is because the effect of temptation on our desires and evaluative judgments can alter our reasons for action. In contrast, it rarely if ever alters our normative reasons for belief; it affects a thinker's perception of her reasons, but not the reasons themselves. Of course, one could exercise doxastic self-control to maintain a belief that is inadequately supported by one's reasons, but nothing in the phenomenon of epistemic temptation inherently requires this.

Second, assuming that genuine cases of epistemic temptation occur regularly (although how frequent they are will vary from person to person), a disposition to exercise doxastic self-control on at least some occasions will be more conducive to having significant true beliefs than a disposition always to withhold belief or to redeliberate and perhaps change one's mind. In such cases, the thinker's past assessment of the evidence is more likely to have been correct than her current assessment. And the beliefs in question will tend to be significant ones that the thinker has a strong interest in having if true; epistemic temptation is more likely to strike when the stakes are high. I therefore suggest that a rational prohibition on doxastic self-control would be too costly, requiring the thinker to

lose significant true beliefs and transition into what will often be a less accurate doxastic state.

Further, the exercise of doxastic self-control is consistent with at least minimal internalist standards, in that the thinker will still be in possession of her original grounds for believing that P, at least obliquely. It is only possible if she recalls her past belief from memory, and is thus aware that she previously took herself to have sufficient reason to form the belief. She may not be in a position in the moment to grasp her reasons for accepting P as true, but she has access to the fact that by her own lights, there are such reasons.¹⁹ She can thus understand her continence as an expression of her own best judgment, rather than an inexplicable impulse or moment of *akrasia*, in a way that should go some distance toward appeasing the internalist about rationality.

Most importantly, as claimed in Section 3, it is consistent with accepting the norm that a belief is correct only if it is true. But it will be objected that the thinker is by hypothesis not in a position to *know* whether or not her past perspective is more likely to be correct. These cases are interesting precisely because the thinker herself cannot rely purely on evidence to resolve what to do. I must therefore deny that deferring to one's past perspective requires having sufficient reason to believe that one will be correct in doing so in order to be rational, or even sufficient reason to believe that it is likely. This is not as troubling as it might seem. We are generally entitled to form beliefs on the basis of our own first-order assessment of the evidence, without needing some additional reason to believe that we are likely to be correct in doing so. The only option one can be certain will not result in a false belief is withholding judgment. And while withholding judgment is rationally permitted in these cases, I do not think it is rationally required for any but the most Cliffordian risk-averse among us. It is a way of avoiding error, but at the cost of losing many a true and important belief.

Finally, between the options of maintaining a prior belief or abandoning it, a compelling rationale can be given at the level of having a general disposition to maintain. In my view, we form beliefs in order to settle questions for ourselves precisely because we do not have the cognitive resources or the leisure constantly to be reassessing and updating credences. We need to close our deliberations even when we are not certain so that we can move forward in our investigations and act on our conclusions.²⁰ And beliefs do not play this cognitive role simply in the way that a heavy object is subject to inertia; they settle questions for us in part because we treat them as doing so. Creatures who go in for belief must therefore already have the cognitive disposition to maintain beliefs that have been formed. If this is right, then I think the burden is on the opponent to explain why it is rationally impermissible for this disposition to extend to cases in which the thinker herself is unsure which option would be more accurate.

That we are rationally permitted to stick with our beliefs if there has been no change in evidence is a claim that many will already be convinced of. For

instance, it is entailed by the Bayesian principle of Conditionalization, which holds that a thinker's credences at t_1 should match her credences at t_0 conditional on any new evidence she has learned. If she has learned no new evidence, her credences should not change. Conditionalization has its discontents, however, since it conflicts with even mild forms of internalism about epistemic rationality (e.g. Hedden 2015). For instance, it makes no (explicit) provision for forgetting evidence one once possessed; it requires that the evidence on which one's credences are based be cumulative, even if some of that evidence is no longer mentally accessible (Williamson 2000; Titelbaum 2013). More fundamentally, one might think that to accept Conditionalization is just to assume without argument that there must be some rational requirement on how one's present credences should relate to one's past credences, where Conditionalization states what that relation is.²¹ The view I have articulated here addresses only a specific subset of cases to which Conditionalization applies, does not assume that one's credences should in general match up in any way over time, and is one I think even a proponent of internalism can accept.

It might be objected that doxastic self-control as I have characterized it is merely an exercise of practical rather than epistemic rationality, and that all I have succeeded in doing is to point to a particular way in which we can employ our agency to shape our beliefs indirectly. It is uncontroversial that we can control the gathering of evidence, the duration of deliberation, and other mental activities that have bearing on what we end up believing. It may seem that doxastic self-control as I have characterized it is simply one more activity of this type. But I think this would be a misconstrual of what is going on. The question facing a thinker in the grips of epistemic temptation is 'whether to continue to believe that P'. This is just as much a question of doxastic deliberation as the question of whether to form the belief that P; the former, like the latter, is governed by the norm of truth. If the thinker elects not to abandon her belief, this is precisely a matter of continuing to take P to be true. It is not in light of considerations that show belief maintenance to be worthwhile in any other respect. This is the purview of epistemic rationality.²²

Finally, must a thinker who has successfully exercised doxastic self-control in the way I have described stand in an alienated relation to her belief that P? The difficulty with addressing this question is that it is rarely made precise exactly what is problematic about alienation. One thing that is commonly said is that when we are alienated from a belief, we stand in a relation to it that could not be the paradigm case (e.g. Boyle 2011, 9). If this is all that is meant, then I am happy to grant that bracketing one's current perspective on the truth could not be a rational thinker's default stance. One can only exercise doxastic self-control once a belief has been formed, and one can only deliberately form the belief that P if one now takes P to be true. Deliberation about what to believe *ipso facto* accords one's present perspective epistemic authority. It follows that the capacity for doxastic self-control is parasitic on the default authority of the present.

However, as I observed a moment ago, it is equally impossible for thinkers like ourselves to function as though our beliefs are justified only by what strikes us as true in the current moment. We simply have nothing like the cognitive resources to make good on this ideal. And I have argued that even if we did, it would leave many of us epistemically worse off. Further, in exercising self-control to maintain her belief that *P*, the thinker still represents *P* as true. She elects not to change her mind, thereby recommitting herself to *P* and putting herself in a position to avow it as her belief. She is not forced to discover that she still believes, nor does it follow that she stands in a merely spectatorial relationship to her belief. True, she cannot avow her belief on the basis of the evidence as it currently strikes her. But if anything, this is because she is alienated from her present judgment, not from her belief.

It is also important to remember that the notion of alienation has at least one foot in the domain of the political. If I am correct that exercising doxastic self-control is an important route to achieving autonomy and self-governance in our lives, then the burden is on those who deploy the charge of alienation as a critique of the stance I have advocated here to explain further why it is something we should avoid. Let me emphasize that this is not merely an academic issue; I think the phenomenon in question is extremely common, and frequently implicated in events that are misdescribed as exemplifying a lack of willpower. There are numerous examples of differential behaviors between genders, ethnic groups, and socioeconomic classes with respect to 'stick-to-it-iveness' in certain contexts. Girls are far more likely to quit sports teams they have joined than are boys; women more likely to leave STEM field jobs they have started; those in poverty more likely to abandon long-term savings plans and end up in borrowing traps. This is not best understood on the model of supposing that some groups have stronger will-muscles than others, or that those in question suddenly decide that these activities are not in themselves worthwhile. Rather, many such cases might be explained by fluctuation in the agents' beliefs concerning their abilities, the likelihood that the world will cooperate, and so forth. The capacity for doxastic self-control in such situations might be precisely what is needed to maintain confidence that continuing to pursue the activity is a sensible choice. And if this requires a kind of alienation, the alternative is much worse.

5. Belief and volition

I have argued that staying within the first-personal perspective does not conceptually, psychologically, or rationally require that one's belief as to whether *P* be settled by one's present take on *P*. Insofar as I remember having had a different take in the past, either take is a rationally permissible candidate for being *my* settled view on *P*, as long as I have not acquired significant new evidence in the

interim. I will conclude by reflecting on the extent to which this claim makes space for an element to belief that is genuinely volitional.

First, although I have been critical of the extent to which transparency is a rational ideal, a central insight I wish to accept is that a thinker cannot voluntarily believe *P* if she does not take *P* to be true. If she is convinced that not-*P*, she cannot form or sustain the belief that *P* except by manipulating herself. This is the extent to which it is true that belief is not under voluntary control.

This does not mean that we must be passively compelled by the mere appearance of truth, however. A diachronic thinker often has more than one candidate standpoint available to constitute where *she* stands. When she has multiple standpoints available, it is up to her to determine which one constitutes her stance on what is true. In making this determination, she must take herself to be pursuing truth by believing what her best reasons now support; it cannot be done for any other practical reason. But nor can it be done for any particular epistemic reason. It is in part a volitional matter, in whatever sense the choice between multiple normatively underdetermined practical options is volitional.

This claim does not apply to all instances of belief, but it does apply widely. In cases where the thinker has taken no previous stance on the truth of *P*, or in which she has forgotten any stance she took, she cannot but be compelled by the evidence as she now sees it. On the other hand, it may not only be in the exercise of doxastic self-control that volition plays a role in determining what one believes. Whenever we are in a position to recognize that our evidence does not pin down a uniquely rational response, then whatever doxastic state we form will be a matter of going beyond what the evidence seems to us to require.²³ Focusing on the diachronic case is helpful because it is easier to see how the thinker might be able to have multiple takes on the evidence in view, but insofar as this is possible synchronically, then there is room for volition to play a role in such cases as well.²⁴

Still, the volitional latitude I am arguing for here has strict limits. For instance, a tempting but mistaken thought is that it is possible for a thinker to have a policy or set of policies concerning which way to go in various situations. The problem is that policies are by nature general, whereas the truth norm on belief is particular. Even if such a policy were adopted solely on grounds of truth-conduciveness, this would license policies that 'traded off' a few false beliefs for many true beliefs. But trading off in this way violates the conceptual constraint that one must take *each* of one's own beliefs to be true. It is therefore impossible to adopt general policies that directly govern one's beliefs, even with respect to negotiating conflicts in perspective that are not settled (from a subjective point of view) by the evidence. To be sure, one can have policies that govern one's evidential standards, or one's deliberative and assertoric behavior, but such policies concern our beliefs only indirectly. The volitional latitude that arises from being a diachronic believer can have only a particular belief as its object.

That said, it is enough to ground the possibility of a legitimate form of doxastic autonomy or self-governance. Insofar as my beliefs are shaped by my own activity in responding to my evidence, they can be attributed to *me* rather than to some passive, subpersonal mechanism. And while in the moment of epistemic temptation, I will not be in a position to evaluate whether or not I have gotten it right in exercising doxastic self-control, I will often be able to see it retrospectively as an event of correctly resisting the influence of corrupting forces on my beliefs (although I will sometimes come to recognize it as a mistaken exhibition of dogmatism). If I manage to maintain my well-supported belief that completing a PhD in philosophy is a reasonable goal to pursue because I have the ability to write a successful dissertation, even though I am periodically assailed with self-doubt as a result of fear combined with stereotype threat, this will be in part due to me and not just to a fortunate amount of stability in my evidence. It will have been an instance of being not only self-governing with respect to my actions over time, but with respect to my beliefs.

Notes

1. This paper builds on Paul (2015) and is in part intended to be a further development of the view I sketched there, though it departs from the earlier view in some ways.
2. In particular, although discussions of doxastic voluntarism often take Williams (1973) as their starting point, the conception of 'believing at will' at work in that paper is not the one that I think should be defended, for reasons given in Shah (2002).
3. The argument here largely owes to Hieronymi (2013), though I have put some of these claims in ways that I am not sure she would accept.
4. This is importantly different from the notion Williamson (2000) terms 'luminosity'. To claim that a mental state is luminous is to claim that if you are in that state, you are in a position to know that you are in that state. Transparency is more limited; it is meant to hold only in cases where the thinker's belief is appropriately responsive to evidence. Further, it often holds not because the thinker is in a position to detect the mental state she is in, but rather because she is in a position to be in whatever state she thinks she is in.
5. To be clear, Hieronymi is not addressing questions about self-knowledge or even specifically investigating transparency; her interest is in understanding what beliefs would have to be in order to explain the kind of answerability we in fact exhibit.
6. An important exception is Lawlor (2013). Lawlor draws on a case originally offered by Foley (2001).
7. Although the word 'temptation' might suggest that it always involves finding some other belief more desirable than one's current beliefs, this is not an essential feature of the phenomenon. The parallel with practical temptation concerns the distortion of judgment, not the conative aspects.
8. I offer some other examples in Paul (2015). One might object to this example on the grounds that we should rarely if ever believe that a philosophical thesis is true. Perhaps what we ought to do is merely 'accept' the view in the context of

arguing for it at a conference, for instance, or act as if it is true. Though I cannot argue for this here, I think we should not settle for these simulacra of belief. For instance, it is plausible that there is a belief norm of assertion, such that I ought not go to the conference and assert that P unless I genuinely believe that P. The example is admittedly messy, but the real-world examples that I think should interest us generally are.

9. It might also be that after substantial experience with certain types of situations, we have a kind of evidence about how these situations can influence our judgment. This kind of insight can certainly help us deal with epistemic temptation, but I think it is a substantial achievement, and my aim here is simply to deny that it is necessary for weathering epistemic temptation in a rationally permissible way.
10. I make this argument in slightly more detail in Paul (2015).
11. Nagel (1969) argues for a similar view with respect to our desires and values over time.
12. The proponent of transparency might object here that I am interpreting the constraint too narrowly, and that the lesson to be learned from cases of epistemic temptation is simply that we should sometimes take the fact that we believed that P in the past to be a consideration that *now bears* on the truth of P. But if this suggestion is understood as claiming that the fact that I believed that P is evidence for P, in that it makes P more likely, then I think the evidence will generally be far too weak to make a difference – I have had any number of false beliefs in the past. On the other hand, if the suggestion is that it is one of the things that it is relevant to consider in thinking about what to believe, I could not agree more; my objection is to characterizing the relevance in terms of ‘bearing on the truth’ of P. If the bearing-on relation is understood this broadly, then the Transparency View becomes trivial.
13. Talk of ‘taking’ here might suggest a problematic multiplication of psychic entities. But I do not intend to be reifying a little ‘taker’ in the head; all we need here is the ordinary phenomenon of concluding deliberation by answering a question for oneself, however that normally happens. Thanks to Michael Bratman for raising this concern.
14. Weatherson (2008) has also discussed a notion of doxastic self-control that he thinks grounds a kind of voluntarism with respect to some beliefs. But Weatherson’s sense amounts to what I would call *indirect* control of belief, and is in that sense distinct from (though compatible with) mine; it concerns the thinker’s control over which possibilities she considers as live before proceeding to form a belief.
15. There is clearly more to be said here about what is involved in negotiating such situations. I am tempted to say that it involves a kind of cognitive skill or virtue, but I have no specific account to offer. I only wish to deny that it is a matter that is settled by one’s evidence.
16. The memory condition is important here. I do not think it follows from what I have argued here that learning about a past perspective in some other way – by reading about it in what happens to be one’s own diary, say – would suffice. The perspective has to be accessible first-personally.
17. Richard Pettigrew rightly objects that this would be the wrong response if confronted with two other people who disagree with one another about whether P on the basis of the same evidence, and where there is no univocal higher-order evidence about which person is more likely to be right. In this case, it is not permissible to simply pick one of the two people and adopt that belief. But

a central theme of my project is that we are not required to treat our own past and future ‘selves’ as if they were other people.

18. Nefsky and Tenenbaum ([manuscript](#)) advance a similar argument concerning the question of whether one can autonomously act on past intentions.
19. Tenenbaum’s notion of ‘oblique cognition’ is helpful here: ‘An oblique cognition is a representation of a claim or an object ... through which one understands (or seems to understand) that there are reasons to accept that the object is as one represents it, or that the claim is true’ (1999, 894).
20. Holton (2014) accentuates these points.
21. Thanks to Mike Titelbaum for emphasizing this to me in conversation.
22. One might have the worry here that what I am describing is higher-order reasoning, and that this could not rationally result in a first-order attitude. But I am in agreement with the Transparency View in thinking that the deliberative higher-order question cannot be answered by appeal to reasons bearing on whether to have the attitude, and thus that it is in some sense transparent. I am simply arguing that the first-order question need not be settled by one’s current assessment of the evidence. Thanks to John Broome for raising this question to me.
23. If the Uniqueness Thesis is true, then the thinker will never be *correct* in supposing that there are multiple permissible attitudes she can rationally take toward a body of evidence. Nevertheless, it might still be possible for her to do so.
24. Thanks to Richard Pettigrew for helping me to frame this point.

Acknowledgements

I’m grateful to the Massachusetts Institute of Technology, where I was employed as a visiting professor while writing this paper. Audiences at the 2015 Pacific APA, the BRAT conference at UW-Madison, UT Austin, the University of Vermont, the MIT work-in-progress seminar, the Midwest Epistemology Workshop at the University of Missouri, and Flickers of Freedom were all of great assistance. I’m especially indebted to Mike Titelbaum for written comments, countless conversations, and general inspiration. Thanks also to Richard Pettigrew, Matthew Silverstein, Nishi Shah, Berislav Marušić, Richard Moran, Matt Boyle, and Alex Byrne for very helpful comments and conversations.

References

- Alston, William. 1988. “The Deontological Conception of Epistemic Justification.” *Philosophical Perspectives* 2: 257–299.
- Boyle, Matthew. 2009. “Two Kinds of Self-Knowledge.” *Philosophy and Phenomenological Research* 78 (1): 133–164.
- Boyle, Matthew. 2011. ““Making up Your Mind” and the Activity of Reason.” *Philosophers’ Imprint* 11 (17): 1–24.
- Boyle, Matthew. 2013. “Active Belief.” *Canadian Journal of Philosophy* 39 (supplemental issue): 119–147.
- Christensen, David. 2010. “Higher-Order Evidence¹.” *Philosophy and Phenomenological Research* 81 (1): 185–215.
- Feldman, Richard. 2001. “Voluntary Belief and Epistemic Evaluation.” In *Knowledge, Truth, and Duty*, edited by Matthias Steup, 77–90. Oxford: Oxford University Press.

- Foley, Richard. 2001. *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press.
- Hedden, Brian. 2015. "Time-Slice Rationality." *Mind* 124 (494): 449–491.
- Hieronymi, Pamela. 2013. "Believing at Will." *Canadian Journal of Philosophy* 39 (supplemental issue): 149–187.
- Hieronymi, Pamela. *Minds That Matter*. Manuscript.
- Holton, Richard. 2014. "Intention as a Model for Belief." In *Rational and Social Agency*, edited by Manuel Vargas and Gideon Yaffe, 12–37. Oxford: Oxford University Press.
- Lawlor, Krista. 2013. "Exploring the Stability of Belief: Resiliency and Temptation." *Inquiry* 57 (1): 1–27.
- Moran, Richard. 2001. *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton: Princeton University Press.
- Nagel, Thomas. 1969. *The Possibility of Altruism*. Oxford: Oxford University Press.
- Nefsky, Julia and Sergio Tenenbaum. "Extended Agency and the Problem of Diachronic Autonomy." Manuscript.
- Paul, Sarah. 2015. "Doxastic Self-Control." *American Philosophical Quarterly* 52 (2): 145–158.
- Pettit, Philip and Michael Smith. 1996. "Freedom in Belief and Desire." *The Journal of Philosophy* 93(9): 429–449, p. 488.
- Raz, Joseph. 2011. *From Normativity to Responsibility*. Oxford: Oxford University Press.
- Shah, Nishi. 2002. "Clearing Space for Doxastic Voluntarism." *The Monist* 85 (3): 436–445.
- Shah, Nishi. 2003. "How Truth Governs Belief." *Philosophical Review* 112 (4): 447–482.
- Shah, Nishi, and J. David Velleman. 2005. "Doxastic Deliberation." *Philosophical Review* 114 (4): 497–534.
- Tenenbaum, Sergio. 1999. "The Judgment of a Weak Will." *Philosophy and Phenomenological Research* 59 (4): 875–911.
- Titelbaum, Michael. 2013. *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford: Oxford University Press.
- Weatherston, Brian. 2008. "Deontology and Descartes's Demon." *Journal of Philosophy* 105: 540–569.
- Williams, Bernard. 1973. "Deciding to Believe." In *Problems of the Self*, 136–151. Cambridge: Cambridge University Press.
- Williams, Bernard. 2006. *Ethics and the Limits of Philosophy*. 2nd ed. Oxford: Clarendon Press.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.