

How We Choose our Beliefs

Gregory Salmieri, UNC/Chapel Hill

Benjamin Bayer, Colorado College

June 5, 2009

ABSTRACT: Recent years have seen increasing attacks on the "deontological" conception (or as we call it, the "prescriptive conception") of epistemic justification, the view that epistemology guides us in forming beliefs responsibly. Critics challenge an important presupposition of the prescriptive conception, doxastic voluntarism, the view that we choose our beliefs. We assume that epistemic prescriptions are indispensable, and seek to answer objections to doxastic voluntarism, most prominently William Alston's. We contend that Alston falsely assumes that choice of belief requires the assent to a specific propositional content. We argue that beliefs can be chosen under descriptions which do not specify their propositional content, and that these descriptions—which concern the method of inquiry whereby a belief is to be formed—nonetheless specify the features of the belief that make it epistemically responsible to adopt. More generally, we urge that the identity of a belief is not exhausted by its content.

1.

Modern epistemology traces back to Descartes and Locke, who each set out to discover a method—a set of prescriptive norms—by which one could arrive at truth and avoid error, and by which one could assess one's cognitive activities and existing beliefs. This *prescriptive conception* of epistemology and the associated conception of epistemic justification, now sometimes referred to as the "guidance" (or "deontological") conception of justification (Alston 1988a, Plantinga 1992, Goldman 1999), enjoyed near universal acceptance until the 20th century, when it came progressively under attack.¹ Many of these attacks have drawn their strength from the prescriptive conception's (seeming) dependence on a thesis that has fallen into disrepute: *doxastic voluntarism*, "the doctrine that we have extensive control over what we believe"—that "we choose what to believe."² Since "ought" is taken to imply "can," the prescription that we *ought* to adopt a belief only if we have adhered to certain rules seems to presuppose that we can directly control whether or not we adopt the belief—that we *can* "believe it or not." On this view, if doxastic voluntarism is false—if we do not choose our beliefs—then prescriptions about our beliefs would be inappropriate. And doxastic voluntarism is widely thought to have been refuted by Bernard Williams (1970) and William Alston (1988a).

Yet both the prescriptive conception and voluntarism are hard to give up, and we think that it is important that both be maintained. There is a widespread conviction that some beliefs are more appropriate than others, and that this is something which we as believers are in some way responsible for. For example, a popular book on finance tells its readers that "once a stock reaches your price target, unless you get new information [...] you shouldn't think the stock is going much higher";³ in an essay about lesbianism, a teacher reports telling her students that "you shouldn't think that people are bad just because they're gay";⁴ and a character in a 19th

¹ The "deontological conception" is misleadingly named because epistemic guidance need not be deontological in the sense of that term familiar from ethics. Deontological ethical theories hold that the "right" is prior to the "good"—i.e. that certain principles of action are binding on us in themselves, independent of the consequences to which they lead (either in general or in any particular case). But someone could hold the "deontological" conception of justification while still maintaining that we are only obligated to follow epistemic principles because they help guide us to the truth and away from falsehood. Such a view would be teleological rather than deontological, in the senses of these terms familiar from ethics, but still guidance-oriented. Often the "deontological" conception is associated with such concepts as "praise," "blame," "obligation" and "prohibition," but the view of justification in question is broader than that described by these socially-oriented concepts: it involves any kind of "ought" vs. "ought not," whether spoken to others or to oneself. For more on the distinction between deontological vs. teleological conceptions of epistemic guidance, see Zagzebski (1996, 32-47), Goldman (1999, 273-4).

² Foley (1999), 964.

³ Cramer and Mason (2006), 59.

⁴ Gillespie (1999), 171.

Century serial who wants his daughter to find a husband tells his friend, “That can’t be you, Duke [...] I was wrong to think it could. One can’t dispose of other people’s hearts or indulge in cut-and-dried schemes for their futures.”⁵ In each of these cases, a belief is censured and we are given a reason for the censure, which points to some principle to which the authors hold people accountable in their believing. This is a normal and, we would argue, indispensable part of ordinary thought. Epistemology as traditionally conceived aims to identify the overarching principles that guide thought as a whole and license such evaluations.

Consequently, whether or not we retain voluntarism and the prescriptive conception has deep implications for how we understand ourselves as knowers and for every aspect of epistemology as a discipline. To take just one example concerning a theoretical question in epistemology, voluntarism is a key motivation for internalism about justification: if epistemic justification requires that we choose our beliefs rightly, then epistemic justification requires that we must know the justifiers, the grounds of our choice (Alston 1986/1988b, 189). If voluntarism is false, internalism is no longer required. For this reason, leading opponents of voluntarism typically embrace some form of externalism (Alston 1986/1988b; Plantinga 1998).⁶

Several recent writers have attempted to defend the prescriptive conception by decoupling it from any commitment to doxastic voluntarism, which is now widely regarded as a liability. It is important to distinguish between genuine and merely apparent examples of this strategy. In the latter category we place Richard Feldman’s (2001; 2008) advocacy of a “modest deontologism,” according to which one is subject to certain obligations (“role ‘oughts’”) insofar as one fills a role, regardless of whether it is within one’s power to discharge these obligations. A teacher, for example, ought to clarify, and a parent ought to nurture; and he or she can be blamed for failing to fulfill these roles, even if success is impossible for him or her. Likewise, Feldman argues, a believer ought to follow the evidence, and can be blamed for failing to do so, even if doing so is not a matter of choice. We do not find Feldman’s position convincing, though we can only briefly indicate why here. Surely the reason we blame people who are incapable of teaching well for teaching badly is that they did not have to become teachers at all. If it were not up to them whether or not to teach, then they would remain bad teachers, but they would be pitiable, rather than blameworthy.⁷ They could not be blamed (at least as we think that word should be used), because, though the standards by which we assess teachers apply to them in some respects, they do not apply as *prescriptions by which they could guide their teaching*, since (*ex hypothesi*) they are unable to conform to them. We could broaden our usage of the terms “praise” and “blame” to admit such cases, but only by detaching it from the idea of *prescriptive norms*. And this is what Feldman does with these terms and the term “ought.” He argues that “ought” does not imply “can”; and, while this may be correct for

⁵ Cook (1874), 159.

⁶ Sometimes this is called the “normative” or “teleological” conception of justification. We regard this as a misnomer because the (so called) “deontological” conception is at least as normative in a stronger sense as teleological theories, and it can be at least as teleological.

⁷ In response to this sort of objection, Feldman (2008) says there are ways we ought to chew and breathe, even though we do not choose our role as chewers or breathers. But chewing and breathing are under our voluntary control. An example of a “role ought” in which we neither choose the role nor have voluntary control over the way we perform it would be the way in which we “ought” to circulate our blood. There are, of course, norms concerning blood circulation, and we might use “ought” language (and even occasionally the language of praise and blame in a metaphorical way) to refer to them, but these are clearly norms of a very different sort than both those sought by traditional epistemology and ethics. They are not norms that *one follows* but norms by which one assesses that over which one has no control. It is norms of the first sort that we call “prescriptive.” Feldman may reply that he wishes only to decouple these kinds of non-prescriptive norms from doxastic voluntarism. But then he is left without a basis for epistemic praise and blame. Incidentally, the attempt to isolate logically non-prescriptive norms concerning the teleological functions of living organisms from prescriptive norms may itself be suspect. One author (Binswanger 1992b) argues that we form our concept of non-conscious goals by analogical extension from conscious purposes, the kind of norms from which we typically derive prescriptive norms.

some uses of the word “ought,” it is not the case for the prescriptive use that is central to traditional epistemology and ethics.⁸ In embracing a sense of “ought” that does not imply “can” for epistemic norms, we think Feldman is simply abandoning the prescriptive conception of the discipline.

Genuine attempts to preserve the prescriptive conception while abandoning doxastic voluntarism involve maintaining that there are choices connected with belief that allow for prescriptive norms in epistemology, while denying that beliefs themselves are directly chosen. John Heil (1983) and Robert Audi (2002) distinguish between the *state* of believing something and the *act* of forming or acquiring a belief. In Audi’s case, this distinction itself provides sufficient ground for rejecting doxastic voluntarism, since he claims that only acts—rather than states—can be chosen. So while states of belief themselves may not be chosen, acts of belief-formation can be, and there can be prescriptions concerning the acts of belief-formation which bear indirectly on states of belief. The contours of Heil’s position are different, but he and Audi agree that our epistemic obligations recommend certain *procedures* of gathering evidence, such as paying attention to logic, careful observation, identifying and scrutinizing the data on which a conclusion is based, etc.⁹ The account we will defend also identifies such procedures or methods of inquiry as the ones in virtue of which prescriptive norms apply in epistemology, but we dispute the firm distinction Audi and Heil draw between states of belief and acts of belief-formation. Rather, we argue that these acts partially constitute the beliefs themselves in such a manner that choosing to perform these actions amounts to choosing beliefs.

Our main purpose here will be to defend doxastic voluntarism. A few other recent writers have attempted to defend doxastic voluntarism outright. Most prominently, Matthias Steup (2001a; 2008) argues that a compatibilist account of freedom of action can be extended to apply to beliefs as well. However, he does not directly address Alston’s arguments against doxastic voluntarism, and his position depends on compatibilism about doxastic freedom—a position which we worry is not plausible especially in the case of doxastic freedom (Bayer, unpublished). Other recent defenders of voluntarism also rely on compatibilism (Ryan 2003) or defend the voluntariness of only a limited scope of beliefs (Ginet 2001).

Our defense of doxastic voluntarism turns on a notion of belief choice quite different from the one that Williams and Alston envision when they attack voluntarism. They both focus on cases in which one entertains a proposition and then decides whether to accept or reject it, and we agree that this is impossible (in most cases at least). It is evident to introspection that one cannot simply elect to believe that that Mussolini was a fine man, or even that he was once the leader of Italy if one does not believe that already. But to choose a belief is not necessarily to entertain a specific propositional content and assent to it. We will argue that beliefs can be chosen under descriptions which do not include their propositional content, and that it is precisely insofar as they answer to these descriptions that they are subject to prescriptive epistemological norms. This will permit us to defend a prescriptive conception of epistemology that is intimately connected to doxastic voluntarism.

To make this case, we will first present Alston’s argument, the most influential argument against doxastic voluntarism. We will argue that it makes a false assumption about the candidate forms of belief choice. Once we have articulated an alternative candidate that is most relevant to epistemic prescription—one’s choice of a method of inquiry and related cognitive actions—we will show how this candidate nevertheless counts as a

⁸ Others have underscored and elaborated on this contention that “ought” does not imply “can” See Ryan (2003).

⁹ See also Kornblith (1983) for a similar approach.

form of belief choice, contrary to Audi and Heil. Having defended a robust and epistemologically significant role for choice in believing, we will then turn select objections.

2.

Though his argument was advanced later than Williams', most contemporary discussions of doxastic voluntarism begin with Alston (1988a). (It is now widely accepted that not every characterization of belief-choice need involve the contradiction Williams purports to identify.¹⁰) Alston rejects voluntarism for more modest reasons, claiming on psychological grounds that doxastic voluntarism is simply implausible. His argument is based on a taxonomy of ways in which an agent can be responsible for a state of affairs. The taxonomy begins with a basic distinction between two ways in which an agent can be responsible for a state of affairs: through *voluntary control* (which has *basic*, *non-basic*, and *long-range* forms) or through *indirect voluntary influence*.

The agent has *basic* voluntary control over a state of affairs when he can here and now cause it directly by choice, as one can cause one's index finger to move. Alston claims it is obvious that we have no *basic voluntary control* over beliefs. Like Williams, he assumes that such control would have to consist of entertaining a proposition and then electing to believe it. He asks us to just try to choose to believe that the U.S. is still a colony of Great Britain, as one chooses to move one's finger. We couldn't do it, he points out, even if offered a powerful incentive. Likewise, we cannot choose to *reject* a proposition that seems to us obviously *true*, or even to accept or reject a proposition that we are simply uncertain about.

An agent has *non-basic control* over a state of affairs if he has basic control over something that here and now causes the intended state of affairs (as one can flip off a switch by moving one's index finger). Alston thinks we also have no such control over our beliefs, because there is no obvious "switch" that we can flip in order to bring them about.¹¹ Alston considers whether that switch might be the act of inquiry itself, the act of looking for reasons or evidence, (a point suggested by Chisholm (1968)). Alston concedes that we have basic voluntary control over the act of inquiry, but denies that it allows us to choose beliefs, because we cannot know which beliefs will result from the inquiry (and *a fortiori* cannot intend these beliefs).¹²

An agent has *long-range voluntary control* over a state of affairs when he can accomplish something via a long-range project involving different directly voluntary actions on different occasions. Thus we have long-range voluntary control over our finger dexterity because, though we cannot cause our dexterity to increase or decrease here and now, we can voluntarily engage in exercises which will result, over time, in increased dexterity. Alston thinks we have no control of this kind over our beliefs, and that, in any event, it is not the sort of control that is presupposed by the traditional, prescriptive conception of epistemology. Alston acknowledges that individuals can sometimes place themselves in circumstances they deem likely to bring about a belief in a particular proposition, such as subjecting themselves to certain sorts of peer pressure, and he allows that they may occasionally succeed in acquiring beliefs in this manner. But, he argues that this does not constitute a form

¹⁰ See Shah (2002).

¹¹ Feldman (2001) later points out that if the proposition in question is "The switch has been flipped," then one can quite literally choose to flip a switch in order to believe it.

¹² A similar point is made by Feldman (2008, 345) against a proposal raised by Kelly (2002).

of control over our beliefs because the chances of success are not great enough, and because it would be impossible to pursue such projects for a great number of beliefs. Moreover, he points out that this kind of control could not be what makes epistemic justification possible, because any belief formed on the basis of such peer pressure would be quite *unjustified*.¹³

Finally, an agent has exercised *indirect influence*, when, although he had no intention to bring something about, it would not have come about were it not for some action intentionally taken. Thus a man can be held responsible for being overweight or for being frequently late for appointments even though he did not choose either of these things as such, if they are results of his stopping for fast food on his way to his appointments. It is uncontroversial that we can exert this sort of influence over our beliefs by, e.g., being careful to search for additional evidence for particular beliefs, or by developing general habits of being critical about what we accept on authority. But Alston doubts that such influence as we do have is adequate to account for the range of epistemic prescriptions we are wont to make. He considers a definition of justified belief cashed out in terms of beliefs influenced by particular activities or general habits that are themselves justified, but finds it to be both too wide and too narrow. In particular, he suggests that this definition would imply that the member of a primitive tribe who acquires beliefs about ancient tribal history via the unreliable testimony of tribal elders is to blame for his belief. Alston thinks this is implausible, and that even if this individual failed to engage in otherwise virtuous inquiry, no one could blame him for acquiescing in the overwhelmingly influential standards of his culture.¹⁴

The nature of Alston's argument should now be clear: if beliefs are chosen in an epistemologically relevant way then either they must be voluntarily controlled or (perhaps) indirectly influenced, but they are not voluntarily controlled, and the forms of indirect influence to which they may be subject are not epistemologically relevant. Therefore beliefs are not chosen.

3.

¹³ In a case where an individual put himself in a position to believe a particular proposition, and accidentally happened upon evidence that really did support the belief, his intention to form it would not really be doing the relevant work here.

¹⁴ We will not have occasion to address this type of case in connection with other issues that arise in our main text, so we wish to briefly state that though we agree that cultural factors of the sort Alston mentions may mitigate a believer's culpability, we do not agree that they eliminate it. How we should assess Alston's tribe member depends in part on what role these beliefs play in the tribe member's life. Does he believe them idly? Or does he employ them as a basis for actions or inferences—for example, deciding to wage war on the neighboring tribe in retaliation for some alleged ancient offense? In the former sort of case, his epistemic sin may be quite trivial, but then what he has is not much of a belief (we discuss the relation between action and belief below). In the latter case, the transgression is more significant, and he will have ample reason to reconsider the belief, because in using it to infer or deliberate, he'll necessarily bring it into contact with other beliefs, some of which may weigh against it. And this should prompt him to consider the credentials of beliefs on both sides, to ask himself why he believes them. In the case of the belief accepted on mere authority, he'll find that there's not much of a reason. Of course testimony is an indispensable means of knowledge, and there are things we can come to know in no other way. But a believer needs to have reason to think that the testimony is reliable. The tribesman may have reasons to trust the tribal elders about many matters. But it should become clear after only slight scrutiny that their credibility with regard to the distant past is quite limited. Our position here assumes that testimonial justification has a kind of foundationalist structure, which ultimately depends on nontestimonial justification. Alston may dispute foundationalism about testimony as Thomas Reid did. But part of the grounds for this sort of objection is externalist reliabilism, and one of the presuppositions of that reliabilism—mainly rejection of doxastic voluntarism—is the very thesis that's presently in question.

We have already indicated both that we will identify the choices in virtue of which prescriptive norms apply to beliefs as those over the actions by which the beliefs are formed and sustained. We will later argue that these and other, related actions partially *constitute* beliefs. This point, and, more generally, the fact that states can be partially constituted by actions, raises some questions about Alston's distinction between basic and non-basic voluntary control, which we will address in the next section. First we want to reconsider whether a subject's choice of a method of inquiry that results in a given belief might amount to at least non-basic voluntary control, after all, i.e. to the kind of act over which we have basic control that results in an intended state of affairs.

Alston, who (as we've seen) denies that such choices amount to non-basic voluntary control, argues that they only seem to if one ignores "the difference between doing A in order to bring about E, for some definite E, and doing A so that some effect within a certain range will ensue." In order for an inquirer to have

immediate voluntary control over propositional attitudes in the way we do over the positions of doors and light switches, the search for evidence would have to be undertaken with the intention of taking up a certain particular attitude toward a particular proposition. For only in that case would the outcome show that we have exercised voluntary control over *what* propositional attitude we take up (1988a, 271).

Alston gives the example of someone who looks up Al Kaline's lifetime batting average in a baseball almanac—he has the intention of forming some belief about the average, but not the specific belief that the average is .320.

But surely there is a significant sense in which the person intends to believe the *specific* proposition about the batting average that is attested to by the almanac, whatever it turns out to be (and, as it happens, this is the proposition that the average is .320). Far from intending merely to form *any old belief* about the average and coming to the one he does by chance, the man opens the almanac on the premise that one and only one figure will be written next to Kaline's name and with an intention to believe *of that figure* that it is Kaline's lifetime batting average. Similarly, consider the case of a student asked to divide 276 by 12. Presently he has no belief of the form " $276 \div 12 = X$," but he can easily acquire one by going through the method of long division—a process which he knows will yield some determinate quotient, though he does not know in advance which. In choosing to execute the method, therefore, he chooses to form some determinate belief—namely the belief " $276 \div 12 = 23$," though at present he can only specify this belief through some such description as "The belief that results from long division of 276 by 12."¹⁵

In both cases, the methods of inquiry are known to yield unique beliefs and are enacted precisely for this reason. Granted, the subject cannot specify the content of these beliefs in advance, but why should this prevent him from intending to form *those beliefs*? We all the time have intentions with regard to specific contentful items without knowing their specific content: for example, a job applicant who eagerly tears into an envelope from a prospective employer obviously intends to read the letter therein and equally obviously does not already know what it says; similarly one can pick a specific card from a deck fanned out face down in front of one, though one does not know its face value. More generally: the identity of an object of choice is rarely exhausted by the description under which it is chosen, or even by the whole of the chooser's knowledge of the

¹⁵ Obviously the inquirer's intentions serve to delimit the object of his choice. While he may not know the content of the belief he will acquire, he intends to form *some* belief. But it is not true that every unknown consequence of an act of inquiry also counts as an object of choice. The inquirer may know that there will be eventual unintended consequences of attaining that knowledge which he cannot predict, such as settling an unforeseen bet or writing a sports article, but it is not his present intention that any of these come about.

object. Thus, we all the time choose things that *de re* have characteristics about which we are unaware. Each donut in a bakery window has enumerable properties including a precise mass, chemical composition, and causal history (not to mention such “Cambridge properties” as its distance from the center of London) yet an agent’s ignorance of these does not prevent him from choosing one for his breakfast. Alston, then, has not yet given an argument that shows us that, in choosing a method of inquiry, then, we cannot choose what is *de re* a certain belief with a certain content.

Given that beliefs have other attributes besides content with reference to which they may be chosen, in order to determine when we can be said to have chosen a belief in a way that makes us *epistemically responsible for it*, we need to consider how the different attributes of beliefs relate to our reasons as epistemologists for evaluating some beliefs as better than others. Only if the features in virtue of which we choose beliefs are epistemically relevant can we regard the choice as epistemically responsible or not.

Alston ignores this issue by treating all cases of choosing incompletely known objects as cases of being responsible for merely the range of objects marked out by the features that one does know. But this is clearly wrong even outside of the case of belief. Consider his own example:

I am a servant and I am motivated to bring the door into whatever position my employer chooses. He has an elaborate electronic system that involves automatic control over many aspects of the household, including doors. Each morning he leaves detailed instructions on household operations in a computer. Doors can only be operated through the computer in accordance with his instructions. There is no way in which I can carry out an intention of my own to open or close a door. All I can do is actuate the relevant program and let things take their course. Since the employer’s instructions will be carried out only if I actuate the program, I am responsible for the doors’ assuming whatever position he specified, just as I am responsible for taking up some attitude or other towards a proposition within a given range. But I most emphatically am not responsible for the front door’s being open rather than closed, nor can I be said to have voluntary control over its specific position. Hence it would be idle to apply deontological concepts to me vis-à-vis the specific position of the door: to forbid me or require me to open it, or to blame or reproach me for its being open (1988a, 272).

Alston does acknowledge that the servant is responsible for the door’s succeeding or failing to be in the state intended by his employer. But *de re* this state is either that of being open or being closed. Suppose it is the latter. If, because the servant fails to actuate the program, the door remains open and the dog escapes, then he is to blame for this. His job is to actuate the program, which would have closed the door. Although he didn’t know that the program would close the door *per se*, he did know that it would affect the states of all the doors in the house in one and only one way, and he knows the sorts of things that result from doors’ being open or closed, and that the employer, in making decisions such as whether or not to let the dog out of its cage, may be counting on the doors’ being in the position specified by the program.¹⁶ In short, what the servant chooses and is responsible for is not just that the door be in *some* position or other (that will obtain regardless of his action), but that it be in the very position specified by the employer (whether or not the employee can identify this position

¹⁶ This last point, which is important, is obscured in Alston’s example because we are told nothing about the employer’s reason for writing the program—he is treated as some sort of eccentric; but in fact there are many jobs that are essentially like the one Alston describes. For example, the person who sits behind the lighting console in a Broadway theatre, presses buttons on the console at certain points during the play, the buttons initiate programs that turn on and off dozens of lights and alter their brightness. These programs were written by the lighting designer and the person behind the console likely doesn’t know the effect of his job on each individual light. Nonetheless, he knows that if he doesn’t press the buttons when he should, the lighting will be wrong, and that it will be his fault if, because of this, the lead has to deliver a soliloquy in a shadow.

in any terms other than these). Similarly, what's next to Kaline's name in the almanac is not, like Schrödinger's cat, a range of potentialities that collapses into something determinate only when looked at: it is a specific number. Therefore, we maintain, when someone chooses to look the number up (or to perform a long division) what he chooses and is responsible for is the specific belief he comes to and not merely the range within which he antecedently knows the belief to fall.

Of course we cannot always be held accountable for all the *de re* features of everything we choose. Sometimes one is, as Alston thinks, responsible only for the object's falling within a certain range. Consider the case of a blackjack player who is dealt the Jack of Hearts after hitting on a hard twelve when the dealer is showing an ace. The player chooses to receive the next card in the deck, which *de re* (and unbeknownst to her) is a Jack. But she surely cannot be blamed for drawing the jack and thereby losing the game, because she cannot know its face value, which is the only feature of the card which makes it a value or disvalue within the context of the game. The game is such that, in deciding whether to hit, she must regard it as a matter of chance which card she will receive, and so she can be praised or blamed only for her decision to receive a random card rather than none at all. (In this case, she deserves praise for this decision, because the odds are that she will benefit from receiving another card.)

The cases of the black-jack player and the person choosing to actuate the program are alike in that in both the subject knows that something within a certain range will result, without knowing just what in the range will result. But there is an important difference. In choosing to actuate the program (as in the cases of choosing methods of inquiry) the servant knows that where within the range the outcome will fall is determined by some factor that ensures that the outcome will stand in a certain definite relation to his wider goal in making the choice; in blackjack, by contrast, the face value of the next card to be dealt is not (and is known not to be) determined by a factor that ensures that it will stand in any definite relation to the player's goal. The servant's goal in actuating the program—the goal which *ex hypothesi* he is employed to fulfill—is to bring the doors into the positions specified by his employer; he knows that actuating the program will accomplish this. What he does not know or have any control over is whether the program will open or close a given door, and he also (as the example is set up) does not know which position of that door is consistent with his goal of having the doors in the state specified by his employer. The door's remaining open is bad, not in its own right, but because this is not what the employer has specified, and the servant (if he does not actuate the program) is responsible for the door's remaining open, not in its own right, but precisely insofar as it is contrary to the employer's specifications. In sum: the servant chose (what *de re* is) the door's being open in accordance with precisely the description under which this position of the door is blameworthy.

The same point applies in the cases of believing that Al Kaline's lifetime batting average is .320 or that $276 \div 12 = 23$. Epistemological praise or blame attaches to these beliefs not insofar as they have the contents they do, but rather in virtue of their being (or failing to be) well supported by an authoritative method of inquiry. The subject choosing to consult the almanac or to perform long division does not know which statistic or quotient he will arrive at, but he does know that the resulting belief about the statistic will have the authority of the almanac and that the belief about the quotient will be based on long division, and (let us suppose) he knows that beliefs arrived at in these ways are well supported. If so, then he chooses the beliefs in precisely those respects in which they are epistemologically praiseworthy, and he is responsible for the epistemologically good

characteristics that they have.¹⁷ We can praise the belief about the batting average because it is formed in accordance with the prescription to form beliefs about historical facts on the basis of reliable testimony, and the mathematical belief because it is formed in accordance with the prescription to form beliefs about arithmetic on the basis of memorized arithmetic tables and algorithms derived from them.

This conception of doxastic choice and epistemic responsibility is not limited to beliefs formed through some algorithmic method like long division. We form beliefs about the current weather by looking out the windows, about future weather by studying the clouds or consulting meteorologists, about prehistoric times by studying fossils, etc. In most cases both the choice and implementation of a method include innumerable choices of detail, many of them made implicitly, and there is rarely as well-defined procedure as there is in long division. Nevertheless, in all these cases we make choices and these choices are (or ought to be) made for the sake of deriving a specific conclusion. And the beliefs we form are epistemologically assessable not on the basis of their content as such but on the basis of the methods by which they were formed (owing to their truth-aptness). Though we cannot identify in advance the content of the beliefs that these methods will lead to, we do know in advance that they will lead to the formation of beliefs that are likely to be true. (Later we will address the question of whether this sort of control extends to all of our beliefs, and specifically to perceptual beliefs, which Alston rightly emphasizes as a particularly difficult case for doxastic voluntarism.)

So far, we have mainly addressed Alston's contention that choice of the method of inquiry *cannot* count as the choice of belief, as long as the outcome of inquiry is unknown to the inquirer. We have pointed out that there are many contentful states that are chosen without advance knowledge of the content, by reference to other features, and we have suggested that beliefs have, in addition to their contents, other features by reference to which they might (in some sense) be chosen. This shows how beliefs *might* be chosen in the way Alston rules out. Among the features other than content with reference to which one might choose a belief we focused on beliefs' being results of processes of inquiry, and we argued that it is this feature, rather than their content, that is relevant to epistemological assessment and, therefore, relevant to whether we should count as choosing

¹⁷ It is also important to emphasize that the outcome of the choice is not the primary object of *epistemological* evaluation. This might be obscured by the fact that we can evaluate the subject's beliefs *merely* on the grounds of their truth or falsehood. If a subject forms a well-supported but false belief, it might seem that the subject has not chosen the belief in the primary sense in which it is praiseworthy. But notice that something analogous could be said in the case of the servant employed to actuate the program opening and closing doors. The analogy here to a well-supported but false belief would be a situation in which the servant carried out his instructions properly and actuated the program, but (owing to a mistake in the instructions or a bug in the program) the doors end up in an undesirable configuration and the dog escapes. The dog's escape, like the false belief, is bad. But this is irrelevant to our assessment of the servant as an employee. He did his job impeccably and (let us suppose) could not have prevented the escape. Similarly in the case of the well-supported but false belief, if the subject did his "job" as a thinker impeccably, he is not responsible for the false belief that he could not have prevented (or could not have prevented without taking some special precautions that would have been unwarranted).

To put the point another way: on the prescriptive conception of epistemology, epistemic evaluation is primarily evaluation of the subject's adherence to a method. Method-independent evaluations of the results of his actions are possible, but that they are not directly relevant to epistemological evaluation. This said, the prescriptive conception of epistemology need not be "deontological," even though epistemic assessment is a matter of whether a subject adhered to certain prescriptions. The ultimate reason for following these prescriptions may be (and, according to most prescriptive epistemologists, is) that forming beliefs in the prescribed way maximizes the number of true beliefs in the long run. Consequences still matter, though not immediate consequences.

It is also worth noting that a prescriptive conception of justification that treats the consequences for true and false belief as standards of justification need not be mere reliabilism or "rule consequentialism" about epistemological evaluation. It is need not be *merely* because various policies of belief formation lead to forming other true beliefs in the long run that these policies are epistemically responsible to adopt (though this is still one factor). In the same way that virtue ethicists criticize consequentialism for alienating values from valuers, one could criticize reliabilists for alienating beliefs from believers. A belief is not fully *one's own* belief unless it is integrated with the rest of one's beliefs by means of a cognitive method, by means of the known relevant evidence, rather than by means of arbitrary conjecture. Arguably this is a perspective adopted by virtue epistemologists.

our beliefs (and being responsible for them) in the sense required by the prescriptive conception of epistemology. However, for all we have said, it may seem unnatural to think of the fact of having resulted from a certain process of inquiry as a feature with reference to which we can choose a *belief*. Readers may be tempted to conclude that though we do have the sort of responsibility that the prescriptive conception requires, what we choose is not beliefs, but only acts of belief-formation, as Heil and Audi maintain. To defend the stronger claim, we will need to discuss the nature of beliefs as mental states.

4.

It is possible to have different mental states with the same content. This point is most commonly made in connection with the various “attitudes” one can have toward the same proposition: one can believe that it’s raining, hope that it is, fear that it is, etc. Similarly, one can see a particular elephant, imagine it, love it, etc.¹⁸ Two attitudes towards the same content can differ from one another even when they are of the same broad type. For example, the same object can be desired with different degrees of intensity and in a healthy or a neurotic manner.

Similarly two beliefs can differ as mental states, even if they have the same content. Consider four people who believe that low carbohydrate diets are a healthy short-term way to lose weight. The first believes this because he heard it from a co-worker; the second believes it because he has seen a number of friends lose weight on such diets without noticeable negative consequences; the third is a nutritionist who has studied journal articles on the subject, and the fourth is a researcher who conducted an extensive study upon which several of the journal articles were based.

Obviously there is a similarity between the respective mental states enjoyed by the four people, and in many contexts we may describe them as holding the “same belief.” This may lead us to think of this as a case in which an identical effect, the belief, is caused in four different manners in four different subjects. However, when we take seriously the idea that a belief is a mental state, it becomes clear that the shared content of the four subjects’ beliefs does not exhaust the beliefs’ identities, and these identities are inextricably linked to the processes by which the beliefs are formed, and (we should now add) maintained.

Each of the four believers will differ from the others in his degree of commitment to his belief, in the kinds of implications he will draw from it, in the manner in which he will integrate it with new evidence, in the circumstances under which he would revise it, and (therefore) in the way he will act upon it, etc.¹⁹ In short, the same content will function quite differently in the mental lives of the four different subjects, and these differences in functioning will result from (perhaps amongst other factors) the differences in the manners in which the beliefs were formed.

¹⁸ On certain theories, such mental states reduce to attitudes towards propositions. Though we disagree with such views, nothing here turns on the question of whether all content is propositional and all contentful states propositional attitudes. The point is merely that a mental state’s identity is not exhausted by its content. This point is compatible with a variety of views of the nature of states and their content.

¹⁹ Of course, the differences in the formation of the beliefs is not the *sole* cause of these other differences. Such factors as other mental states, character traits, and abilities may also play a role. For example, the person’s desires and his willpower will factor in to his loyalty to the belief in action, and his rhetorical prowess will factor in to his ability to persuade others. But differences in method of approach do not necessarily reduce to differences in background beliefs about methodology. As Lewis Carroll showed us in “What the Tortoise Said to Achilles,” it can’t be that “Whatever Logic is good enough to tell me is worth *writing down*”: formalizing as a belief every aspect of our methodology would, of course, imply an infinite regress.

We can see this if we contrast the person who accepts the low-carb diet belief on uncritical authority from his coworker with the scientist who has researched the issue. Since the basis for the former's belief is merely someone's say-so, not any scientific evidence, his understanding of the requirements of the diet will be vaguer than the scientist's, and his decisions about what to eat will be correspondingly different.²⁰ The uncritical believer will also be less able to recognize counter-evidence if any should present itself and therefore less liable to be dissuaded from his belief by such evidence, but, since he is prone to believe whatever his coworkers say, he will be more likely to be dissuaded by the random comments of others. Moreover, the scientist will likely be more self-conscious of his degree of confidence in the belief and will take this into account when acting on it (or drawing further conclusions from it), whereas the less critical believer will not do this or will do it in a far more approximate and uncritical way.

Though two believers' beliefs may share the same content, the role that content plays in their mental lives will be quite different.²¹ One needn't be a functionalist to acknowledge that beliefs function in our cognitive economies and that their functional properties are either partially constitutive of them, or else follow from differences that are constitutive; in either case, beliefs that function differently must be different beliefs.²² In fact, beliefs with the same content can function quite differently, and one of the factors that both constitutes and determines the way in which the belief functions is the cognitive action by which it was formed, i.e. the method of inquiry employed. The manner in which one comes to believe a proposition has a direct impact on the manner in which one believes it.

Once we notice that, whenever one believes a proposition, there is a *manner* in which one believes it, it becomes obvious that many of our cognitive actions, not only those by which we initially form the belief, directly impact this manner and (therefore) the identity of the belief state. We have already mentioned the issue of whether different believers will notice counterevidence and how they will deal with it; there is also the question of how sensitive they will be to confirming evidence. There is a range of policies one can adopt towards the evidence bearing on a belief: one can remain open to new evidence, and reconsider a belief in light of it when it arises; one can actively seek out new evidence; or, one can actively work to insulate a belief. With respect to each of these policies there are many possible differences of detail and degree. The policy of action one adopts with respect to such evidence constitutes a continuation of the action by which one initially formed the belief (or else, in some cases, a reversal of course). Engaging periodically in such actions is part of believing something.

²⁰ This point holds regardless of the theory of meaning we adopt. On a crude verificationist theory of meaning, the evidence just is the meaning and since the evidence differs, so will the meaning. On an analytic definition view of meaning, we can suppose that fleeting workplace scuttlebutt would be unlikely to convey fully respectable definitions of "low-carb diet" or "healthy short-term weight loss." On an inferentialist theory of meaning, the same workplace gossip would also be unlikely to convey the variety of interconnected inferences about ways to tell if something is low-carb, if one is losing weight in a healthy way, what are the consequences of not following such a diet, etc., all of which inferences collectively constitute its meaning. (Externalist theories of meaning, or of "wide content," can be left aside here. By design these theories are not intended to supplement an account of psychological explanation, but of the semantics of agreement and disagreement, and of theoretical continuity or change. Our focus here is on "narrow content," which is the sort of meaning relevant to understanding what it is to *assent* to a proposition, since assent is clearly a psychological act.)

²¹ And it is not even clear that the content is exactly the same, for in the contrast the scientist's belief the uncritical believer's seems vague and indeterminate even in its content, for the scientist has a more robust conception of what it is to be "low-carb" (and perhaps even to be "healthy").

²² Whether or not one is a functionalist about beliefs (or mental states more generally), the fact that beliefs formed in different ways function differently (even if they share the same content) shows that they are different mental states. However, whereas a functionalist theory would take this difference between the two states as brute, a non-functionalist theory would seek to explain it in terms of some more basic difference between the mental states.

Consider again a belief in the efficacy of a low-carb diet. Many of one's other beliefs will have some bearing on this belief, as will any new evidence that one might acquire. Suppose our believer encounters a published study indicating that subjects losing weight quickly from the low-carb diet have difficulty keeping it off. Advocates of confirmation holism have rightly stressed how, in the face of new evidence such as this, the believer can either reject his previous belief in the efficacy of the diet, or adjust some auxiliary hypothesis to accommodate his original belief.²³ One of the believers in the low-carb diet theory was a researcher who conducted some of the experiments which are cited to confirm the theory. Let's suppose now that his belief in the theory is based not wholly on this data but also in large part on the desire to make a name for himself as an early advocate of a radical new theory. If so, he may be motivated to maintain his pet low-carb diet theory in the face of the new study and might, for example, decide that the measurements of carbohydrate levels involved in the new research were systematically inaccurate ("Really they were eating more carbs than they thought"). (Assuming an understanding of the standards of peer review) a less biased researcher would be less likely to dismiss the accuracy of the data, and more likely either to take it as evidence against the low-carb diet theory or as requiring qualification of the theory. In either case, new mental action is needed to keep the belief: either to dogmatically insulate it from the evidence, or to hold it critically and analyze the specific circumstances under which it holds true.

Action is needed not only to keep or revise a belief in the face of occasional discoveries that challenge it. The very *possession* of a belief requires regular action of another sort which results in frequent occasions for testing it. To count as a belief at all, a mental state must serve (or at least be available to serve) as a premise in one's reasoning (whether practical or theoretical).²⁴ Williams and Alston rightly point out that there is a difference between believing something and merely acting as if one believes it, but, particularly if we include mental actions as actions, it is possible to carry this point to absurdity. If someone consistently acts on a premise, not just in his intercourse with others, but also in the privacy of his own thought—if he regularly accepts conclusions drawn from it, and it figures into the motivation of his actions, then what can it mean to deny that he *believes* it?²⁵

²³ This is not necessarily to endorse the kind of confirmation holism which leads to the underdetermination of theory by evidence thesis. To hold that, one would need to accept not only that multiple revisions were *possible* in the face of new evidence, but that of these multiple revisions, all were equally justifiable.

²⁴ We sometimes speak of believing that P as "taking it that P," and this is no doubt because to believe that P is to take that P *for granted* for the sake of one's reasoning. Timothy Williamson thinks "belief" might be defined along these very lines:

[T]o believe *p* is to treat *p* as if one knew *p*—that is, to treat *p* in ways similar to the ways in which subjects treat propositions which they know. In particular, a factive propositional attitude to a proposition is characteristically associated with reliance on it as a premise in practical reasoning, for good functional reasons; such reliance is crucial to belief. (Williamson 2000, 46-7).

Whether or not one accepts Williamson's acceptance of "knowledge" as a primitive concept in terms of which "belief" is to be defined, the idea that we hold beliefs insofar as we *use* them in inferences is less controversial, and part of the same broad functionalist understanding of belief we have already invoked. Interestingly, Williamson applies his understanding of the nature of belief to the question of doxastic voluntarism:

One attraction of [an account of belief in terms of knowledge] is that it opens the prospect of explaining the difficulty, remarked by Hume, of believing *p* at will in terms of the difficulty of knowing *p* at will (Williamson 2000, 46).

This point cuts both ways, however. If there is a sense in which one *can* know at will, there may also be a sense in which one can believe at will. Doubtless Williamson has in mind the point that we cannot create facts at will, and since knowledge is a "factive mental state," it cannot be willed insofar as facts cannot be willed. We can acknowledge this, but also note that there is such a thing as willful attention to or ignorance of the facts. As we have stressed, the ability to attend to the evidence or ignore or evade it is relevant to the case for voluntarism, whether or not one accepts Williamson's account of knowledge.

²⁵ Both Williams and Alston appeal to the distinction between choosing to believe something and choosing to act as is one does, and Williams distinguishes between choosing to believe something and choosing to deceive oneself (which presupposes believing the opposite). While there is no doubt something right about the distinction between actually believing

These actions, which proceed from the belief and which we would argue partially constitute the state of believing are closely connected with the same actions by which the belief is formed and sustained. The use of the belief generates a stream of new evidence bearing on the belief which the subject will need to deal with, and in light of which he may need to reconsider the belief. To use an earlier example, if one believes that a low-carb diet is a healthy way to lose weight quickly, then one should infer that dieters would suffer no increased risk of osteoporosis. Inferring that dieters should suffer no osteoporosis, we either strengthen or weaken our belief about the low-carb diet depending upon what we discover. Most notably, conclusions drawn from the belief will either be consistent or inconsistent with one another and with other beliefs, and so they will serve either to confirm or disconfirm the belief. Every inference from a belief generates a potential challenge to it, so that in arguing we regularly need to look back to our reasons for our premises as well as forward to their implications.

To summarize: since the possession of a belief is the reliance on some claim as a premise in both theoretical and practical reasoning, and since the act of using it as this premise generates evidence that continues to test the belief itself, holding a belief *means* constantly reevaluating the belief, which requires ongoing choice. There is, then, a whole complex of mental activity which is integral to the possession of a belief. Far from being the mere assent to a content, a belief has a robust identity that cannot be divorced from the actions by which it is formed and sustained. The belief is the assent to and *holding of* a content over time, and this consists (at least partially) in the performance of a host of mental acts on different occasions.²⁶

Heil and Audi's distinction between states of belief and actions of belief-formation is artificial, as, indeed, is the attempt to draw a firm general distinction between actions and states. There are many states other than beliefs what are also partially constituted by processes of action. Consider an army's state of readiness, which is maintained only through acts of preparation such as training exercises, drills, inspections, etc., and partially consists in such actions as certain soldiers' patrolling the periphery of the camp on the lookout for signs of enemy action. Paradoxically, states need not be static—at least not absolutely. The army's readiness counts as a state (and is static) when considered relative to the action for which it is ready (*viz.* fighting), but considered in its own right it may involve constant activity. Generally states of living organisms (and of organizations of organisms, like armies) do involve constant action. Health for example, which is a bodily state, consists in the performance of myriad bodily actions (circulating the blood, engaging in cellular metabolism, etc.) The same is true of mental states, and (we've argued) of beliefs specifically.

a proposition and feigning belief in it (to oneself or others) by acting in certain ways, the points we have made about the relation between beliefs and mental acts suggest that there are other alternatives in addition to these two, and that these alternatives may differ in degree more than in kind. We acknowledged early on that one cannot normally choose a belief in the sense of entertaining a proposition and then deciding whether to believe or reject it. Williams and Alston reject such choices on the grounds that people appearing to make them are really only choosing to act as if they believe something new, not because they are really choosing to believe, relying on the distinction we are now suggesting is less than strict. And there are cases in which people report making direct choices to believe, particularly in cases in which they change their minds about religious views, and other deeply held convictions. Williams and Alston's accounts of what is going on in such cases do not ring true to the phenomena and we have already seen reason to question some of the distinctions on which they are based. Because of this, we think our account has greater resources with which to deal with the sort of cases in question, and exploring such a case will call attention to some of the nuances of our view.

²⁶ This position should not be confused with a pragmatism on which the content of the belief is somehow constituted by the activities in which the believing consists. For all we have said the content may exist mind-independently, and indeed (for the most part) we think that it does. Our point is only that the state of *believing* this content—of standing in a certain relation to it—is a complex one and is (at least partially) constituted by actions taken over time. The state of belief is like that of standing on a balance beam, which involves a succession of small actions. These actions are partially constitutive of the standing, but (of course) they are in no way constitutive of the beam.

At least some of the actions in which the state of believing consists are voluntary, and when these are the ones with reference to which the belief is epistemologically assessable, then there is a significant sense in which we can be said to have chosen the belief and to be responsible for it—the very respect in which the traditional, prescriptive conception of epistemology is committed to saying these things.

In making this argument, we have spoken of individual beliefs as fully distinct mental states, each of which involves various voluntary actions distinct from those involved in other beliefs. This is, of course, a simplification. Our beliefs are woven together into a fabric, such that it can be difficult to individuate one from the others; and, however we individuate the beliefs we will often find that the very same mental actions are involved in multiple beliefs. There are cases where we decide, e.g., to look up a statistic in an almanac, and thereby to form a single belief, but when we decide to read and rely on an encyclopedia's account of a battle, we thereby form innumerable, interconnected beliefs about the battle. Because these interconnected beliefs were not each the result of a special act of inquiry, they might not seem to have been chosen themselves. But if these beliefs are holistically interconnected, if they can't make sense as the kind of belief they are without reference to other beliefs, then they *have* been chosen through an act of inquiry in the same way that any long-range commitments are chosen.²⁷

As we shall see in our next section, considerations of this kind may account for how we choose many of our beliefs, even when we don't think of ourselves as intending to form each and every one of them—for even when we are not seeking to answer specific questions, we are nevertheless constantly confronted by mental managerial choices. Each of us, in every moment and in regard to every thought that crosses his mind, is confronted with choices about whether to raise or lower our level of attention and where to direct it, about when to initiate and sustain processes of inquiry and how to deal with tensions between beliefs (or between old beliefs and new information) when they occur to him.²⁸ The cognitive actions chosen in any given moment, with any given content in mind, are parts of the process of holding any beliefs on which that content bears.²⁹

We cannot normally individuate and enumerate all the beliefs that will be affected by a given choice or identify precisely how they will be affected, but we can characterize the sort of effects that choices of a certain kind will have, and this is all we need to do for the purposes of the prescriptive conception of epistemology, because it is only insofar as they reflect choices of the relevant kinds that our beliefs admit of epistemological evaluation. For example, if a person chooses to push some new piece of information out of mind because it gives him an uncomfortable feeling, he is choosing to hold all the beliefs on which that information bears in willful ignorance of relevant data. Thus, *de re*, he chooses to hold specific mal-informed beliefs. He doesn't choose to hold these beliefs by selecting a content, or even a subject matter, but rather by selecting a way of thinking. Nevertheless, this choice amounts to the selection of specific mal-informed beliefs, and it is *qua*

²⁷ Consider a similar point from W.K. Clifford: "If a belief is not realized immediately in open deeds, it is stored up for the guidance of the future. It goes to make part of that aggregate of beliefs which is the link between sensation and action at every moment of all our lives, and which is so organized and compacted together that that no part of it can be isolated from the rest, but every new addition modifies the structure of the whole. No real belief, however trifling and fragmentary it may seem, is ever truly insignificant; it prepares us to receive more of its like, confirms those which resembled it before, and weakens others; and so gradually it lays a stealthy train in our inmost thoughts, which may someday explode into overt action, and leave its stamp upon our character forever" (1901: 168-9).

²⁸ See Harry Binswanger (1992a).

²⁹ Alston classifies all of the mental managerial choices we've mentioned as cases of "indirect voluntary influence." But, on the understanding of belief that we've developed above, the relationships between these choices and our beliefs is far more intimate than he allows. Indeed, these choices are partially constitutive of our beliefs (considered as mental states).

beliefs held in the light of evaded evidence (not *qua* beliefs on some particular subject matter) that these beliefs are worthy of epistemic criticism.

5.

The preceding amounts to an answer to the question raised by our title. It is an account of how we choose our beliefs—that is, of the choices in virtue of which we are responsible for our beliefs and the respect in which these choices qualify as choices of beliefs rather than merely of actions from which beliefs result. But a significant question remains: do we choose *all* of our beliefs as this account suggests? We have already argued that more cases of belief are like the long division example than might have been obvious at first. While few of our beliefs are formed individually through chosen algorithms, the acquisition and maintenance of at least many of our other beliefs depends on how we choose to conduct our cognition—a realm in which we do (our should) make the choices we do because of the sorts of beliefs that they generate and sustain. But is this true of all beliefs? The most likely class of exceptions are perceptual judgments, e.g.: “This is table,” “It’s brown,” “It’s solid”, etc. (Alston, who emphasizes such judgments, claims that they constitute the majority of our beliefs.) Other plausible exceptions are various everyday inductive inferences (judging “A fire truck is driving by outside” from hearing the siren), the acceptance of testimony on face value (believing a passerby who tells us “It’s raining outside”), simple introspective judgments (like “I’m nauseous”) or the belief in certain mathematical or logical truths (e.g. “Twice two is four” and the principle of non-contradiction). What are we to make of these beliefs?

If there is a class of beliefs that is unchosen and a class that is chosen, then this distinction will be fundamental to epistemology. The chosen beliefs will be subject to justification in the traditional (prescriptive) sense, whereas the unchosen ones will be warranted (if at all) only in some other sense. Epistemologies might develop in different directions from this starting point, perhaps inducing us to some higher order skepticism about the unchosen beliefs until they can be founded on chosen ones, or else treating them as unjustified justifiers (unchallengeable foundations to support other beliefs). For our part, however, we do not think that any of the beliefs in question is entirely unchosen. For reasons of space, we will confine our comments here to testimonial, perceptual, and introspective judgments, leaving discussion of inductive, logical and mathematical beliefs for some other occasion.

So far we have spoken only of beliefs formed by deliberate inquiry, in which a method is self-consciously applied in search of new knowledge. But in (at least many) other cases, beliefs are (or ought to be) formed with some sensitivity to the need for method and standards. This can be seen in the sorts of cases we have already examined, and considering them will shed light on the cases of apparently automatic beliefs currently under consideration. Even if one does not open an almanac with the goal of discovering a certain batting average but merely finds the almanac open on the table and notices the average, one is likely to believe what one reads because one knows that such almanacs are reliable, whereas one is less likely to believe something one notices on the cover of *The Enquirer*. Or, at least a *reasonable* person is more likely to spontaneously form beliefs this way, and this is because he has a general goal of forming (only) beliefs with a

certain character—ones that are supported in a certain way—and so has made habitual certain belief-forming behaviors. Anyone who does not proceed in this way—who believes *The Enquirer*, for example—deserves epistemic blame not just for being credulous in general but also for the specific beliefs he forms in this way. And these beliefs are blameworthy not insofar as they have the specific contents they do, but rather insofar as they were formed in an epistemologically bad way—viz. on the basis of unreliable testimony..

We can understand the epistemological responsibility involved here by noting the previous cognitive processes that might condition a person's current degree of sensitivity. At some earlier stage, one may have formed *general* beliefs—responsibly or not—about what counts as reliable testimony, and because of his commitment to these beliefs, he may now be either sensitive or not to the reliability of sources he comes across passively. These earlier beliefs may themselves have resulted from active inquiry: having discovered his mother was right when she predicted the hot stove would burn his hands and that friends were wrong when they spread rumors that he knew to be false, he wondered about what made the difference between the two types of report, about the qualifications of different types of people to speculate on subject matters of different types. If he made this earlier inquiry, his sensitivity to different types of testimony is improved, and he continues to accept testimony on the basis of this standing policy of sensitivity. In this sense, his sensitivity could itself be chosen, and carry with it holistic commitments that condition further beliefs we form.

Our sensitivity to the reliability of testimony can be seen as analogous to our sensitivity to the reliability of our perceptual judgments. It has often been pointed out that our normal propensities to form perceptual judgments can be affected by our beliefs about the perceptual circumstances. For example if we know we are at a magic show, we would not form the beliefs we otherwise might about solidity of solid-looking hoops or the continuity of seemingly bifurcated women. Again, if we believe that we have ingested hallucinogens or are in the presence of holographers we will form perceptual judgments differently than we normally do.³⁰ Part of the formation and perpetuation of our perceptual beliefs involves the choices we make concerning whether or not we make ourselves alert to all the available data that may bear on the present perceptual situation, and so our perceptual beliefs involve all the choices that are involved in the formation and perpetuation of a host of beliefs about our perceptual situation. In the same way that sensitivity to the reliability of testimony might result from earlier generalizations and processes of inquiry, sensitivity to the reliability of our perceptual judgments might also result from inquiry prompted earlier by the recognition of mistakes in our judgments.

Of course not every perceptual belief is conditioned by sensitivity to conditions of perception. Young children, for example, know little about these conditions and form numerous perceptual beliefs in spite of this. But even in these cases, it is worth considering additional cognitive processes that condition children's acceptance of perceptual beliefs, and the extent to which they might also count as voluntary. In an underappreciated portion of his *Essay*, Locke makes an important point about the voluntariness of our knowledge—which we can just as easily describe as a point about the voluntariness of our beliefs:

There is also another thing in a man's power, and that is, though he turns his eyes sometimes towards an object, yet he may choose whether he will curiously survey it, and with an intent application, endeavour to observe accurately all that is visible in it. . . . Just thus is it with our understanding: all that is voluntary in our knowledge is the employing

³⁰ Alston counts checking to make sure a perceptual situation is normal as a type of voluntary influence that we have over our beliefs. But we have already argued that many of these "influences" are actually constitutive of believing.

or withholding any of our faculties from this or that sort of objects, and a more or less accurate survey of them: but, they being employed, our will hath no power to determine the knowledge of the mind one way or another; that is done only by the objects themselves, as far as they are clearly discovered (Locke 1979/1690, Book IV, Ch. XIII, §1-2).

Even a young child who is not yet sensitive to the conditions of perception still has a choice about “whether he will curiously survey” the objects of perception, and whether he will “with an intent application, endeavour to observe accurately all that is visible in it.” The scope of choice over our perceptual beliefs this permits depends on the extent to which it is in our power to “employ or withhold our faculties.” Locke believes that knowledge consists of the “perception of the agreement and disagreement of ideas,” and later in this passage suggests that as long as the mind receives some ideas from the senses, it cannot help but perceive this agreement or disagreement, and acquire knowledge. But if one disagrees with Locke’s theory of ideas, if one thinks that abstractions do not enter the mind ready-made from perceptual data—if one thinks that the process of abstracting concepts from perceived similarities and differences is itself an active process of cognition—there is even more room for voluntariness than Locke might allow.

The extent to which one makes choices in the process of conceptualization itself is relevant even in cases where norms concerning perception might not apply. The case of an introspective belief like “I am nauseous,” which is arguably indefeasible, is a bit more difficult to analyze than cases of perceptual belief. But note that this belief depends on the possession of the concept “nausea.” When someone asserts that he is nauseous, he is asserting that he is experiencing a sensation of the same sort as he and others have experienced on previous occasions. In relating the state he presently introspects to other states he is relying on observation and memory of himself and others in the past. And, depending on one’s theory of concepts, in grouping his present state together with those past ones, he may even be committing himself to some more substantive theses.

When someone is nauseous or sees a rabbit pulled out of a hat there surely is *something* intimately related to his beliefs over which he has no choice. The relationship between that unchosen thing and beliefs formed is epistemologically significant. Is the relationship one of identity? Of justification? Of exculpation?³¹ This is a substantive question for epistemology and it is well beyond the scope of this short paper. But however this question is to be answered, it is clear that the relationship between perceptual judgments and the unchosen is quite close, and that significantly less choice is involved in these judgments than is involved in, e.g., scientific theories or verdicts in murder trials. This should not surprise us. Given that perceptual judgments form (at least a large part of) the justificatory basis for these other beliefs, it is to be expected that their epistemic standing would be both simpler and more secure. Choice, where present, leads to questions of justification that don’t arise in its absence. So, we should expect that our comparatively simple and secure perceptual judgments would involve less choice than beliefs inferred from them.

However limited and qualified the control we have over our perceptual judgments may be though, we’re suggesting that it is the type of control required by traditional epistemology and thought impossible by Alston and others. But perceptual judgments are a limiting case of beliefs, and even if our suggestion about them fails, we think we have shown how many other beliefs can be, and indeed *are* chosen. It is these other

³¹ To borrow a term from McDowell.

beliefs about which we are most in need of epistemological advice, and we can be held responsible for following or failing to follow this advice, just as traditional epistemologists assumed.

References

- Alston, W. (1988a). The Deontological conception of epistemic justification. *Philosophical Perspectives*, 2, 257-99.
- Alston, W. (1986/1988b). "Internalism and externalism in epistemology." *Philosophical Topics*, 14, 179-221. Reprinted in *Epistemic justification: Essays in the theory of knowledge*. Ithaca: Cornell University Press, 185-226.
- Audi, R. (2001). Doxastic voluntarism and the ethics of belief. (In M. Steup (Ed.), *Knowledge, truth and duty* (pp. 93-114). New York: Oxford University Press.)
- Bayer, B. The Elusiveness of doxastic compatibilism. Unpublished manuscript. <http://www.benbayer.com/doxastic-compatibilism.pdf>
- Binswanger, H. (1992a). Volition as cognitive self-regulation. *Organizational Behavior and Human Decision Processes*, 50 (2), 165-178.
- Binswanger, H. (1992b). Life-based teleology and the foundation of ethics. *The Monist*, 75 (1), 84-103.
- Chisholm, R. (1968). Lewis' ethics of belief. (In P.A. Schlipp (Ed.), *The philosophy of C.I. Lewis* (pp. 223-42). La Salle, IL: Open Court Publishing Company.)
- Clifford, W.K. (1901). *Lectures and essays*, Vol. II. London: Macmillan and Co.
- Cook, E.D. (1874). *Young Mr. Nightingale*, Vol. III. London: Sampson, Low, Marston, Low, and Searle.
- Cramer, J. and Mason, C. (2006). *Mad Money: Watch TV, get rich*. New York: Simon and Schuster.
- Feldman, R. (2001). Voluntary belief and epistemic evaluation. (In M. Steup (Ed.), *Knowledge, truth and duty* (pp. 77-92). New York: Oxford University Press.)
- Feldman, R. (2008). Modest deontology in epistemology. *Synthese*, 161, 339-355.
- Foley, R. (1999). Voluntarism. (In R. Audi (Ed.), *The Cambridge dictionary of philosophy*, 2nd edition (p. 964), New York: Cambridge University Press.)
- Gillespie, P. (Ed.) (1999). *Love makes a family: Portraits of lesbian, gay, bisexual, and transgender parents and their families*. Amherst, MA: University of Massachusetts Press.
- Ginet, C. (2001). Deciding to believe. (In M. Steup (Ed.), *Knowledge, truth and duty* (pp. 63-76). New York: Oxford University Press.)
- Goldman, A. (1999). Internalism exposed. *Journal of Philosophy*, 96 (6), 271-93.
- Heil, J. (1983). Doxastic agency. *Philosophical Studies*, 43, 355-364.
- Kelly, T. (2002). The rationality of belief and some other propositional attitudes. *Philosophical Studies*, 110, 163-196.
- Kornblith, H. (1983). Justified belief and epistemically responsible action. *The Philosophical Review*, 92 (1), 33-48.
- Locke, J. (1979/1690). An essay concerning human understanding. Peter H. Nidditch (Ed.). Oxford: Clarendon Press.

- Plantinga, A. (1992). Justification in the 20th century. *Philosophical Issues*, 2, 43-77.
- Plantinga, A. (1998). *Warrant: The current debate*. New York: Oxford University Press.
- Ryan, S. (2003). Doxastic compatibilism and the ethics of belief. *Philosophical Studies*, 114, 47-79.
- Shah, N. (2002). Clearing space for doxastic voluntarism. *The Monist*, 85 (3), 436-45.
- Steup, M. (2001). Introduction. (In M. Steup (Ed.), *Knowledge, truth and duty*. New York: Oxford University Press.)
- Steup, M. (2008). Doxastic freedom. *Synthese*, 161, 375-92.
- Williams, B. (1970). Deciding to believe. (In H.E. Kiefer and M.K. Munitz (Eds.), *Language, belief, and metaphysics*. Albany: SUNY Press.)
- Williamson, T. (2000). *Knowledge and its limits*. New York: Oxford University Press.
- Zagzebski, L. (1996). *Virtues of the mind*. New York: Cambridge University Press.