St. Cloud State University

# **Doxastic Voluntarism and Epistemic Deontology**

Epistemic deontology is the view that the concept of epistemic justification is deontological: a justified belief is, by definition, an epistemically permissible belief. I defend this view against the argument from doxastic involuntarism, according to which our doxastic attitudes are not under our voluntary control, and thus are not proper objects for deontological evaluation. I argue that, in order to assess this argument, we must distinguish between a compatibilist and a libertarian construal of the concept of voluntary control. If we endorse a compatibilist construal, it turns out that we enjoy voluntary control over our doxastic attitudes after all. If, on the other hand, we endorse a libertarian construal, the result is that, for our doxastic attitudes to be suitable objects of deontological evaluation, they need not be under our voluntary control.

Keywords: Epistemic deontology, doxastic attitudes, doxastic voluntarism, voluntary control, determinism, compatibilism

# 1. Epistemic Deontology and the Argument from Doxastic Involuntarism

Most people believe that, neglecting exceptions, we have voluntary control over what we do. However, do we have as well voluntary control over what we believe, or more generally, over our doxastic attitudes, the attitudes we take towards propositions: belief, disbelief, and suspension of judgment? This question is interesting in its own right, but to epistemologists it is of particular importance, for how we answer it has important implications with regard to the question of what it means for a belief to be justified. According to a venerable tradition in epistemology, epistemic justification is to be understood as epistemic permissibility. Our beliefs are justified if, and only if, what we believe is epistemically permissible for us to believe. And our beliefs are unjustified if, and only if, what we

believe is epistemically impermissible, that is, something we ought not believe.1

This conception of epistemic justification — I will refer to it as *epistemic deontology* — has its critics. According to these critics, epistemic deontology must be rejected because our doxastic attitudes are not under our voluntary control. The argument goes like this:

### The Argument From Doxastic Involuntarism

- (1) If doxastic attitudes are proper objects of deontology, then it is possible to exert voluntary control over them.<sup>2</sup>
- (2) It is not possible to exert voluntary control over doxastic attitudes.

#### Therefore:

(3) Doxastic attitudes are not proper objects of deontology.

This argument has been rather influential, which is in no small part due to the forceful and sophisticated way in which William Alston has stated and defended it in two well-known papers.<sup>3</sup> Many take it to be a decisive, knock-out type of argument. Nevertheless, there is an effective reply to it, the gist of which can be stated as follows. The concept of voluntary control must be given either a compatibilist or a libertarian sense. If we use the concept in its compatibilist sense, we get the result that we enjoy almost unconstrained voluntary control over our doxastic attitudes. If we use the concept in its libertarian sense, the result will be, provided determinism is false, that voluntary control over doxastic attitudes is possible, though rather limited in its extent. However, voluntary control in that sense is not a necessary condition for the applicability of deontology. Thus, however we construe the concept of voluntary control, at least one of the argument's premises is false.

The argument from doxastic involuntarism enjoys widespread assent because its premises are quite plausible. The first premise, an extension of the widely accepted Kantian dictum that 'ought' implies 'can', tells us that if  $\varphi$ -ing is to have

<sup>&</sup>lt;sup>1</sup> Historically, epistemic deontology originates with Descartes and Locke. For advocates of epistemic deontology among contemporary epistemologists, see BonJour (1985), Chisholm (1968) and (1977), and Ginet (1975). For a good account of the deontological tradition in epistemology, see chapter 1 in Plantinga (1993).

<sup>&</sup>lt;sup>2</sup> The kind of possibility I am referring to is not logical but what we might call *psychological* possibility. What I have in mind is the psychological ability to exert control over one's doxastic attitudes.

<sup>&</sup>lt;sup>3</sup> Alston (1989), 91ff, and 115-42, and Baergen (1995), 24-29.

deontological status, the agent must have voluntary over  $\phi$ -ing, that is, must be able both to  $\phi$  and to refrain from  $\phi$ -ing. This demand is certainly not implausible. Consider such things as aging, hair loss, and the growing of one's finger nails. About processes such as these, we don't have an effective choice. They take place whether or not we want them to. It would thus be rather odd to speak of a permission to age, or a duty to avoid hair loss or finger-nail growth. And so it is quite reasonable to insist that proper objects for deontology must be subject to voluntary control, at least under ordinary circumstances.

The first premise, though, is not sacrosanct. First, while accepting the more specific Kantian dictum that 'ought' implies 'can', one might reject the broader principle that  $\phi$ -ing can have deontological status only if it is within the agent's power both to  $\phi$  and to refrain from  $\phi$ -ing.<sup>4</sup> Second, one might not even agree with the weaker Kantian principle.<sup>5</sup> Third, one might accept the first premise in the area of morality, but reject it for epistemic context, a strategy Richard Feldman has employed for the purpose of defending epistemic deontology.<sup>6</sup> However, although I will not exempt the first premise from scrutiny, the main thrust of my argument will be directed against the second premise.

The second premise, then, will be my main target, but I admit that it, too, enjoys a high degree ofplausibility. Many philosophers take the thesis of doxastic involuntarism to be true because actions present us with choice in a way beliefs — or, more generally, doxastic attitudes — do not. I can now effectively choose between raising and not raising my arm. But I have no such choice with regard to my doxastic attitudes. For example, I cannot now choose between believing and disbelieving that four is an even number, that cats are insects, or that there is an odd number of coins in my wallet. I believe that four is an even number, I disbelieve that cats are insects, and I suspend judgment as to the question of whether the number of coins in my wallet is odd. I am stuck with these doxastic attitudes. There is nothing I can now do to change them. Thus William Alston writes:

<sup>&</sup>lt;sup>4</sup>See, for example, Wolf's paper "Asymmentrical Freedom," in Fischer (1986), 151-66, in which she argues that a good action can be praiseworthy even when the agent could not have refrained from performing it, whereas for a bad action to be blameworthy, it must be the case that the agent could have done otherwise.

<sup>&</sup>lt;sup>5</sup> See, for example, Stocker (1971). Alston notes that a case has been made for exceptions to Kant's principle. He claims, however, that the scope of these objections does not include doxastic attitudes. See Alston (1989), 118, note 6.

<sup>&</sup>lt;sup>6</sup> See Feldman (1988).

When I see a car coming down the street, I am not capable of believing or disbelieving this at will; ... when I look out my window and see rain falling, water dripping off the leaves of trees ... I form the belief that rain is falling willy-nilly. There is no way I can inhibit this belief.<sup>7</sup>

I do not wish to deny that Alston makes a good point. When we consider the kind of perceptual beliefs Alston has mentioned, then the phenomenology of belief does look distinctly non-voluntary. Nevertheless, I will argue that the second premise is false.

## 2. Three Concepts of Voluntary Control

The origin of the term 'voluntary' is of course the Latin word 'voluntas', which means 'the will'. Voluntary control, then, is control through the will, control one exerts by exerting one's will. Thus we might initially say the following: I have voluntary control over  $\phi$ -ing if, and only if, I can control whether or not to  $\phi$  by exerting my will. And by 'the will' what is meant here need not be a metaphysically suspect entity, but rather our faculty of making choices and decisions, a faculty that can be analyzed in as many ways as there are theories of our cognitive and mental make-up. Voluntary control thus understood can easily be illustrated with examples. We take ourselves to have voluntary control over raising one's arm, having Pizza for lunch, and holding one's breath for 20 seconds, but not over processes such as aging, hair loss, or holding one's breath for 20 minutes.

This seems pretty straightforward, but getting a grip on the concept of voluntary control is not as simple as enumerating things over which we take ourselves to have, or not have, voluntary control. For there is, after all, the following problem: do we really have voluntary control over our actions (or at least a significant number of them) if we do not have control over the will itself, that is, if we can't control our choices and decisions themselves? Consider arm-raising. Suppose I decide to raise my arm and then execute that decision: I raise my arm. But suppose further that my decision to raise my arm was authored, or controlled, not by me but by you. You happen to have a ray-gun with which you can force me to make certain decisions, and you just used that gun on me: you forced me to decide to raise my arm. Alternatively you could have forced me to refrain from making that decision had you been suspecting that I might make that decision. Now if that were the case it wouldn't be *me* but *you* who exerted voluntary control over

<sup>&</sup>lt;sup>7</sup> Alston (1989), 91f and 129.

whether or not I raise my arm. And thus it is not at all implausible to say that, in order to have voluntary control, we must have what Richard Double calls *dual* control: control not only over the execution of our choices and decisions, but also over our choices and decisions themselves.<sup>8</sup> However, precisely how are we to understand the concept of dual control? This question is intimately connected with the problem of free will and determinism, which presents us with the choice between libertarianism and compatibilism.<sup>9</sup> Next I will briefly state how I conceive of the difference between these two views, and then set forth two compatibilist and one libertarian concept of voluntary control.

Let's begin with determinism. Determinism is the view that all events, including actions and decisions, are caused by earlier events and the laws of nature. But for the purposes of this paper I will work with a weaker notion. Allowing for indeterminacy at the micro-level of physical events, I will take determinism to be the view that the physical context within which actions, choices and decisions take place is one of complete causal determination. Within such a context whatever happens had to happen because it was necessitated by what happened earlier. If determinism is true the action of raising my arm was caused by events that happened before I raised my arm. Given these antecedent events, it had to happen that I raised my arm. One of these antecedent events was my decision to raise my arm. But my decision was in turn just another event caused by antecedent events, and thus it, too, had to occur in exactly the way it did. This means that under the circumstances that prevailed at the time I made my decision I could not have made an alternative decision. Under these circumstances I decided to raise my arm, and I could not have decided to refrain from raising my arm. But if I could not have decided to refrain from raising my arm, the question arises of whether the action of raising my arm was free. A natural response to that question is to deny that it was free. Determinism, then, seems to be in conflict with our prephilosophical belief that we are free agents.

Compatibilits believe that this conflict can be resolved. Some compatibilists reconcile the possibility of free action with determinism by insisting that free action does not require free will. I will refer to this as *simple* compatibilism.

<sup>8</sup> See Double (1991), 10ff.

<sup>&</sup>lt;sup>9</sup> Here I take libertarianism to mean the same as incompatibilism simpliciter: the view that there is no genuine freedom of action unless determinism is false. A libertarian, then might hold that determinism is in fact false, and that freedom of action exists, or alternatively, that determinism is in fact true, and that, consequently, free action is not possible.

According to simple compatibilism, free action is *hypothetical*. An action of mine is free if, and only if, had I decided to do otherwise I could have done otherwise. Other compatibilists, however, accept that something like what we mean by 'free will' is a necessary condition of free action, and try to reconcile freedom with causal determination by insisting that even a causally determined will can be free in the relevant sense. According to this view, if one arrives at a decision 'freely' because it was not the result of things such as neurosis, paranoia, manipulation, brain washing, or other things of this nature, then the 'free will' condition is met in the relevant sense, in spite of the fact that the decision was determined by antecedent conditions and the laws of nature. I will refer to this approach as *refined* compatibilism.<sup>10</sup>

Simple compatibilism was advocated by Hobbes, Locke, and Hume. According to Locke the power to act freely is the "power in any agent to do or forbear any particular action, according to this determination or that of the mind." <sup>11</sup> And Hume said: "By liberty, then, we can only mean a power of acting or not acting, according to the determinations of the will; that is, if we choose to remain at rest, we may; if we choose to move, we also may. Now this hypothetical liberty is universally allowed to belong to everyone who is not a prisoner and in chains." <sup>12</sup> If we apply this idea of free action to the concept of voluntary control, we get the following account:

# Hypothetical Voluntary Control

I have hypothetical voluntary control over doing x if, and only if, I can do x if I decide to do x, and I can refrain from doing x if I decide not to do x.

The corresponding account of voluntary control over doxastic attitudes can be put as follows:

# Hypothetical Voluntary Control Over Doxastic Attitudes

I have hypothetical voluntary control over my doxastic attitude toward p if, and only if, were I to decide to take an alternative attitude toward p, I could take that attitude.

<sup>&</sup>lt;sup>10</sup> For examples of refined compatibilism, see Dennet (1984), Frankfurt (1986), Levin (1979), Neely (1974), Watson (1975), and Wolf (1987) and (1990).

<sup>11</sup> John Locke, Essay Concerning Human Understanding, Book 22, Chapter 21.

<sup>&</sup>lt;sup>12</sup> David Hume, Enquiry Concerning Human Understanding, Section VIII, Part I.

Simple compatibilism succumbs to counterexamples in a fairly uncontroversial way. Consider someone who is brainwashed into handing over half of his monthly income to a sect. Unless we make further relevant assumptions, we have no reason to deny that he enjoys hypothetical voluntary control: had he decided to keep all of his income to himself, he could have. But given that his decision to give up one half of his income was due to brainwashing, it is rather counterintuitive to say that he has voluntary control over the disposal of his income, or that in handing over half of it, he is acting freely. Thus sophisticated compatibilists insist on a deeper account of free action: one that imposes, in addition to hypothetical control, internal conditions intended to ensure that free action is incompatible with brainwashing, hypnosis, manipulation, paranoia, and the like. I will refer to voluntary control that is both hypothetical and internally autonomous in the required sense as *deep* voluntary control. The following formula may serve as a gloss of how the concept of voluntary control can be construed in analogy to the refined compatibilist's conception of free action.

## Deep Voluntary Control

I have deep voluntary control over doing x if, and only if, (i) I can do x if I decide to do x, and I can refrain from doing x if I decide not to do x, and (ii) with regard to deciding whether or not to do x, I enjoy deep internal control.

When applied to doxastic attitudes, a gloss of such theories reads as follows:

# Deep Voluntary Control Over Doxastic Attitudes

I have deep voluntary control over my doxastic attitude toward p if, and only if, (i) were I to decide to take an alternative attitude toward p I could take that attitude, and (ii) with regard to deciding which doxastic attitude to take towards p, I enjoy deep internal control.

Next, let us consider libertarianism. According to this approach, freedom of action requires not merely the ability to do as we will, but also freedom of the will, where freedom of the will is understood not as having deep control over one's will as suggested by refined compatibilism, but as having a will that is *uncaused*. Since the will can be causally determined even if it is internally unimpeded in the sense refined compatibilism requires, a will that is free in that sense does not, according to libertarians, suffice for genuine freedom. Consider once again arm-raising. Libertarians would say I enjoy under the present circumstances freedom with

regard to arm raising if, and only if, I enjoy hypothetical control over it, and in addition under the present circumstances can do both: decide to raise my arm *and* decide not to raise it. And this condition is met only if it is not now causally determined by past events and the laws of nature whether or not I decide to raise my arm.

If we conceive of voluntary control in analogous fashion, then voluntary control is not hypothetical but *categorical*. I now enjoy categorical voluntary control over arm-raising only if I can now decide in favor of *and* against raising my arm, that is, only if which decision I make is not now causally determined. Let us sum up this libertarian concept of voluntary control as follows:

### Categorical Voluntary Control

I have categorical voluntary control over doing x if, and only if, (i) I can decide to do x and decide to refrain from doing x; and (ii) I can do x if I decide to do x, and I can refrain from doing x if I decide to refrain from doing x.

The application of this to doxastic attitudes yields the following account:

# Categorical Voluntary Control over Doxastic Attitudes

I have categorical voluntary control over my doxastic attitude toward p if, and only if, (i) I can decide to take an alternative doxastic attitude toward p; (ii) if I decide to take an alternative doxastic attitude toward p, I can take that attitude.

In preparation for setting forth my reply to the argument from doxastic voluntarism, I will discuss, for each of the three concepts I have distinguished, whether we have voluntary control over our doxastic attitudes, and if we do, to which extent we enjoy such control. First, however, I need to do further conceptual groundwork. In my approach to the issue of voluntary control, two problematic concepts figure prominently: that of making a decision about which doxastic attitude to take, and that of executing such a decision. Next, I will argue that both of these concepts have legitimate application within doxastic contexts.

# 3. Making and Executing Doxastic Decisions

To begin with, let us note that there is a kind of practical decision that does not have an analog in doxastic contexts. Practical decisions can be made in two rather

different ways: by mentally flipping a coin, and on the basis of deliberation. Suppose the menu at the restaurant where you are having dinner specifies vanilla or chocolate ice-cream for dessert, and you like both flavors equally. I will refer to cases of this kind as 'nothing-at-stake' situations. In such situations, the making of a decision is a clearly recognizable, momentary event. Here freedom of the will, if there is such a thing, finds its most definite expression.

To this kind of decision making, there exists no analog in the area of belief-formation. This is so because with beliefs, there always is something important at stake: truth. A false belief, no matter how trivial, is an epistemic failure, a transgression for which, before the court of epistemic reason, a belief's triviality is no excuse. And this is not just an attitude we may or may not adopt, but rather an attitude we have deeply internalized. It is therefore psychologically impossible for us to believe, either for practical reasons or just for the heck of it, what we take or even know to be false. Nor is it possible for us to refrain from believing what we take to be true, for to take something to be true is already to believe it. And thus we cannot, in the way we can choose between vanilla and chocolate ice-cream, choose between believing and not believing that four is an even number, that milk is an alcoholic beverage, or that there is an odd number of coins in one's wallet. Consequently, if mental coin flipping is all there is to voluntary control, voluntary control over our doxastic attitudes cannot be had.

However, there is more to voluntary control than deciding to do a certain thing as the result of mental coin flipping. Sometimes we do something because we deliberated about it and came to the conclusion that, all things told, there are better reasons for than against it. Suppose you are figuring out the route for a weekend trip. Should you choose the shorter route along cross-country roads, or the longer one taking the interstate? You roughly calculate that the interstate, allowing you to travel at a higher speed, will get you to your destination quicker. Taking the country roads will be prettier, but it will also be less safe. Weighing all of these factors, you feel attracted to the safety and convenience of the interstate, and thus render your decision: you're going to take the interstate. This, certainly, is an example of *one* way in which we can make a decision.

Now, just as we can weigh the reasons for and against an action, we can weigh the evidence for and against a proposition. So in addition to practical deliberation, there is epistemic deliberation. But if by engaging in practical deliberation we can decide what to do, shouldn't we by weighing the evidence, that is, by engaging in epistemic deliberation, be able to decide what to believe? It seems to me we should. Consider the O. J. Simpson case and the book Vincent Bugliosi has written about it.<sup>13</sup> In this book, Bugliosi argues (among other things) that Simpson did in fact kill his former wife. Let's suppose that, initially, he simply was baffled by the case and suspended judgment. However, at a certain point, Bugliosi concluded that the evidence pointing to Simpson's guilt left no room for reasonable doubt, whereas the counterevidence the defense team was appealing to could easily be explained away. In arriving at this conclusion, Bugliosi decided to believe that Simpson was guilty. My justification for this claim is an appeal to parity of argument: if you think that practical deliberation enables us to make decisions about what to do, but deny that epistemic deliberation enables us to make decisions about what to believe, then you had better come up with a relevant difference. Now, various ways of establishing a relevant difference have been suggested to me, often with clearly communicated signs of being flabbergasted by the idea that there should be such a thing as doxastic decision making. I will discuss these suggestions shortly, but first I need to indicate briefly what I mean when I talk of 'executing' a doxastic decision.

Consider the moment at which Bugliosi came to a conclusion about what the evidence in the Simpson case supports. At that moment, he arrived at a firmly held belief: Simpson is guilty.  $^{14}$  In which sense can we speak of 'executing a decision' in that context? I execute a doxastic decision, it seems to me, if I take a doxastic attitude *because* I brought an episode of epistemic deliberation to a conclusion with a verdict about what my evidence supports.  $^{15}$  The meaning of 'because' is intended to be causal: I execute a decision, say, to believe that p, if I form the belief that p, and my having concluded that my evidence supports p is a contributing factor in the causal history of that belief.

The execution of epistemic decisions is both analogous and disanalogus to the execution of practical decisions. It is analogous since, just as I execute a doxastic decision if I take a doxastic attitude because I have decided to take it, I execute a practical decision if I act in a certain way because I have decided to act in that

<sup>&</sup>lt;sup>13</sup> Vincent Bugliosi. *Outrage. The five Reasons Why O. J. Simpson Got Way With Murder*. Dell Publishing Company, 1997.

<sup>&</sup>lt;sup>14</sup> I'm simplifying here. Sometimes one tends to have a view on the matter already while sifting through the ebvidence. Of course an epistemicaly responsible agent would, if that were the case, be ready to revise her beliefs should her examination of the evidence indicate the need for this.

<sup>&</sup>lt;sup>15</sup> I am offering only a sufficient condition for the execution of a doxastic decision, for it seems to me that doxastic decisions can be executed as well in ways that do not involve prior deliberation about one's evidence.

way. But there is also disanalogy since, in the practical domain, the execution of a decision often involves at a minimum a muscular effort, and often more than that. Forming a belief involves no such effort. Furthermore, there is disanalogy because a significant amount of time might elapse between the making of a decision and its execution. Ordinarily, this is different in doxastic contexts. Once we have judged that our evidence supports believing that p, we believe that pautomatically and instantaneously. No effort is required, and we certainly are not aware of any act such as 'executing' our decision. However, these differences do not undermine my main point: that one has executed a decision to  $\varphi$  if on  $\varphi$ -s because one has concluded that one's reasons or one's evidence supports φ-ing. Why should temporal distance between the decision and its execution, as well as expending effort, be necessary to decision making? There is nothing conceptually incoherent about cases where a decision is made and executed, but neither of these two conditions are met. For example, it is hardly contentious to claim that you can decide to think, for the next few minutes, about your last summer vacation. And surely this decision might be executed without any lapse in time or the dispensation of effort.

Moreover, it should be noted that the execution of a doxastic decision *can* run into difficulties and thus be accompanied by a temporal interval as well as effort. Consider, for example, someone who grew up in a racist environment but, while at college, for the first time seriously examines his racist beliefs in light of the relevant evidence, and concludes that these beliefs are utterly unsupported. This conversion need not automatically and instantaneously result in appropriate belief revision. Old beliefs may be hard to shake off, and thus our subject might have to exert considerable effort, over a long period of time, to get rid of the beliefs in question.

Let us now turn to objections to my claim that we can make and execute doxastic decisions. One objection I have encountered goes as follows. In some cases of voluntary action, we have introspective awareness of a mental event that can be called *deciding to act*. But we have never any introspective awareness of a mental event that could be called *deciding to believe*. Consequently, there is no reason to suppose that there really is such a thing as deciding what to believe. This objection rests on an unduly narrow conception of decision making. I agree that, in practical contexts, we can be introspectively aware of a decision. This is possible, for example, in what I have called *nothing-at-stake* situations, or when we decide to do something that goes against our inclination, or when an act of practical

deliberation is compressed into a short moment at which we are clearly struck by the necessity of performing a certain action. However, certainly practical deliberation can be a long, arduous, and frequently interrupted process. Under such circumstances, it might not be possible to have introspective awareness of a distinct mental event of deciding to act. Nevertheless, deliberation has taken place, and a decision was made. Thus it appears doubtful that introspective awareness of one's decision as a distinct mental event is a necessary condition of decision-making. A weaker condition is that, after the decision is made, one can then become introspectively aware of the fact that one has made a decision. But that condition is one doxastic decisions can meet. Having concluded that the evidence clearly indicates Simpson's guilt, surely Bugliosi introspectively could have become aware of having made the decision to believe that Simpson was guilty. Nor is it clear to me that the point in question should be granted: that it is impossible to be introspectively aware of an act such as *deciding to believe*. I will get back to this question momentarily.

Secondly, it might be objected that reference to doxastic decisions can easily be explained away. According to the objection, we can restate

- (1) Bugliosi decided to believe that Simpson was guilty more elegantly, and philosophically more accurately, as
- (2) Bugliosi decided that Simpson was guilty.

There is no need, then, to use a locution such as 'deciding to believe.' Consequently, an appeal to doxastic decisions amounts to doing epistemology by terminological fiat: to bring into apparent existence, through a mere terminological maneuver, a phenomenon that does not actually have a basis in reality.

This objection is based on confusion. A moment's reflection should suffice to see that (2), understood literally, cannot be true. Whether Simpson is guilty depends on whether or not he killed is former wife, not on what anybody decides. So Bugliosi cannot strictly speaking decide Simpson's guilt. What he can decide is to *believe* that Simpson is guilty. Thus (2) is elliptical: an economical way of expressing (1), which is philosophically the more accurate way of expressing the point in question. It is true that, for the sake of rhetorical elegance and convenience, ordinarily we don't talk about decisions to believe. The reason for that is not, however, that we never make doxastic decisions, but rather that it is not normally necessary to state the obvious: that when after considering our evidence we decide that such-and-such is the case, we do not take ourselves to have made

it the case that such-and-such by deciding it to be the case. Rather, we merely take ourselves to have decided to believe that such-and-such is the case.

Let us now reconsider the claim on which the previous objection was based, the claim that we cannot become introspectively aware of a mental event that could be called 'deciding to believe.' Suppose, at some point prior to the Simpson case, Bugliosi reflected on the meaning of propositions (1) and (2), and decided that when we say something like (2), what we really mean is something like (1). Suppose further that, after the DNA evidence was presented, Bugliosi reviewed its impact in light of all the other evidence and concluded that this evidence tipped the scale and conclusively established Simpson's guilt. This act, it seems to me, could be condensed in a short moment with accompanying introspective awareness of it. And since, given our first assumption, Bugliosi takes the act of deciding that Simpson is guilty to be the act of deciding to believe that Simpson is guilty, I see no reason to deny that, at that moment, Bugliosi could have become introspectively aware of the act of having decided to believe that Simpson is guilty.

A third objection goes as follows. In the case of an alleged doxastic decision, there is nothing left to do. Once the 'decision' has been made, that is, once one has concluded that one's evidence for p is better than one's evidence against p, one believes that p. This is different in the case of practical decisions: once one has decided to do a certain thing, there is something that is left to do: the thing one has decided to do. But this is essential to decision-making. If the decision to  $\varphi$  coincides with  $\varphi$ -ing, then one has not decided to  $\varphi$  but simply  $\varphi$ -s. Consequently, it is misleading to speak of deciding to believe something. What appears to be a decision to believe something is merely the formation of a belief.

This objection exploits the fact that, when we make doxastic decisions, we usually automatically and instantaneously adopt the attitude we take to be supported by our evidence. However, this fact is not a good reason to conclude that making a doxastic decision, and forming the doxastic attitude one has decided to take, are not distinct mental events. That bringing an episode of epistemic deliberation to a conclusion, and forming the belief one takes to be supported by one's evidence, are indeed two distinct mental events can be seen as follows. Let's suppose you decide your evidence supports believing that p. If so, you have formed a belief: the belief that your evidence supports believing that p. Of course, that belief has a different content from the belief that p, the belief that you take to be supported by your evidence. Let us call the former the verdict belief, and the latter the object belief. According to the objection, to form the verdict belief is to form

the object belief. But that cannot be true, for there is a different content in each case. Furthermore, note that making a doxastic decision need not result in the formation of an object belief. Suppose you decide that your evidence supports suspension of judgment with regard to p. If so, you have formed a verdict belief about what your evidence supports, but you have not formed an object belief at all. So in that case, the doxastic decision obviously is not to be identified with the formation of an object belief. But if the doxastic decision is an independent mental event in its own right in that case, then surely it is also an independent mental event in its own right when an object belief is formed.

Moreover, recall the example of the racist I mentioned above. Usually, making a doxastic decision naturally and instantaneously results in adopting the attitude one takes to be supported by one's evidence. But this need not be the case, for as the example of the racist illustrates, there may be emotional resistance to forming the doxastic attitude one has selected as appropriate. Thus the absence of a temporal interval between the decision and its execution is only an accidental feature of doxastic decision making. It would be a mistake, therefore, to assert that in doxastic decision-making, there is nothing left to do. Logically, there always is something left to do: to adopt the attitude that one takes to be supported by one's evidence. Of course, in practice adopting that attitude usually requires no additional effort. But it is nevertheless a mental act that is distinct from the act of deciding what one's evidence supports.

This concludes the review of objections to my claim that there are such things as making and executing doxastic decisions. Next, I will examine the extent to which we have voluntary control over our doxastic attitudes for each of the three concepts of voluntary control I have distinguished.

# 4. Are Doxastic Attitudes under Hypothetical Voluntary Control?

If we can execute a decision to take a doxastic attitude toward p that is different from the one we have actually taken, then we have hypothetical voluntary control over our doxastic attitude toward p. But can we execute decisions to take alternative doxastic attitudes? The answer to this question is: of course we can, at least under ordinary circumstances. What indeed should stand in the way of executing a decision to take an alternative doxastic attitude? Suppose you believe that p, but you weigh your evidence and decide your evidence contradicts p. What should then prevent you from disbelieving that p? And if you weigh your evidence and

decide your evidence neither supports nor contradicts p, what should then prevent you from suspending judgment about p? I do not mean to suggest that the execution of one's doxastic decisions never runs into problems. For example, a woman suffering from anorexia might compulsively continue to believe she is overweight even though she came to the conclusion that this belief is not supported by any evidence. Normally, however, neither compulsions nor other abnormal conditions stand in the way of executing our doxastic decisions. So under ordinary circumstances, we execute our doxastic decisions spontaneously and automatically. This is different in the area of action. The execution of our practical decisions can meet with resistance in a way the execution of our doxastic decisions does not. If I decide to steal your wallet, you will have a pretty good idea how to prevent me from executing that decision. But if I told you of my decision to believe that there is an alien from outer space among us, I don't think you would have any idea how to prevent me from executing that decision short of killing me, or perhaps merely knocking me out. 16 Doxastic attitudes, then, are under hypothetical voluntary control to an even greater extent than our actions. For the purpose of illustration, let us consider three examples.

My first example is that of an unjustified belief. Suppose Charles believes the check he just wrote won't bounce. His belief is rooted in wishful thinking and unsupported by his evidence. Does Charles have hypothetical voluntary control over this belief? Our answer to this question depends on our assessment of what will happen if he reflects on the epistemic credentials of his belief. Suppose he asks himself whether he has any good reason to believe the check won't bounce. He concludes that he has no idea how much money there is in his checking account, that there might very well not be enough in it to cover the check, and he notices that he fell victim to wishful thinking. So he tells himself: I'd better suspend judgment as to the question of whether this check will bounce. Is Charles going to encounter an execution problem? We could introduce in Plantinga-style a brainlesion, a mad scientist, or an alien from Alpha Centauri with a ray gun, causing Charles to be stuck with his belief in spite of his determination to get rid of it. But neglecting such bizarre circumstances, there simply is no question as to what will happen after Charles has decided to suspend judgment: he will suspend judgment. And this is precisely what will happen almost all of the time when, having iden-

<sup>&</sup>lt;sup>16</sup> You could argue with me. But in that case you don't interfere with the execution of my decision, but are trying to make me revise my decision.

tified an unjustified belief, we decide to get rid of it: we will get rid of it, that is, execute our doxastic decision. I take it then that unjustified beliefs are normally under hypothetical voluntary control: normally if we decide to drop them we can.

Second, let's consider a situation in which we have conflicting evidence. Suppose I am on a jury, having to decide whether the defendant is guilty. Until the final arguments are made I suspend judgment, but when it comes to reaching a verdict I ask myself whether the evidence leaves any reasonable doubt as to the defendant's guilt. Suppose I decide that it does not; that is, I decide to believe the defendant is guilty. What should stand in the way executing this decision? Or suppose I decide there is reasonable doubt, and thus to refrain from believing the defendant is guilty. What should prevent me from executing that decision? Again there are conceivable bizarre circumstances under which execution will run into a problem. But normally jury members do not encounter such problems. It therefore seems to me that we normally enjoy hypothetical voluntary control over those doxastic attitudes we decide to take when we consider conflicting bodies of evidence.

The example of conflicting evidence is likely to invite the following objection. As soon as I have determined that the evidence leaves no reasonable doubt as to the defendant's guilt, I have no choice. At that point I am stuck with the belief that the defendant is guilty, and that belief is beyond my voluntary control. This objection certainly needs to be considered, and further below I will come back to it.<sup>17</sup> However in the present context it misses its target. What I am claiming is that I have hypothetical voluntary control over whether to believe the defendant guilty or not. And I do have such control if whatever attitude toward the proposition "The defendant is guilty" I decide to take I can take — a condition that is ordinarily met. Of course there is a further problem here. It is hard to see how, after I decide that the evidence proves the defendant guilty, I could then decide to take an attitude other than believing what I think the evidence proves. That is why it is plausible to suppose that, after I have made my decision about the defendant's guilt, I'm stuck with my doxastic attitude. But if this is taken to be a reason for denying voluntary control, we are concerned no longer with the hypothetical but rather the categorical sense of that concept.

Third, let's consider what makes up a huge bulk of a person's overall doxastic system: doxastic attitudes justified by perceptual experiences. When I teach, I see

<sup>&</sup>lt;sup>17</sup> See my discussion of Alston in section 7.

and thus believe that there are other people in the class room. And I do not see, and thus disbelieve, that there is an elephant in the room. Do I have hypothetical voluntary control over these doxastic attitudes? I clearly do. What should prevent me from executing the decision to disbelieve that there are other people in the room, or from executing the decision to believe that there is an elephant present, should I decide, after an examination of my evidence, to take these alternative attitudes?

Again brain lesions and the like can serve as conceivable obstacles. But I do not suffer from any brain lesion, nor have I ever encountered a mad scientist pointing a ray gun at me. The fact is that I don't recall having ever experienced anything remotely like the circumstances under which I would be prevented from changing my doxastic attitudes should I decide to change them. Thus, if while teaching a class I decided, for whatever reasons, that the students I seem to be perceiving are just clever imitations, say robots or holograms, and thus decided to disbelieve that there are other people in the class room, nothing would prevent me from executing that decision. And if while teaching a class I had a clear and vivid elephant-like experience and because of that experience decided to believe an elephant was in the class room, nothing would stand in the way of executing that decision. Thus, when I teach my classes, I enjoy hypothetical voluntary control over the doxastic attitudes I take toward the propositions that there are other people in the room, and that there is an elephant present.

Some such story can be told for most of our doxastic attitudes. Probably sometimes a doxastic attitude remains in place even after the subject has decided in favor of another one. But there is little reason to suppose that such situations occur frequently. And so I conclude that we have hypothetical voluntary control over most, if not almost all, of our doxastic attitudes.

# 5. Are Doxastic Attitudes Under Deep Voluntary Control?

If we replace hypothetical with deep voluntary control, the extent of voluntary control over our doxastic attitudes shrinks, but not significantly. Deep voluntary control is the combination of hypothetical voluntary control and a further, internal requirement the satisfaction of which ensures that one's will is not subjected to phenomena such as brainwashing, drug addiction, paranoia, hypnosis, and the like. Beliefs that have their origin in such phenomena are not under deep voluntary control. However, most likely they would not be under hypothetical control either.

Moreover, our beliefs are not ordinarily caused by such phenomena. The move from hypothetical to deep voluntary control should not, therefore, make a significant difference with regard to the extent of doxastic voluntariness. And in fact, if we look at particular versions of refined compatibilism, it turns out that the constraints by which they discriminate between free and unfree decision-making are not the kind of constraints that pose a problem for exerting voluntary control over doxastic attitudes. For an illustration of this point, let us briefly consider the accounts proposed by Harry Frankfurt and Richard Double.

Frankfurt construes free will in analogy to the traditional, compatibilist construal of freedom of action as (roughly) the freedom to do what one wants to do. He writes: "Analogously, then, the statement that a person enjoys freedom of the will means (also roughly) that he is free to want what he wants to want. More precisely, it means that he is free to will what he wants to will, or to have the will he wants." And for a person to have the will he wants to have is for that person to enjoy conformity between first-order and second-order desires. For example, an unwilling drug addict is a person whose will is to take the drug he is addicted to, without, however, wanting to have that will. Unlike a willing drug addict, such a person does not secure "conformity of his will to his second-order volitions" and thus, in wanting to take the drug, does not exercise free will. 19

Can Frankfurt's account of free will be applied to doxastic attitudes? Perhaps. Just as we can distinguish between a willing and an unwilling drug addict, we might distinguish between, for example, a willing and an unwilling theist. The willing theist believes in God, *wishes* to believe in God, and this wish conforms to the second order desire that he have that wish.<sup>20</sup> The unwilling theist also believes in God and wishes to believe in God. Unlike the willing theist, however, he has the second order desire that he not wish to believe in God. The willing theist, then, exercises free will in believing that God exists, whereas the unwilling theist does not, or so the proposal would go. I am not confident that this approach gives us a plausible account of deep voluntary control over doxastic attitudes. For it seems to me that what is both necessary and sufficient for such control is the

<sup>&</sup>lt;sup>18</sup> Frankfurt (1986), 75.

<sup>19</sup> Ibid.

<sup>&</sup>lt;sup>20</sup> By saying that he wishes to believe in God, I do not mean to suggest that his belief is the result of wishful thinking. Rather, let us suppose his belief is justified. The point I am making is that he would like to be the sort of person who believes in God, and in this sense wishes that he believe in God. And perhaps it is part of being a devout theist that one wishes this belief to be as firm as possible.

ability to assess one's evidence without being irrationally constrained by paranoia, neurosis, hypnosis, manipulation, and the like. And if one's ability to deliberate about one's evidence is thus unconstrained, then, I suspect, one enjoys deep voluntary control even if one does not enjoy the kind of conformity Frankfurt demands. However, in the present context, my concern is not whether a Frankfurt-type account of deep voluntary control over our doxastic attitudes is plausible, but whether it would significantly constrain its extent. And the answer to this question is that it would not. The example of the unwilling theist illustrates that occasionally we can suffer from a lack of conformity between our first-order and second-order desires with regard to our doxastic attitudes. But it would be preposterous to suggest that, in relation to our doxastic attitudes, most of the time our first-order desires do not conform to our second-order desires. It is hard to see, then, that modeling deep voluntary control after Frankfurt's account of free will is going to supply us with a reason to endorse the thesis of doxastic involuntarism.

Of course, Frankfurt's account of free will does not exhaust the range of options compatibilists can exercise. There are two standard objections to it.<sup>21</sup> The first problem with Frankfurt's account of free will is that, with regard to a person's second-order desires, we must ask how they came about. If they came about in a way that is inconsistent with the subject's autonomy, that is, through paranoia, neurosis, hypnosis or other means of external manipulation, then it is hard to see why conformity of lower level desires to them should ensure freedom of will. Second, according to some philosophers, free will requires rationality. But for rationality, conformity of first-order to second-order desires is not enough. For an illustration of both of these problems, consider a member of Hamas who volunteers for a suicide bombing. Even on the supposition that his first-order desires conform to his second-order desires, it is doubtful that his decision was a manifestation of free will, for his decision is, as we may suppose, the result of both manipulation and manifest irrationality.

Richard Double has proposed an account of free will that is designed to avoid these problems.<sup>22</sup> It involves the following individually necessary and jointly sufficient conditions: (i) Self-knowledge: the subject must know what her mental states are. (ii) Reasonability: the subject must be motivated to engage in critical evaluation of her beliefs, desires, and choices. (iii) Intelligence: the subject must

<sup>&</sup>lt;sup>21</sup> See Double (1991), 31ff, Slote (1985), and Van Inwage (1983), 109.

<sup>&</sup>lt;sup>22</sup> Double (1991), chapter 2.

be able to "learn facts about the world, retain that information, and draw inferences from it," and she must have the "ability (rather than the desire) to reason well about ourselves and our mental states." [iv) Efficacy: the subject must have the power to makes choices that are indicative of (i)-(iii), that is, choices that express self-knowledge, and are reasonable as well as intelligent. (v) Unity: The subject must be a single agent, "as opposed to the possibilities that there are no agents, or multiple agents, or multiple recalcitrant cognitive subsystems." The demands for reasonability, intelligence and efficacy ensure that a Hamas suicide bomber will not count as a reasonable subject. He might be willing to examine his beliefs, desires, and choices, but this examination would not be critical, intelligent, and efficacious in the required sense. Is Double's account of free will plausible? This question I leave open, for my present concern lies not with its plausibility, but rather with the concept of voluntary control we can derive from it. [25]

Let us, then, think of deep voluntary control as hypothetical control plus freedom of the will as conceived of by Double: we have voluntary control over our doxastic attitudes if, and only if, we can execute alternative doxastic decisions, and second, our doxastic decisions meet Double's five conditions. If that is the way we think about voluntary control in the doxastic domain, we must conclude that we enjoy a good deal of it. For unless excessively strengthened, none of Double's five conditions raises much of a problem for ordinary agents in ordinary situations. Certainly self-knowledge and unity pose no serious problem for ordinary agents. The conditions of reasonability and intelligence might be harder to meet, but even here there is no reason to suppose that most people are unable to assess their beliefs critically or do not, as Double puts it, "desire to perform a critical and nondogmatic evaluation [of their choices] and the mental states that bring [them] about in cases where such evaluation is appropriate." 26

Of course, there are further versions of refined compatibilism to consider. Perhaps it is possible to develop a plausible version of refined compatibilism according to which we can exercise deep voluntary control in the practical, but not in the doxastic domain. However, until such an account lies before us, construing the concept of voluntary control as does refined compatibilism leads to

<sup>&</sup>lt;sup>23</sup> Ibid, 43.

<sup>&</sup>lt;sup>24</sup> Ibid, 38.

<sup>&</sup>lt;sup>25</sup> In the end, Double himself judges the account he developed unsatisfactory. See ibid, chapter 5.

<sup>&</sup>lt;sup>26</sup> Ibid. 48

the consequence that we enjoy voluntary control over our doxastic attitudes, and do so to a large extent.

### 6. Are Doxastic Attitudes Under Categorical Voluntary Control?

According to the second premise of the argument from doxastic involuntarism, it is not possible to exert voluntary control over *any* of our doxastic attitudes. The idea is that when I see a car coming down the street, when I see rain falling, and when I see other people nearby, I am stuck with, and thus do not have voluntary control over, the beliefs that arise in me. Yet if we work within the compatibilist framework and thus construe voluntary control as either hypothetical or deep, it is a mistake to say that we lack voluntary control over beliefs such as these. We can classify these beliefs as being beyond our voluntary control only if we conceive of voluntary control in libertarian fashion and take it to be categorical.

Consider one of Alston's examples: you believe a car is coming down the street because you see one coming down the street. Although we have both hypothetical and deep voluntary control over this belief, I agree with Alston when he asserts that in such a situation we cannot drop this belief, at least not ordinarily. But why can't we? The reason why we can't is not that we are unable to execute a decision to refrain from believing this. Rather we cannot help believing a car is coming down the street because, given the excellent perceptual evidence we have when we see one coming down the street, we cannot *decide* to drop the belief that has arisen in us, that is, conclude that in light of our evidence we ought not believe that a car is coming down the street. It would be outright crazy to make such a decision, and rational people are not, at least not normally, capable of such craziness. Thus if we use the concept of voluntary control in its categorical sense, Alston turns out to be right: we do not enjoy voluntary control over whether or not to believe a car is coming down the street when we see one coming down the street.

I do not claim, however, that I could not ever decide to refrain from believing that a car is coming down the street when I take myself to perceive one coming down the street. We can imagine circumstances under which I must consider certain things that might justify me in suspending judgment on this matter. I might suddenly remember that I suffer from a brain lesion that makes me hallucinate oncoming cars, or that I am on the premises of a Hollywood studio where a special effects team projects holograms of approaching cars. I see no reason to suppose

that, under such circumstances, I could not decide to refrain from believing that there is a car coming down the street. My present concern, however, is not whether there are hypothetical circumstances in which I can decide to drop this belief, but whether, *under normal circumstances* when I take myself to be perceiving a car coming down the street, I can resist forming the belief that a car is coming down the street. Under normal circumstances, such an experience is excellent evidence — compelling, decisive evidence, we might want to say — for the belief in question. How then could I, as long as I remain within the confines of rationality, decide to drop it? The general point is that, under normal circumstances, excellent evidence is intellectually irresistible. Thus when we have such evidence, we do not normally have categorical voluntary control over the beliefs our evidence supports.<sup>27</sup>

An analogous point can be made for the domain of action.<sup>28</sup> First, consider an omission that is compellingly supported by practical reason. Ask yourself whether you can now decide to stick a knife in your hand. The phenomenology of this matter is ambiguous, and the ambiguity is precisely that between hypothetical and categorical voluntary control. It seems clear that sticking a knife in your hand is something over which you might have hypothetical voluntary control. If you were to decide to do it, say after I offer you one million dollars as an incentive, you might overcome your reluctance to injure yourself and execute that decision. It also seems clear that, under the present circumstances in which neither I nor anybody else is offering you a nice monetary incentive, you cannot decide to stick a knife in your hand, at least as long as your mind operates within the bounds of reason. It is highly unlikely, therefore, that you now have categorical voluntary control over whether to stick a knife in your hand.

Let's consider a second example. Ask yourself whether you can skip one of your classes when you have no reason at all to do so. There is little doubt that if you were to decide to cancel class because you have what you take to be a good

<sup>&</sup>lt;sup>27</sup> H. H. Price seems to acknowledge this point when he writes: "If you are in a reasonable frame of mind, you cannot help preferring the proposition which the evidence favors, much as you may wish you could." Price (1954), 16. Of course we should investigate further exactly how we are to conceive of this inability. Is a reasonable person unable to believe contrary to the evidence because her evidence causally prevents her from doing so, or is the constraint in question of a non-causal nature? This issue certainly deserves further exploration, but it goes beyond the scope of this paper.

<sup>&</sup>lt;sup>28</sup> For an excellent defense of the view that libertarianism severely constrains the extent of free action, see Van Inwagen (1995). For a critical response to VanInwagen's defense of the view in question, see Fischer and Ravizza (1995).

reason, you could execute that decision. Nobody would drag you against your will into the classroom and force you to teach. So you have hypothetical voluntary control over whether or not to teach your class. But now ask yourself whether you can decide to cancel class when you do not have any reason at all to cancel it, and, as it normally is the case, plenty of reasons in favor of teaching it. Under such circumstances to decide not to teach it would be crazy, moronic, plain dysfunctional, and I take it that normally we are not capable of such dysfunctional decision making. I don't think, therefore, that you have normally categorical voluntary control over whether or not to teach your classes.

My general point is that categorical voluntary control, if possible at all, is limited in its extent because practical decision making is constrained by practical reasons, and epistemic decision making by epistemic reasons, that is by our evidence. When our reasons are decisive or intellectually compelling, then normally they effectively constrain what we can decide to do and believe. Of course, the qualifier 'normally' is important here. I do not wish to suggest we can never decide to do what is opposed by practical reason, or to believe what is opposed by epistemic reason.

That categorical voluntary control is significantly limited by reasons is a point worth mentioning only on the assumption that determinism is false. If we assume that determinism is true, then categorical voluntary control is not merely limited but indeed nonexistent. If on the other hand we assume determinism to be false, then categorical voluntary control is, within the confines of reason, attainable. Frequently it is not clear which course of action, or which doxastic attitude, our reasons demand. And when that is the case, then, assuming determinism is false and no other effective constraints are in place, the opportunity arises to exercise categorical voluntary over what we do and believe. For example, the opportunity for categorical voluntary control might arise when there is both evidence for and against a proposition, and we attempt to sort out, in light of the total evidence available, which attitude toward that proposition would be appropriate. The opportunity for categorical voluntary control might also arise when we believe something without justification, for there might then be a time during deliberation at which it is within our power both to drop the belief and to retain it.

I do not, however, claim that we do in fact enjoy categorical voluntary control over either actions or beliefs, for whether such control is a coherent notion is far from clear. The problem is that categorical voluntary control appears to be both in need of, and inconsistent with, causal determination. If I make a decision and

that decision is not caused by anything, and thus under exactly the same circumstances I could have decided this way and that way, in which sense can the decision be said to be *my* decision, and to be a rational decision? It would appear the decision simply happened. It is hard to see, then, how I can really make a decision, and how it can be rational, unless it is causally determined. On the other hand, if it was causally determined, then my control over it was not categorical. The concept of categorical voluntary control, then, is deeply problematic.<sup>29</sup> I therefore leave it open whether we can make sense of categorical voluntary control at all. And thus the claim I have argued for in this section is merely conditional: if the concept of categorical voluntary control can be coherently worked out, and if determinism is false, then categorical voluntary control over doxastic attitudes is possible, albeit rather constrained in its extent.

## 7. Alston's Argument Against Doxastic Voluntarism

In his well-known paper "The Deontological Conception of Epistemic Justification," Alston presents the following argument. For epistemic deliberation to enable us to have

voluntary control over propositional attitudes in basically the way we do over the positions of doors and light switches, it would have to be the case that the [evaluation of the] evidence was 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.<sup>30</sup>

It is not clear exactly how the argument Alston has in mind is supposed to go. Perhaps the following is a fair construal of it.

(1) If epistemic deliberation enables us to exert voluntary control over our doxastic attitudes, then it enables us to exert voluntary control over what doxastic attitude we take up.

<sup>&</sup>lt;sup>29</sup> For a clear statement of this problem, see, for example, Double (1991), chapter 8.

<sup>&</sup>lt;sup>30</sup> Alston p. 131. The revisions in brackets were made because Alston speaks here of *the search for evidence*, rather than *the evaluation of the evidence*. This is misleading. In this passage Alston responds to Chisholm's claim that, because we can deliberate, that is, assess the epistemic credential of our beliefs, ethical predicates can be applied to both actions and beliefs. See Chisholm (1968), 224. But Chisholm's point, as well as the claim I am defending in this paper, is not that we bear responsibility for our beliefs at a time t because we can search for evidence we do not have at t, but rather because we can assess whether our beliefs are justified by the evidence we possess at t.

(2) If epistemic deliberation enables us to exert voluntary control over what doxastic attitude to take up, then we can undertake an episode of epistemic deliberation with the intention to take up a particular doxastic attitude.

(3) But we cannot undertake an episode of epistemic deliberation with the intention to take up a particular doxastic attitude.

#### Therefore:

(4) Epistemic deliberation does not enable us to exert voluntary control over what doxastic attitude to take up.

### Consequently:

(5) Epistemic deliberation does not enable us to exert voluntary control over our doxastic attitudes.

The first two premises state what Alston says in the passage I cited, and the third what is implicit in it. It is the second premise where this argument runs into a serious problem, and it is the third, which is obviously true, from which it derives its force. Deliberation is not a means of *doing* something intentionally, but rather a means of *forming an intention*. Thus we cannot deliberate about whether to  $\phi$  with the intention to  $\phi.^{31}$  If we already have an intention, there is no logical space for deliberation left. For there to be such space, it must be the case that we are not yet clear about our intentions. And this applies, of course, to both epistemic and practical deliberation.

The second premise, however, makes what is impossible for us to do — to deliberate whether or not to  $\phi$  with the intention to  $\phi$  — a necessary condition of exerting voluntary control over doxastic attitudes through deliberation. Now, if this were true for epistemic deliberation, it would have to be true for practical deliberation as well. But when we consider the second premise in the context of deliberating about what to do, it becomes clear that it cannot be true. Consider:

- (2\*) If practical deliberation enables us to exert voluntary control over whether to keep the light on or turn it off, then we can deliberate about whether to keep the light on or turn it off with the intention (say) to turn it off.
- (2\*) is obviously false. Nonetheless, it is what we get when we apply Alston's reasoning to a context of practical deliberation. Let us, then, straighten things out. Suppose I deliberate about whether to keep the light on or turn it off. Since

<sup>&</sup>lt;sup>31</sup> And this, it seems to me, is not merely a psychological impossibility, but a conceptual one.

deliberation is a means, not of carrying out an intention, but of forming one, at that point I neither intend to keep the light on, nor intend to turn it off. However, as soon as my deliberation is concluded, my decision is made and an intention formed: say the intention to turn the light off. And so I go ahead and turn off the light. And in doing so, I exert voluntary control over something by way of having deliberated about it. Quite clearly, this did not require of me anything like the sort of thing demanded by  $(2^*)$ .

Now, if  $(2^*)$  is false, then so is (2). Exerting voluntary control over my doxastic attitudes by engaging in epistemic deliberation need not meet the condition premise (2) imposes. Suppose I deliberate about whether or not to believe that p. I begin by considering my evidence with regard to p. Eventually I conclude, say, that my evidence supports p. In so doing, I formed the intention to believe that p, and subsequently believe that p. And just as deliberating about whether to  $\varphi$  with the intention to  $\varphi$  is not required in the previous case, it isn't required in this one either. So premise (2) is false. Alston has not, then, provided us with a sound argument against the thesis that epistemic deliberation affords us voluntary control over our doxastic attitudes.

It might be thought we can salvage Alston's argument by throwing out the premise about intentional undertakings and retaining only the point that, after epistemic deliberation, we merely end up, or get stuck, with the doxastic attitude we selected as the one our evidence supports. To this move, there are two replies. First, if we work within the compatibilist framework, after deliberation we are 'stuck' with our doxastic attitudes in pretty much the same way we are after deliberation stuck with our actions and omissions, which is to say: not at all, at least not normally. It is only within the libertarian approach, then, that the phenomenon of being 'stuck' with a doxastic attitude turns into a worry. However, and this is my second reply, within the libertarian approach, the phenomenon of being stuck with a doxastic attitude is compatible with its being the result of having exerted categorical voluntary control. What we need to recognize is that, even if there was a time t<sub>1</sub> at which I had categorical voluntary control over believing that p, I might not have categorical voluntary control over believing that p at a later time  $t_2$ . Suppose I form the belief that p after deliberation, and during deliberation there was a point at which I could have decided to take an alternative attitude toward p. So I believe that p because I have decided to believe that p. Suppose further that, at the end of my deliberation, the conclusion that my evidence supports p strikes me as compelling. In the situation under consideration, there was

a time  $t_1$  at which I had categorical voluntary control over whether or not to believe that p, but after I conclude my deliberation, that is, at a later time  $t_2$ , I do not have categorical voluntary control over this belief any longer. Consequently, the point the objection rests on — after deliberation we are usually stuck with the doxastic attitude we adopt as a result of our deliberation — fails to undermine the claim in question: that deliberation can afford us categorical voluntary control over our doxastic attitudes. Rather, what the objection establishes is that we cannot sustain categorical voluntary control over doxastic attitudes beyond the end of our deliberation. Once our deliberation is concluded, the doxastic attitude we have selected is locked in, and we can regain categorical voluntary control over it only if our evidential situation changes.

### 8. The Argument from Doxastic Involuntarism Revisited

It is time to return to the argument we considered at the outset: the argument against doxastic involuntarism, according to which epistemic deontology is doomed. Let us consider it again:

The Argument From Doxastic Involuntarism

- (1) If doxastic attitudes are proper objects of deontology, then it is possible to exert voluntary control over them.
- (2) It is not possible to exert voluntary control over doxastic attitudes.

#### Therefore:

(3) Doxastic attitudes are not proper objects of deontology.

If we distinguish, as I have argued we should, between hypothetical and categorical voluntary control, there is an effective rebuttal to this argument. The initial appearance of soundness it conveys is due to an equivocation which becomes apparent only if the difference between these two types of voluntary control is taken into account. The first premise asserts the needed link between epistemic deontology and voluntary control. It is most plausibly interpreted as employing the concept of voluntary control in its compatibilist sense. For consider the claim that categorical voluntary control is a necessary condition of deontological status. If this were true, deontology across the field would require an indeterministic world. Compatibilists would certainly see no reason to agree that deontology should be subjected to such a stringent requirement. Thus, if the argument is to

retain the broad appeal it is intended to have, the first premise must be read as making either hypothetical or deep voluntary control a necessary condition of deontology.

However, the second premise, which asserts doxastic involuntarism, derives its plausibility from employing the categorical sense of voluntary control. If this equivocation, and thus invalidity, is to be avoided, both premises must employ the concept of voluntary control in the same sense. Let us therefore distinguish between a compatibilist and a libertarian reading of the argument. First, let's restate the argument within the framework of compatibilism. In that case the argument's premises must be understood thus:

- (1<sub>sc</sub>) If doxastic attitudes are proper objects of deontology, then it is possible to exert hypothetical voluntary control over them.
- (2<sub>sc</sub>) It is not possible to exert hypothetical voluntary control over doxastic attitudes.

As I have argued in section 4, it's not so easy to see why  $(2_{sc})$  should be true. Here, the burden of proof lies squarely with those who wish to deny that we have hypothetical voluntary control over our doxastic attitudes.

Second, let's restate the argument's premises within the framework of refined compatibilism:

- (1<sub>rc</sub>) If doxastic attitudes are proper objects of deontology, then deep voluntary control over our doxastic attitudes is possible.
- (2<sub>rc</sub>) Deep voluntary control over our doxastic attitudes is not possible.
- $(2_{\rm rc})$  is no more plausible than  $(2_{\rm sc})$ . The point of replacing simple with refined compatibilism is to secure the relevance of free will in a way that is compatible with determinism. This is accomplished by imposing conditions that ensure the absence of things that undermine personal autonomy, phenomena such as paranoia, neurotic compulsions, hypnosis, drug addiction, and the like. But advocates of doxastic involuntarism do not suggest that beliefs are, on a broad basis, subjected to such forces, and that that is the reason why we cannot exert voluntary control over them. After all, this would be an exceedingly implausible claim. Once again, the burden of proof lies with those who oppose doxastic voluntarism. If the present version of the argument is what they have in mind, they need to make a persuasive case for the claim that deep voluntary control essentially involves some

feature that makes voluntary control over doxastic attitudes impossible. Until that case has not been made, I see no reason to accept  $(2_{rc})$ .

Third, let's consider what happens when we give the argument a libertarian reading and take voluntary control in its categorical sense.

- (1<sub>1</sub>) If doxastic attitudes are proper objects of deontology, then categorical voluntary control over our doxastic attitudes is possible.
- (2<sub>1</sub>) Categorical voluntary control over our doxastic attitudes is not possible.

As I have already mentioned, compatibilists would be adamantly opposed to  $(1_l)$ . They would see no reason to make the viability of deontology dependent on what the possibility of categorical voluntary control requires: an indeterministic world.  $(2_l)$  is problematic as well. What I will have to say about it depends on whether we take determinism to be true or false.

On the assumption that determinism is true, (2<sub>l</sub>) is true *ex hypothesi*. So philosophers who insist that voluntary control must be categorical, accept determinism, and accept the first premise, will consider this version of the argument sound. However such philosophers would be unlikely to be concerned about the viability of epistemic deontology *in particular*. If determinism is true then categorical voluntary control is impossible for doxastic attitudes and actions alike. Consequently such philosophers would have to have doubts about the viability of deontology as such. But to call deontology as such into doubt is not the original intent of the argument, and thus I suspect that critics of epistemic deontology would not want to endorse this particular version of it.

If determinism is false, then the plausibility  $(2_l)$  depends on whether there is a relevant difference between practical and epistemic deliberation, a difference that would explain why, in an indeterministic world, deliberation affords us categorical voluntary control over actions, but not over doxastic attitudes. As long as a plausible account of such a difference has not been put forth,  $(2_l)$  remains implausible.

In defense of the idea that epistemic deontology can be rejected on the ground of doxastic involuntariness, there is one further gambit available. It could be argued that, though assuming the falsity of determinism commits us to accepting the *possibility* of categorical voluntary control over our doxastic attitudes, this possibility is rarely *realized*. But this, involuntarists might argue, means that most of the time our doxastic attitudes are not suitable for deontological evaluation. This particular version of the argument from doxastic involuntarism proceeds from a revised version of the second premise:

The Argument from Doxastic Involuntarism: Final Version

- (1<sub>1\*</sub>) What can be subjected to deontology must be under categorical voluntary control.
- $(2_{1*})$  Most of our doxastic attitudes are not under categorical voluntary control.

#### Therefore:

(3<sub>1\*</sub>) Most of our doxastic attitudes cannot be subjected to deontology.

The conclusion of this argument is weaker than the conclusion we considered previously. It does not assert the impossibility of epistemic deontology altogether. It merely restricts its scope. However, in spite of this scope restriction,  $(3_{l^*})$  is a significant conclusion. If it were true it would be a mistake to conceive of epistemic justification in terms of epistemic permissibility.

Since I accept  $(2_{1*})$ , my reply to this version rests solely on rejecting the first premise: (1<sub>1\*</sub>). Once again, compatibilists would be opposed to it. However, it seems to me that libertarians should not accept it either. Within the libertarian framework, the extent of voluntary control shrinks, as I have argued, in a rather striking fashion. However, does the applicability of deontology shrink correspondingly? I do not think libertarians need to concede that it does. Consider some actions and omissions over which we lack categorical voluntary control because we cannot make the requisite decisions. You do not normally have, as I have argued above, categorical voluntary control over actions and omissions that are decisively opposed by reason, such as: sticking a knife in your hand, sticking a knife in my hand, throwing yourself in front of an oncoming truck, making your mortgage payments, and stopping at red traffic lights. But these are nevertheless things that are within the scope of deontology. You ought not now stick a knife in your own hand, and you have now an obligation to refrain from sticking a knife into my hand, you ought to make your mortgage payments, and you are obliged to stop at red traffic lights. These examples strongly suggest what at least compatibilists would see no reason to deny: that actions and omissions can enjoy deontological status even when we do not have categorical voluntary control over them.

If we accept this point for actions we have no reason not to concede it for doxastic attitudes as well. When I see other people before me, I do not have categorical voluntary control over the belief that there are other people in the room. Nor do I have such control over believing that a car is coming down the street when I see one coming down the street. However, if we accept that categorical

voluntary control is not a necessary condition for the application of deontology, we have no reason to deny that these beliefs are suitable for deontological evaluation. And thus we may after all say that, when I see other people before me, I ought to believe that there are other people before me, or that when I see a car coming down the street, it is epistemically permissible for me to believe that a car is coming down the street. Consequently, libertarians who think that deontology is applicable to actions to the extent we ordinarily think it is, have no reason to accept  $(1_{1*})$ .

This concludes my discussion of the argument from doxastic involuntarism. I have examined four different versions of it. None of them withstand critical examination. The common failure of these versions can be summed up by saying that, depending on how we think of voluntary control, voluntary control over doxastic attitudes is either surprisingly easy to have, or not required by epistemic deontology. It seems warranted to conclude, therefore, that considerations having to do with doxastic voluntarism are less of an obstacle to epistemic deontology than its opponents have claimed. Note, however, that I have endorsed neither the compatibilist nor the libertarian construal of voluntary control. Nor have I asserted that determinism is false. In the end, then, I have not argued for any claim about the extent to which we actually enjoy voluntary control over our doxastic attitudes. My chief claim has been merely conditional: if we begin with the premise that practical deliberation affords us voluntary control over our actions, then, unless a relevant difference can be established between practical and epistemic deliberation, we must conclude that epistemic deliberation affords us voluntary control over our doxastic attitudes. Which extent we ascribe to it will depend, of course, on whether we favor the compatibilist or the libertarian construal of voluntary control 32

arrived at the redaction: February 2000

Matthias Steup • Department of Philosophy • Brown Hall 123 • St. Cloud State University • 720 4th Avenue South • St. Cloud, MN 56301-4498 • USA • steup@stcloudstate.edu

<sup>&</sup>lt;sup>32</sup>A previous version of this paper was read to the Philosophy Department at Wayne State University. I am indebted to the participants in this discussion to valuable comments and objections. I also benefited from discussing the issues discussed in this paper with Robert Audi, John Bahde, Richard Feldman, Noah Lemos, and Bruce Russell.

### References

Alston, William. 1989. *Epistemic Justification*. *Essays in the Theory of* Knowledge. Ithaca: Cornell University Press.

BonJour, Laurence. 1985. *The Structure of Empirical* Knowledge. Cambridge: Harvard University Press

Baergen, Ralph. 1995. Contemporary Epistemology. Orlando, Florida: Hartcourt Brace.

Bennett, Jonathan. 1990. "Why is Belief Involuntary?" Analysis 50: 87-107.

Chisholm, Roderick. 1968. "Lewis's Ethics of Belief," in Paul Schilpp ed, *The Philosophy of C. I.* Lewis. La Salle, Illinois: Open Court, 1968.

— 1978. Theory of Knowledge, 2nd. edition. Englewood Cliffs: Prentice Hall.

Dennett, Daniel. 1984. Elbow Room. Cambridge: MIT Press.

Double, Richard. 1991. The Nonreality of Free Will. Oxford: Oxford University Press.

Feldman, Richard. 1988. "Epistemic Obligations," Philosophical Perspectives 2: 235-256.

Fischer, John Martin, ed. 1986. Moral Responsibility. Ithaca: Cornell University Press.

Fischer, John Martin, and Mark Ravizza. 1991. "When the Will is Free," in O'Connor (1995): 239-69.

Frankfurt, Harry G. 1986. "Freedom of the Will and the Concept of a Person," the *Journal of Philosophy* 68, 1971: 5-20, reprinted in Fischer (1986): 65-80.

Ginet, Carl. 1975. Knowledge, Perception, and Memory Dordrecht: D. Reidel.

Heil, John. 1993. "Doxastic Agency," Philosophical Studies 43: 355-64.

Levin, Michael. 1979. Metaphysics and the Mind-body Problem. Oxford: Oxford University Press.

Neely, Wright. 1974. "Freedom and Desire," Philosophical Review 83: 231-59.

O'Connor, Timothy, ed. 1995. Agents, Causes, & Events. Oxford: Oxford University Press.

Plantinga, Alvin. 1993. Warrant: The Current Debate. Oxford: Oxford University Press.

Price, H. H. 1954. "Belief and the Will," Proceedings of the Aristotelian Society, Supplementary Volume: 1-28.

Schoeman, Ferdinand, ed. 1987. Responsibility, Character, and the Emotions Cambridge: Cambridge University Press.

Slote, Michael. 1985. "Understanding Free Will," in Fischer (1986), 124-139.

Stocker, Michael. 1971. "Ought' and 'Can'," Australasian Journal of Philosophy 49: 303-316.

Van Inwagen, Peter. 1983. An Essay on Free Will. Oxford: Oxford University Press.

— 1995. "When Is the Will Free?" in O'Connor (1995): 219-238.

Wolf, Susan

- 1986 "Asymmetrical Freedom," in Fischer (1986): 225-240, first published in the *Journal of Philosophy* 77, 1980: 151-66.
- 1987. "Sanity and the Metaphysics of Responsibility," in Schoeman (1987): 46-62.
- 1990. Freedom Within Reason. Oxford: Oxford University Press.

Watson, Gary. 1975. "Free Agency," *Journal of Philosophy: pp. 205-20*, reprinted in Fischer (1986): 81-96.