CHAPTER 4

DEONTOLOGICAL DESIDERATA

i. Preliminaries

I now begin the detailed treatment of the items on my initial list of alleged epistemic desiderata. I will be concerned with clarification of the nature of each desideratum, how it should be construed. Where there are serious questions as to the viability of an item, those will be addressed. I discuss the deontological group first because it gives rise to crucial problems about viability, as a result of which I postponed consideration of it in Chapter 3 until and unless they can be resolved.

Here are the deontological candidates for epistemic desiderata of belief (B) that were listed in Chapter 3.

9. B is held *permissibly* (one is not subject to blame for doing so).
10. B is formed and held *responsibly*.
11. The causal ancestry of B does not contain violations of intellectual obligations.

First a word about my terminology. 'Deontology' and 'deontological' come from the Greek *deon*—'what is binding' or 'duty'. In ethics, deontology is the study of duty or obligation, and a deontological theory of ethics is one that takes duty or obligation to be the most basic ethical concept and treats it as an intrinsic ethical value of an act rather than in terms of the consequences of the act. My use is broader. I use it to range over any kind of requirement, not restricted to moral obligation, and not excluding requirements that are based on consequences of what is required. And I identify deontological considerations as having to do with the triad of statuses—*required*, *forbidden*, and *permitted*. Thus any way in which it would be epistemically desirable (desirable from the standpoint of an aim at true belief) for a belief to be required or permitted (i.e., not forbidden) would count as a deontological desideratum in my terminology.

Back to the above list, I think it will suffice to concentrate on 9 and 11. Each of these can be construed as focusing on something's being permitted, not being in violation of any intellectual requirements. Desideratum 9 is matter of the having or the acquiring of the belief being permitted. Desideratum 11 is a matter of the permissibility or lack thereof of what one did that led to the acquisition of the belief. Although 10, the formation in terms of responsibility, is familiar in the literature, I think it is ambiguous between 9 and 11 and so does not require separate treatment. The basic difference between 9 and 11 is what is said to be permitted—either the believing itself or what led up to it. Thus, to foreshadow a major point in the ensuing discussion, 9 gives rise to problems about voluntary control of belief whereas 11 does not.

I have already pointed out in Chapter 1 that it is plausible to suppose that 'justified' came into epistemology from its more unproblematic use with respect to voluntary action. I am justified in doing something, for example, appointing someone to a Teaching Assistantship on my own, provided my doing so is in accordance with the relevant rules and regulations, provided it is *permitted* by those rules and hence that I could not rightfully be *blamed* or *held to account* for it, and was acting *responsibly* in doing so.[1] The rules could be institutional, as in the above example, or legal or moral. Thus I would be morally justified in failing to make a contribution to a certain organization provided my doing so doesn't violate any moral rule. Because of this provenance it is natural to think of believing, when taken to be subject to being justified or unjustified, as subject to requirement, prohibition, and permission. We say things like "You shouldn't have supposed so readily that he would not return", "You have no right to assume that", "You shouldn't jump to conclusions", and "I ought to have trusted him more than I did". Locutions like these seem to be interchangeable with speaking of a belief as being, or not being, justified. These considerations were introduced in this book prior to the abandonment of a justification-based epistemology of belief, and in the new dispensation they have no force. Since we are thinking of 9 and 11 simply as states of affairs that are, or may be thought to be, important goals of cognition, the fact that they have often been thought to constitute a belief's being justified, with all the associations that brings from talk of the justification of actions, has lost whatever meta-epistemological significance it had under the old dispensation. The idea of a belief's being required, permitted, or forbidden will have to swim or sink on its own, without support from the etymology of 'justified'. I will now enter onto the elucidation of 9 and a critical discussion of its credentials as an epistemic desideratum. The criticism will mostly hinge on whether we have effective voluntary control of believings. I will argue that we do not.

---

[1] I don't suggest that doing what is not permitted by the rules is coextensive with being subject to blame for doing it. One might have a valid excuse for doing it despite the rules. When I speak of violating a rule as being blameworthy, it is presupposed that there is no such excuse.

It seems clear that the terms of the deontological triad, *permitted*, *required*, and *forbidden*, apply to something only if it is under effective voluntary control. By the time-honored principle "Ought implies can", one can be obliged to do A only if one has an effective choice as to whether to do A. It is equally obvious that it makes no sense to speak of S's being permitted or forbidden to do A if S lacks an effective choice as to whether to do so. Therefore, the most fundamental issue raised by the claim of 9 to be an epistemic desideratum is whether believings are under effective voluntary control. If they are not and hence if deontological terms do not apply to them, alleged epistemic desiderata like 9 do not get so far as to be a candidate for an epistemic desideratum. It suffers shipwreck before leaving port. I will argue that believings are not subject to voluntary control. But before that, there are some preliminary points to be made.

First, if I considered the possibility of deontological ED for beliefs to be a live one, I would need to consider a belief's enjoying the stronger deontological status of being a case of complying with an epistemic obligation, doing what is required, as well as the weaker status of merely being something that is epistemically permitted. But since I hold that no deontological status is possible for beliefs, I will not need to go into the different statuses separately. And since justificationists of a deontological bent have concentrated on a belief's being epistemically permitted, I will go along with that focus.

Second, although the discussion in the book thus far has been solely in terms of belief, we need to include consideration of other propositional attitudes that are contrary to belief. Chisholm (1977, chap. 1) speaks in terms of a trichotomy of 'believe' (or 'accept'), 'reject', and 'withhold' that p. Since rejecting p is identified with believing some contrary of p, at least not-p, it brings in no new kind of propositional attitude, but withholding p, believing neither it nor any contrary, does. The basic point here is that one has control over a given type of propositional attitude only if one also has control over some field of incompatible alternatives. To have effective control over believing that p is to have control over whether one believes that p or takes on some alternative thereto. Therefore, to be strictly accurate we should say that our problem about 9 concerns voluntary control over intellectual propositional attitudes generally. Though my formulations will mostly be in terms of belief, they should be understood as having this more general bearing.

Third, something must be said about the relation between the voluntary control of actions and of states of affairs. Thus far I have been oscillating between the two. A belief is a more or less long-lived state of the psyche that can influence actions and reactions of the subject so long as it persists. And the same holds for other propositional attitudes. Thus, in speaking of voluntary control of beliefs, we have been speaking of the control of states. But couldn't we just as well speak of the voluntary control of the action of bringing about such states: accepting, rejecting, or withholding a proposition? If the two are strictly correlative, we could equally well conduct the discussion in terms of either. Whenever we are responsible for a state of affairs by virtue of having brought it about, we may just as well speak of being responsible for the action of bringing it about. There are reasons, however, for proceeding in terms of states.

The main reason is this. If we hold that beliefs are subject to deontological evaluation because they are under voluntary control, we need not restrict ourselves to beliefs that are formed intentionally by a voluntary act. I could be blamed for believing that p in the absence of adequate evidence, even if the belief was formed automatically, not by voluntarily carrying out an intention to do so. Provided believing in general is under voluntary control, it is enough that I could have rejected or withheld the proposition by a voluntary act had I chosen to do so.

The final preliminary note is this. Our issue does not concern free will or freedom of action, at least in any sense in which that goes beyond one's action being under the control of the will. On a "libertarian" conception of free will this is not sufficient; it is required also that both A and non-A be causally possible, given all the causal influences on the agent. A libertarian will, no doubt, maintain that if deontological concepts are to apply to believings in the same sense in which they apply to overt actions, then all the libertarian conditions will have to apply to believings. Here, however, I am concerned only with whether believings are under voluntary control.

ii. Basic Voluntary Control of Believing

Locutions like the ones cited earlier as encouraging the application of deontological terms to believing—"You shouldn't jump to conclusions", "I had to accept his testimony; I had no choice"—also strongly suggest that belief is under voluntary control. Else why could we speak of what beliefs one should or shouldn't form, or that one did or did not have a choice as to whether one forms a certain belief? Though this view is distinctly out of favor today, it still has its defenders.[2] Such locutions also naturally suggest not only that believing is under voluntary control but that this control is of the maximally direct sort that we have over the motions of our limbs, the voluntary movements of which constitute *basic* actions. A basic action is one that we perform "at will", just by an intention, volition, choice, or decision to do so. It is something we "just do", not by doing something else. Let's call the kind of control we have over states of affairs we can bring about by basic actions *basic voluntary control.* If we do have voluntary control of beliefs, we have the same reason for supposing it to be basic control that we have for supposing ourselves to have basic control over movements of our limbs, namely, that we are hard pressed to specify any action by doing which we get the limbs moved or the beliefs acquired. Hence it is not surprising that the basic voluntary control thesis has had distinguished proponents throughout the history of philosophy. Augustine, Aquinas, Descartes, Kierkegaard, and many others have usually been read this way.[3] And discussions pro and con of the voluntary control of beliefs have mostly focused on the basic control version. Nevertheless, as the subsequent discussion will show, there are other forms of voluntarism about belief that need to be taken into account in a complete treatment.

But for now I am concerned to give a critical examination of the basic voluntary control thesis. Those who have attacked it are divided between those who hold that believing at will is logically impossible and those who hold that it is only psychologically impossible, a capacity that we in fact lack though one we conceivably could have had.[4] I cannot see any sufficient reason for the stronger claim, and I shall merely contend that we are not so constituted as to be able to take propositional attitudes at will. My argument for this, if it can be called that, simply consists in asking you to consider whether you have any such power. Can you, at this moment, start to believe that the Roman Empire is still in control of western Europe, just by deciding to do so? If you find it incredible that you should be sufficiently motivated to even try to believe this, suppose that someone offers you $500 million to believe it, and that you are much more interested in the money than in believing the truth. Could you do what it takes to get that reward? Remember that we are speaking of believing *at will*. No doubt, there are things you could do that would increase the probability of your believing this, but we will get to that later. Can you switch propositional attitudes toward that proposition just by deciding to do so? It seems clear to me that I have no such power. Volitions, decisions, or choosings don't hook up with propositional attitude inaugurations, just as they don't hook up with the secretion of gastric juices or with metabolism. There could conceivably be individual differences in this regard. Some people can wiggle their ears at will, but most of us cannot. However, I very much doubt that any of us are endowed with the power of believing that p, for any given p, at will. The temptation to suppose otherwise may stem from conflating that power with others that are clearly distinct. If I were to set out to bring myself into a state of belief that p, just by an act of will, I might assert that p with what sounds like conviction, or dwell favorably on the idea that p, or imagine a sentence expressing p emblazoned in the heavens with an angelic chorus in the background intoning the Gloria of Bach's Mass in B Minor. All this I can do at will, but none of it amounts to forming a belief that p. It is all show, an elaborate pretense of believing. Having gone through all this, my propositional attitudes will remain just as they were before; or if there is any change, it will be as a *result* of these gyrations.[5]

---

[2] See, e.g., Ginet 1985 and Meiland 1980.

[3] On the basis of a distinction between believing that p and "accepting" that p, according to which the latter but not the former is a voluntary action, I have argued that these philosophers and others are best construed as ascribing voluntary control to accepting, not believing. See Alston 1996b.

[4] The best-known defense of the logical impossibility is Bernard Williams's "Deciding to Believe", in Williams 1972. It has been criticized in, inter alia, Govier 1976 and Winters 1979.

Don't suppose that our inability to believe at will is restricted to what is obviously false. It also extends to beliefs that are obviously true. I have already made the point that voluntary control attaches to sets of contraries. To take the simplest case, if the sphere of my voluntary control does not extend both to A and to not-A, then it attaches to neither. If I don't have the power to choose between A and not-A, then we are in no position to say that I did A at will, rather than just did it, accompanied perhaps by a volition. Thus, even if I willingly, or not unwillingly, form perceptual beliefs in the way I do, it by no means follows that I form those beliefs *at will*, or that I have voluntary control over such belief formation. It would have to be true that I have voluntary control over whether I *do or do not* believe that the tree has leaves on it when I see a tree with leaves on it just before me in broad daylight with my eyesight working normally. And it is perfectly clear that in this situation I have no power at all to refrain from that belief. So it is with everything that seems obvious to us. We have just as little voluntary control over ordinary beliefs formed by introspection, memory, and simple uncontroversial inferences from uncontroversial premises.

---

[5] A bit later in the discussion I will present other tempting conflations.

The above discussion may suggest to the voluntarist that he can still make a stand on propositions that do not seem clearly true or false and hold that there one often has the capacity to adopt whatever propositional attitude one chooses. In religion, philosophy, history, and high-level scientific inquiry it is often the case that, so far as one can see, the relevant arguments do not definitively settle the matter one way or the other. I engage in a prolonged study of free will or causality. I carefully consider arguments for and against various positions. It seems to me that none of the positions have been decisively established, though there are weighty considerations that can be urged in support of each. There are serious difficulties with all the positions, though, so far as I can see, more than one contender is left in the field. So what am I to do? I could just abandon the question. But, alternatively, I could, so it seems, simply *decide* to adopt one of the positions. Is that not what I must do if I am to make any judgment on the matter?

There are also practical situations in which we are confronted with incompatible answers to a certain question, none of which we see to be clearly true or false. Here we often do not have the luxury of leaving the field; since we must act in one way rather than another, we are forced to form and act on some belief about the matter. It would be a good idea for me to plant these flowers today *iff* it will rain tomorrow. But it is not at all clear to me whether tomorrow will be rainy. I must either plant the flowers today or not, and if I just ignore the issue, that will be equivalent to assuming that it will not rain tomorrow. Hence the better part of wisdom would be to make a choice between the alternative predictions. On a larger scale, a field commander in wartime is often faced with questions about the current disposition of enemy forces. But often such information as he has does not tell him just what that disposition is. In disposing his own forces he must act on some assumption about the enemy's forces. Hence he is forced to decide on a hypothesis as to that disposition and act on that basis. What else can he do?[6]

---

[6] Even if beliefs can be formed at will in these kinds of cases, there still remain vast stretches of our belief, including all the cases discussed above, where it seems obvious what is the case, where we have already seen believing at will not to be a possible move. And so it would still be true that believing permissibly would not be generally viable as an epistemic desideratum.

Despite the intuitive appeal of the idea that beliefs are formed at will in these cases, there are several alternative construals, one or another of which is a better reading of each. Begin with the philosopher who really does come to believe the libertarian account of free will or the epiphenomenalist position on the mind-body question. Where that happens it is presumably because at least for the moment the considerations in favor of the position seem to be conclusive, even though previously they did not. And at that time the belief follows automatically from that momentary seeming of conclusiveness, just as it does in cases where it always seems obvious what the truth of the matter is whenever one turns one's attention to it. At that moment, S is no more able to accept a compatibilist account of free will or a hard-nosed materialism on the mind-body problem than he would be if the positions he comes to believe had seemed obviously true from his first consideration of the problem. If, at a given time, it still seemed to the philosopher that libertarianism and compatibilism were approximately equally well supported, how could she simply decide to believe one rather than another? How could we do that any more than, lacking any reasons at all for one alternative rather than the other, we decide to believe that the number of ultimate particles in the universe is even rather than odd?

The above account in terms of a momentary sense of conclusive support for one alternative could also apply to our practical cases. It could be that the military commander, at a certain point in his deliberations, comes to think the reasons for a particular hypothesis concerning the disposition of enemy forces are conclusive. But I believe that there are other construals for both the theoretical and practical cases. For one thing, the subject may be *resolving to act as though it is true that p*, adopting it as a basis for action without actually believing it. This could well be a correct description of the military commander. He may have said to himself: "I don't know what the disposition of enemy forces is. I don't even have enough evidence to consider one hypothesis much more likely than any other. But I have to proceed on some basis or other, so I'll just assume that it is H and make my plans accordingly". If that's the way the land lies, it would be incorrect to describe the commander as believing that the disposition of enemy forces is H or having any other belief about the matter. He is, self-consciously, proceeding on an assumption concerning the truth of which he has no belief at all. One may also make an assumption for theoretical purposes, in order to see how it "pans out" in the hope that one will thereby obtain some additional reasons for believing it to be true or false. A scientist can adopt "as a working hypothesis" the proposition that the atomic nucleus is positively charged, draw various consequences from it, and proceed to test those consequences. He need not believe that the atomic nucleus is positively charged in order to carry out this operation. Indeed, he would be doing this because he does not yet know what to believe about the matter. Likewise a philosopher might take materialism as a working hypothesis to see how it works out in application to various problems.

Working hypotheses may also be involved in activities that are a blend of the theoretical and the practical. One may accept the existence of God, or some more robust set of religious doctrines, as a guide to life, trying to live in accordance with them, seeking to act and feel one's way into a religious community, in order to determine how the doctrines work out in the living of them, both in terms of how satisfactory and fulfilling a life they enable one to live and in terms of what evidence for or against them one acquires. Again, at least in early stages of this process, one does not yet believe the doctrines in question.

There are other possibilities as well. S may be seeking, for whatever reason, to bring herself into a position of believing p, and she, or others, may confuse this activity, which can be undertaken voluntarily, with believing the proposition to be true. Or S may align herself with some group—a church, a political party, a group of thinkers—that is committed to certain doctrines, and this, which can be done voluntarily, may be confused with coming to believe those doctrines. Finally, there is the distinction between *acceptance* and *belief* that was briefly mentioned earlier. The basic distinction is that belief is something that one *finds* oneself with, something that springs into consciousness spontaneously when the question is raised. Whereas acceptance of a proposition is, at least in the first instance, a deliberate voluntary act of accepting a proposition as true. It differs from the "working hypothesis" or "assuming that p as a basis for action" in that, unlike these cases, S does commit himself to p's being true. He "takes it on board" as one of the things he acts on and draws consequences from. It is, we might say, just like belief except that the commitment to p's being true doesn't arise spontaneously but, at least at the outset, has to be kept in activation by a deliberate voluntary act. Thus the philosopher and the religious seeker might accept, in this sense, a position on the free-will issue or the mind-body problem or various religious doctrines. The philosopher, even though libertarianism does not seem to him to be conclusively established, might accept it—take it as his position on the issue, defend it, draw various consequences for it, while seeking for conclusive evidence pro or con, and not yet finding himself believing it. And there is an analogous possibility for religious doctrines.[7]

---

[7] The latter application is explored in detail in Alston 1996b.

Thus I take it that the analysis of a wide variety of supposed cases of believing at will reveals that in each case coming to believe that p may well have been confused with something else. Hence I think that there is a considerable case for the position that no one ever acquires a belief at will.

iii. Other Modes of Voluntary Control of Believing

The demise of basic control of belief is by no means the end of voluntarism about belief. Many deontologists, after avoiding any commitment to what they call "direct voluntary control of belief" (what I have called "basic voluntary control"), insist that beliefs are subject to what they term "indirect voluntary control".[8] They generally use this term in an undiscriminating fashion to cover any sort of voluntary control that is not basic. Hence they fail to distinguish the three kinds of nonbasic control I will proceed to enumerate.[9] Some of their examples fit one of my three types and some another.

First, note that we take many nonbasic overt actions and their upshots to be under voluntary control in a way that is sufficient for their being required, permitted, or prohibited. Consider opening a door, turning on a light, and informing someone that p. Succeeding in any of these requires more than a volition; in each case I must make one or more bodily movements, and these movements must have certain consequences. In order for me to open a door, I must pull it, push it, kick it, or put some other part of my body into suitable contact with it (assuming that I lack telekinetic powers), and this must result in the door's coming to be open. In order to inform H that p, I must produce various sounds, marks, or other perceivable products, and the product in question must fall under linguistic rules in such a way as to constitute a vehicle for asserting that p. Thus actions like these are not immediately consequent on a volition and are not strictly done "at will". Nevertheless, I might be blamed for my failure to open the door when it was my obligation to do so and I was not prevented from performing basic bodily movements sufficient to bring it about that the door was open. In typical cases we take the extra conditions for success for granted. We suppose that if the agent will just voluntarily exert herself in a way that is open to her, the act will be done. Here we can say that the action and its upshot are subject to the *immediate voluntary control* of the agent

---

[8] See, e.g., Alvin Goldman 1980, Plantinga 1983, Wolterstorff 1983, Moser 1985, Steup 1988.

[9] Even the extended treatment in Pojman 1986 fails to make any distinctions within "indirect control".

(more strictly, *nonbasic* immediate voluntary control), even though more than an act of will is required. I call this control "immediate" since the agent is able to carry out the intention right away, in one uninterrupted intentional act, without having to return to the attempt a number of times.[10] I will use the term 'direct control' for both basic and nonbasic immediate control. If beliefs were subject to one's direct control in either way, that would suffice to render them susceptible to deontological evaluation.

But are beliefs always, or ever, within our immediate nonbasic voluntary control? As in the discussion of basic control we can first exempt most beliefs from consideration. Where it is perfectly clear that a certain proposition is true or false, as with typical perceptual, introspective, memory, and simple inferential beliefs, it is absurd to think that one has any such control over whether one accepts, rejects, or withholds the proposition. When I look out my window and see rain falling, water dripping off the trees, and cars passing by, I no more have immediate nonbasic control over whether I accept those propositions than I have basic control. I form the belief that rain is falling willy-nilly. There is no way I can inhibit this belief or acquire a contrary belief. At least there is no way I can do so on the spot, in carrying out an uninterrupted intention to do so. What button would I push? I could try asserting the contrary in a confident tone of voice. I could rehearse some skeptical arguments. I could invoke the Vedantic doctrine of *Maya*. I could grit my teeth and command myself to withhold the proposition. But none of these will have the least effect on my doxastic condition. Since cases in which it seems obvious to the subject what is the case constitute an enormously large proportion of propositional attitudes, the above considerations show that immediate nonbasic voluntary control cannot be the basis for the application of deontological concepts to most of our beliefs and withholdings.

But what about situations in which it is not clear whether a proposition is true or false? Here I can simply refer the reader back to the last section, in which I argued with respect to basic control that the cases in which it may look as if one comes to believe a proposition at will are best construed in other ways. In those cases involving the philosopher, the general, and the gardener, it is, I claim, implausible to suppose that the subject acquired a belief voluntarily, whether by a mere act of will or by a series of basic or more nearly basic actions that led right away to the intended result. Here, as with the obviously true or false cases, we are at a loss to think what button to push, what bodily movements to make so as to bring about the formation of an intended belief. Until some plausible story can be told as to what one can do voluntarily to result in a belief's being formed immediately, we can ignore the possibility of treating voluntary control of beliefs on the model of nonbasic but immediate voluntary control of doors being open and lights being on.

---

[10] Of course, opening a door or turning on a light may, in special cases, require repeated attempts with intervals between. In the above I was speaking of the simple unimpeded cases.

This brings us to a second grade of what is commonly called "indirect voluntary control", what I will call *long-range voluntary control.* It will be noted that the types of voluntary control I am considering are arranged in an order of increasing indirectness, increasing distance from the most immediate control. Here, as with immediate nonbasic voluntary control, we think of the belief as being produced by the carrying out of an intention by one or more actions that are designed to produce the belief rather than as being produced by a mere act of will, choice, or decision. But unlike the last case, the belief production is not carried out in one uninterrupted action. It involves a series of actions spread out over a greater or smaller period of time, the smallest period of which is too extended to accommodate a single uninterrupted act.[11] A number of voluntarists seem to be thinking in these terms of the cases in which it is not immediately obvious whether a given proposition is true or false. After all, they say, that is what inquiry is for, to resolve such issues. One certainly has voluntary control over whether to keep looking for evidence or reasons, and voluntary control over where to look, what steps to take to find relevant considerations, and so on. It is suggested, in effect, that since we have voluntary control over these intermediate steps, this amounts to what I call *long-range voluntary control* of a propositional attitude. Chisholm, for example, says:

> If self-control is what is essential to activity, some of our beliefs, our believings, would seem to be acts. When a man deliberates and comes finally to a conclusion, his decision is as much within his control as is any other deed we attribute to him. If his conclusion was unreasonable, a conclusion he should not have accepted, we may plead with him: "But you needn't have supposed that so-and-so was true. Why didn't you take account of these other facts?" We assume that his decision is one he could have avoided and that, had he only chosen to do so, he could have made a more reasonable inference. Or, if his conclusion is not the result of a deliberate inference, we may say, "But if you had only stopped to think", implying that, had he chosen, he could have stopped to think. We suppose, as we do whenever we apply our ethical or moral predicates, that there was something else the agent could have done instead. (1968, 224)

---

[11] Obviously, there is no precise boundary between a "single uninterrupted act" and a "series of temporally extended acts" where the temporal extension is small. But there are enough clear cases on either side of the distinction to make it usable.

To be sure, the mere fact that one often looks for evidence to decide an unresolved issue does not show that one has voluntary control over one's propositional attitudes. That would also depend, at least, on the incidence of success in these enterprises. And sometimes one finds decisive evidence and sometimes one doesn't. But let's ignore that complexity and just consider whether there is a case for long-range voluntary control of belief in the successful cases.

No, there is not, and primarily for the following reason. Claims like those in the quote from Chisholm ignore 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 that the phenomenon of looking for more evidence would show that we have voluntary control over propositional attitudes, it would have to be the case that the search for evidence was undertaken with the intention of taking up a *certain* attitude toward a *specific* proposition. For only in that case would it have any tendency to show that we have exercised voluntary control over *what* propositional attitude we come to have. Suppose that I can't remember Al Kaline's lifetime batting average, and I look it up in the baseball almanac. I read there the figure .320, and I thereby accept it. Does that show that I have voluntary control (of any sort) over my belief that Kaline's lifetime batting average was .320? Not at all. At most it shows that I have long-range voluntary control over whether I take up *some* propositional attitude toward *some* proposition ascribing a lifetime batting average to Kaline. So this is not at all parallel to cases where we definitely do have some (albeit fallible) long-range voluntary control over other sorts of affairs. Suppose that I can perform voluntary actions that will result, subject to the usual chances that infect all human endeavor, in my losing twenty pounds. Here there is a completely definite and unique result toward which my voluntary efforts are directed, and success, or at least repeated success, will show that I do have long-range voluntary control (within limits) of my weight.

What the situation described by Chisholm is closely analogous to is the following. 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 of many aspects of the household, including doors. Each morning he leaves detailed instructions on household operations in a computer. Doors can be operated only through the computer in accordance with his instructions. There is no way I can carry out an intention of my own, no matter how long range, to open or to close a particular door at a particular time. All I can is to 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 positions he specified, just as in the Kaline case I was responsible for taking up some attitude or other toward some proposition within a given range. But I definitely am not responsible for the front door's being open rather than closed at a particular time, nor can I be said to have voluntary control over its specific position. Hence it would be idle to apply deontological concepts 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. I had no control over that; it was not subject to my will. And that's the way it is where the only voluntary control I have over my propositional attitudes is to enter onto an investigation that will eventuate in some propositional attitude or other on what is being considered.

Or consider propositions concerning what is visible. I have the power to voluntarily open my eyes and look about me, thereby putting myself in a position, when conditions are favorable, to reliably form propositions about the visible environment. Again, with respect to past experiences, I can "search my memory" for the details of my experiences of the middle of yesterday, thereby usually putting myself in a good position to form beliefs reliably about my experiences at that time. No one, I suppose, would take these facts to show that I have voluntary control over what I believe about the visible environment or about my remembered experiences. What I can control voluntarily is whether I form (or am in a position to form) some accurate beliefs or other about my current visible environment or about my experiences of yesterday. And yet this is the same sort of thing as the search for additional evidence of which Chisholm speaks, differing only in the type of belief-forming mechanisms involved.

I suspect that those who take positions like the one in the passage just quoted from Chisholm secretly suppose that the additional evidence, rather than "automatically" determining the propositional attitude, simply puts the subject in a position to make an informed choice of an attitude. That is, they really locate the voluntary control in the moment of attitude formation rather than in the preliminary investigation, thereby in effect taking the (basic or nonbasic) immediate-control position. But then, faced with the implausibility of those positions, they think to save the application of deontological concepts to beliefs by pushing the voluntary control back to the preliminary search for decisive considerations. But their undercover attachment to the immediate-control thesis prevents them from seeing that voluntary control of the investigative phase has no tendency to ground the deontological treatment of propositional attitudes themselves.

Despite the above arguments against false pretensions to the title of "long-range voluntary control of belief", I have no intention of suggesting that there could not be legitimate claimants. Let's take a fresh start and lay out what it takes for a genuine case of such control in general (not restricted to beliefs). It requires the capacity to bring about a state of affairs, C, by voluntarily doing a number of different things over a considerable period of time, typically interrupted by activity directed to other goals. One has this sort of control, to a greater or lesser degree, over many things: one's weight, cholesterol concentration, blood pressure, and disposition; the actions of one's spouse or one's offspring. One can, with some hope of success, set out on a long-range project to reduce one's weight, improve one's disposition, or get one's spouse to be more friendly to the neighbors. The degree of control one is likely to have varies markedly among these examples. But all these examples and many more illustrate the point that one can have long-range control over many things over which one lacks immediate control. I cannot markedly reduce my weight right away by the uninterrupted carrying out of an intention to—by taking a pill, running around the block, or saying 'Abracadabra'. But that doesn't nullify the fact that I have some degree of long-range control.

To return to our main concern, it does seem that we have some degree of long-range voluntary control over at least some of our beliefs. People do set out on long-range projects to get themselves to believe a certain proposition, and sometimes they succeed in this. Devices employed include selective exposure to evidence and deliberate attention to supporting considerations, seeking the company of believers and avoiding nonbelievers, self-suggestion, and more bizarre methods like hypnotism. By such devices people sometimes induce themselves to believe in God, in materialism, in communism, in the proposition that they are loved by X, and so on. Why doesn't this constitute a kind of voluntary control that grounds deontological treatment of beliefs?

Well, it would if we *do* have sufficient control of this sort. People could properly be held responsible for their attitudes toward propositions in a certain range only if those who set out to intentionally produce a certain attitude toward such a proposition and made sufficient efforts were frequently successful. For only if we were generally successful in bringing about goal G when we try hard enough to do so, do we have effective control over whether G obtains. And if I don't have effective control over G, I can hardly be held to blame for its nonoccurrence. This is a generally applicable principle, by no means restricted to beliefs. If I am so constituted that the most I can do with respect to my irritability is to make it slightly less likely that it will exceed a certain high average threshold, I can hardly be blamed for being irritable.

It is very dubious that we have a reliable long-range voluntary control over any of our beliefs, even in the most favorable cases, such as beliefs about religious and philosophical matters and about personal relationships. *Sometimes* people succeed in getting themselves to believe (disbelieve) something. But I doubt that the success rate is substantial. I know of no statistics on this, but I would be very much surprised if attempts of this sort bore fruit in more than a small proportion of the cases. In thinking about this, let's first set aside cases in which the attempt succeeds because the subject happens onto conclusive evidence that would have produced the belief anyway without deliberate effort on his part to produce that belief. Thus we need to consider cases in which the subject is swimming against either a preponderance of contrary evidence or a lack of sufficient evidence either way. That is, S is fighting very strong tendencies to believe when and only when something seems true to her. Whether these tendencies are innate, or engendered and reinforced by socialization, they are deeply rooted and of great strength. To combat them one must exercise considerable ingenuity in monitoring the input of information and in exposing oneself to nonrational influences. These are tricky operations, and it would be very surprising if they were successful in a significant proportion of cases. I am not suggesting that it is unusual for people to form and retain beliefs without adequate grounds. That is all too common. But in most such cases the proposition in question seems clearly true to the person, however ill supported. The typical case of prejudice, for example, is not one in which S manages to believe something contrary to what seems to him to be the case or something concerning which he has no definite impression of truth or falsity. It is a case in which his socialization has led it to seem clearly true to him that, for example, blacks are innately inferior.

Thus the possibility of long-range voluntary control of beliefs does not provide significant grounding for deontologism, even for the sorts of propositions people do sometimes try to get themselves to believe or disbelieve. Much less is there any such support for deontologism for those propositions with respect to which people don't normally even try to manipulate their attitudes. We have already noted that most of our beliefs spring from doxastic tendencies that are too deeply rooted to permit of modification by deliberate effort. In such cases the project of deliberately producing belief or disbelief is one that is never seriously envisaged. Thus, even if we were usually successful when we set out to produce a propositional attitude, the voluntary control thus manifested would not ground the application of deontological concepts to beliefs generally.

iv. Indirect Voluntary Influence on Believing

Up to this point I have been considering various ways in which believing, rejecting, and withholding propositions might be themselves under effective voluntary control. We have seen that for most of our beliefs we have no such control and that for the others we have at most some spotty and unreliable control of the long-range sort. But this is not the end of the line for the prospects of an epistemic desideratum of the deontological kind. There is still one more way in which subjects might be held responsible for their believings, for believings to be required, forbidden, or permitted, even though they themselves are not under effective voluntary control. To see this, consider the general point that we can be blamed for a state of affairs F, provided something we voluntarily did (didn't do) and should have not done (done) was a necessary condition (in the circumstances) of the realization of F. That is, F would not have obtained had we done (not done) something we should have done (not done). If my cholesterol buildup would have been prevented had I regulated my diet in the way I should have done, but didn't, I can be blamed for that buildup, whether or not I have direct effective voluntary control of my cholesterol level.

Applying this general point to beliefs, we can say that even though believings are not under effective voluntary control, we can be blamed for holding a certain belief, B, if there are things we can voluntarily do such that we should have done (not done) them and if we had done (not done) them we would not have held that belief. Suppose that I accept some idle gossip to the effect that Jim is trying to undermine Susie's position as departmental chair. If I had done what I should have done by way of checking into this matter, I would not have formed that belief or would not have retained it for as long as I did. Hence I could be blamed for holding the belief. To take a case where I am blameless in holding a belief, consider a visual belief where my vision and my belief-forming mechanisms are working normally. There is nothing relevant to that belief formation that I should have done but didn't, and so I am not subject to blame in forming the belief. Note that other deontological terms like 'ought' and 'should' are also applicable to states of affairs not themselves under direct voluntary control in this derivative way. Thus we can say that I *ought* to have a lower cholesterol count and that I *should* not have believed that he did it.

Note that this kind of application of the deontological categories of blameworthiness or the reverse to believings is a derivative one. It is the (actual or possible) voluntary acts in the causal ancestry of the belief to which blameworthiness and other deontological terms of evaluation apply in a primary way. My checking, or failing to check, on the accuracy of the gossip is something that is directly, underivatively blameworthy or the reverse. The propositional attitude that eventuates is blameworthy or the reverse only by derivation from the voluntary acts that give rise to it. Strictly speaking, in thinking of the matter in this way we should not think of the forming of the propositional attitude itself as required, forbidden, or permitted since it itself is not under effective voluntary control. Its blameworthiness or blamelessness stems from the required, forbidden, or permitted voluntary acts in its causal ancestry. Harking back to the initial list of candidate deontological desiderata, we have now moved from

9. B is held permissibly (one is not subject to blame for doing so).

to

11. The causal ancestry of B does not contain violations of intellectual obligations.

I will call this kind of impingement of the voluntary on belief *indirect voluntary influence.*[12]

I will now proceed to put a bit of flesh on the idea of an indirect influence of voluntary actions on propositional attitudes. First, let's note that we do have voluntary control over many actions that can influence our believings, rejectings, and withholdings. These can be divided into two groups: (a) activities that bring influences to bear, or withhold influences from, a particular candidate, or field of candidates, for belief, and (b) activities that affect our general belief-forming habits or tendencies. There are many examples of (a). With respect to a particular issue, I have voluntary control over whether and how long I consider the matter, over whether and where I look for relevant evidence or reasons, reflect on a particular argument, seek input from other people, search my memory for analogous cases, and so on. Here we come back to the activities we saw Chisholm wrongly classifying as the intentional inauguration of an attitude toward a specific proposition. Group (b) includes such activities as training myself to be more critical of gossip, instilling in myself a stronger disposition to reflect carefully before making a judgment on highly controversial matters, talking myself into being less (more) subservient to authority, and practicing greater sensitivity to the condition of other people. It is within my power to do things like this or not, and when I do them with sufficient assiduity I make some difference to my propositional attitude tendencies, and thus indirectly to the formation of such attitudes.

---

[12] This kind of relation of voluntary action to belief is given much less attention in the literature than the kinds I have judged to lack viability. Nevertheless, as an important aspect of cognitive endeavors, it deserves more attention than it receives. I should also mention that in Alston 1988a I argue that a concept of epistemic justification based on this kind of voluntary influence on belief is not adequate, not because there is no such influence but rather because it is not closely enough related to the goal of true belief. Though I am not concerned here with how to construe epistemic justification, I will make a similar point about treating 11 as an *epistemic* desideratum.

There would be no harm in including in (a) attempts to bring about a particular attitude to a specific proposition. For these too would be things that influence our propositional attitudes and over which we have voluntary control. The point of stressing other things is that since the earlier discussion provided reason for thinking such attempts are rarely successful, I want here to emphasize the point that even if we are never successful in voluntarily bringing about a belief that p, there are still many things we can do voluntarily that do have a bearing on what propositional attitudes are engendered.

The next question is whether the deontological triad of concepts applies to activities like those canvassed in the next-to-last paragraph. Is it ever the case that we ought or ought not to engage in an activity of these sorts, such as searching for new evidence or critically examining the credentials of gossip? Is it ever the case that we ought or ought not to strive to make ourselves more (less) sensitive to contrary evidence? Deontologists typically hold that we have intellectual obligations in such matters, obligations rooted in our basic intellectual obligation to seek the true and avoid the false in belief. I accept this view, which seems eminently plausible.

Thus it will sometimes be the case when I believe that p that I would not have done so had I done various things in the past that I could and should have done but failed to do, and it will sometimes be the case that I would not have believed that p had I not done various things in the past that I could and should not have done but did. In either of these cases there is a failure of obligations in the causal ancestry of the belief that renders me blameworthy for having the belief. And if neither of these is the case, then I am blameless, not properly held to blame for the belief. (All this applies equally to rejectings and withholdings). Hence the indirect voluntary influence on the formation of propositional attitudes does have an evaluative bearing on those attitudes, either positive or negative.

I must pause to refine the above formulation. There are certain ways in which dereliction of intellectual duty can contribute to belief formation without rendering S blameworthy for forming that belief. Suppose that I fail to carry out an obligation to spend a certain period in training myself to look for counterevidence. I use the time thus freed up to take a walk around the neighborhood. In the course of doing so I see two dogs fighting, thereby acquiring the belief that they are fighting. There was a relevant intellectual obligation I didn't fulfill, which is such that if I had fulfilled it I wouldn't have acquired that belief. But if that is a perfectly normal perceptual belief, I am obviously not to blame for having formed it.[13]

Here the dereliction of duty contributed to belief formation simply by facilitating access to data. That is not the kind of contribution we had in mind in the above formulations. The sorts of cases we had in mind were those most directly suggested by the two sorts of voluntary activities that affect belief formation: (a) those that involve looking for considerations relevant to the belief in question, or not doing so, and (b) those that affect our general belief-forming habits or tendencies. By revising 11 so as to make this explicit, we can avoid counterexamples like the above.

11A. S is intellectually to blame for believing that p *iff* if S had fulfilled all her intellectual obligations, then S's access to relevant considerations, or S's belief-forming habits or tendencies, would have changed in such a way that S would not have believed that p.[14]

It follows from the above that 11A is a genuine intellectual desideratum, a desirable feature of belief outputs of cognition. And it also follows from the above discussion that, unlike 9, 11A does not fail to qualify as an epistemic desideratum because it is not a real possibility for human beings. But though it is a cognitive desideratum, it will still not be an *epistemic* desideratum if it is not connected in the right way with the truth goal. And how do we determine whether it is? Since we have identified three groups of cognitive desiderata that are, in their several ways, related to the truth goal so as to qualify as epistemic desiderata, an obvious way to proceed is to consider whether 11 is related to the truth goal in one of those ways.

I think we can straightaway eliminate the Group III and the Group V ways from consideration. As for III, it is obvious that where there are no violations of intellectual obligations in the ancestry of a considerable number of S's beliefs, this in no way provides S with resources for determining under what conditions a belief is likely to be true and thereby putting S in a good position to restrict belief formation to true beliefs, in the way higher-level epistemic knowledge or the capacity for such does. Nor does 11A presuppose such a capacity as 8 does. As for Group V, although it could be argued that fulfillment of intellectual obligations has an intrinsic cognitive value that is independent of truth, as 12–15 do, that value is in no way dependent on being associated with a preponderance of true beliefs, as is the case with the Group V desiderata. Even if S's beliefs are preponderantly false, it is a good thing for S to fulfill her intellectual obligations, insofar as this is possible in the face of something less than an outstanding record of achieving the truth.

---

[13] I am indebted to Emily Robertson for calling this problem to my attention.

[14] Another fine-tuning point has to do with the "absoluteness" of the counterfactual involved. Desideratum 11A says S *would not have believed that p* under these conditions. But perhaps S is also blameworthy for believing that p even if it is only much less likely that S would have believed that p under these conditions. I am inclined to accept this weaker interpretation, but I will not press the matter here since I will go on to reject 11 as an epistemic desideratum anyway.

So this leaves us with Group II, the directly truth-conducive desiderata. Though I will argue that 11A is not related to the truth goal in this way either, I do not think that this is obvious on the face of it, as is the case with the Group III and Group V desiderata. Hence this suggestion deserves more serious consideration. It is prima facie conceivable that being formed in a way that does not depend on violations of intellectual obligations should be a way of rendering a belief probably true. Nevertheless, there are many counterexamples to such a claim, some of which I shall now make explicit.

Before we embark on this, something needs to be said about what would constitute failure to fulfill an intellectual obligation. What makes this difficult is that attempts to conform to an obligation might or might not be successful, and where they are not they might be more or less close to achieving it. On a sufficiently rigorous interpretation we are almost always failing in some intellectual obligation or other. This in turn depends on just how it is specified what one is intellectually obliged to do. Consider the obligation to look for relevant considerations pro and con when it is not clear whether the proposition in question in true. How long and how assiduously does one have to look to fulfill the obligation? To require that every conceivably relevant consideration must be taken into account would be a counsel of perfection that is beyond any of our powers, not to mention the fact that we couldn't know whether that limit had been reached. It seems that to make the notion of fulfilling intellectual obligations usable we have to build in a limitation to what could reasonably be expected of a subject, and that is itself is a very imprecise notion and one that is subject to varying interpretations. Moreover, there is the point that what could be reasonably expected along this line will vary for different people in accordance with their abilities, experience, education, propensities, and so on. So the question whether a given subject has done as much as could be expected of him or her is beset with uncertainties, imprecision, and disagreements.

Keeping all this in mind, let's do the best we can by proceeding on the basis of some sense of what could be expected of a given subject and on the basis of some plausible construal of the content of intellectual obligations. I now want to suggest that there are very many sorts of cases in which one does as much as could be reasonably expected of one in the way of voluntary acts leading up to a given belief without the belief's thereby acquiring any considerable likelihood of truth.

First, there are cases of cognitive deficiency. Consider one who forms the belief that socialism is contrary to Christianity for the reasons that are often given for this view by the Christian Right, and is intellectually incapable of figuring out how bad these reasons are, cases that I fear are all too common. Such a person may have done as well as could be expected of him in coming to this belief, but that fact does nothing to make the belief likely to be true. (Thoroughly bad reasons are not truth-conducive). Or consider a college student who doesn't have what it takes to follow abstract philosophical exposition or reasoning. Having read parts of Bk. IV of Locke's *Essay Concerning Human Understanding*, he takes it that Locke's view is that everything is a matter of opinion. He is simply incapable of distinguishing between that view and Locke's view that one's knowledge is restricted to one's own ideas. There is nothing he could do that would lead him to appreciate the difference. Hence he cannot be blamed for interpreting Locke as he does; he is doing the best he can to fulfill his intellectual obligations. But his belief about Locke's view is outrageously ill-grounded, based as it is on his dim-witted impression of Bk. IV of the *Essay*.

Second, consider the innumerable beliefs each of us forms on testimony or authority. Practically everything we believe about science, history, geography, and current affairs is taken on authority. Ideally, we would check out each source to make sure that it is reliable before accepting the testimony. But who has time for that? We can do it in special cases where the matter is of special importance, but it is not a real option for such beliefs generally. If we tried to do so, our doxastic structure would be so impoverished that we would not be able to function in society. Moreover, even if we had time to check up on each authority, in most cases we lack the resources for making an informed judgment. Thus in most cases in which I uncritically accept testimony I have done as much as could reasonably be expected of me. Now consider those cases in which the authority is incompetent or the witness is unreliable. There we are forming a belief on an objectively unreliable basis and hence the belief is not probably true, even though no dereliction of intellectual duty is in the background.

Next consider irresistible beliefs and belief tendencies. If it is impossible for me to alter a certain belief or belief tendency, I can hardly be expected to do so. But some irresistible beliefs are formed in an unreliable fashion in such a way as not to make them likely to be true. The most obvious examples concern strong emotional attachments that are unshakable. For many people their religious or irreligious beliefs have this status, as do beliefs concerning one's country, one's close relations, or one's political party. Such beliefs are often not formed in a truth-conducive way that would render them likely to be true. But S cannot be blamed for doing or having something she can't help doing or having. Here too the believing's not depending on any violation of intellectual obligations is no (even fallible) guarantee of the belief's likelihood of being true.

Of course, showing that 11A is not related to the truth goal in any of the ways the desiderata already approved as epistemic are is not a proof that it is not an epistemic desideratum. Perhaps it is related to the truth goal in some other way that qualifies it as epistemically desirable. That is an abstract possibility. Why shouldn't there be a fourth way, as different from the first three as they are from each other? But I must confess that I have found no fourth way. And until I do, I must rest with the conclusion that 11A does not qualify as an epistemic desideratum.

Thus none of the deontological candidates makes the grade. The first three (alternative versions of 9) fail through the failure of the version of voluntary control of belief presupposed by each to be a real possibility for human beings. And the last (11A) fails through not being connected with the truth goal in the right kind of way. Thus they will receive no further attention in the development of the ED approach to the epistemology of belief.