

Belief is Involuntary

No one has a choice as to what to believe. No one can, simply by an exercise of his will, choose or decide to believe *p*. Belief is involuntary, as I shall say and this means no more, nor less, than that belief is ‘choiceless’, i.e. such that it cannot be acquired simply choosing to do so—beliefs cannot be acquired at will, not by *fiat*. To forestall possible misunderstanding, let me immediately make clear what this proposition does and does not mean to deny. Suppose I now have no beliefs about where the world’s oceans are deepest and want to have that situation rectified. Then I can exercise my will in such a way that I wind up believing that it is at location *L* that the oceans are deepest. All I have to do for this is consult an encyclopaedia, or an authoritative source on the deep seas. Likewise I can decide about any number of topics that I want to have beliefs about them, and by exercising my will, I can make inquiries as a result of which I acquire such beliefs. The thesis that belief is involuntary does not mean to deny anything of what I have just written. What it does mean to deny is that I can believe that at *L* the oceans are deepest just by deciding to do so. More precisely, what it means to deny is that for any specific proposition *p* that we care to think of, we can bring it about, simply by exercising our will, simply by deciding to do so, that we believe that *p*. This chapter aims to offer support for this claim.

But some further initial clarification of the claim is called for, for there are several ways in which the concept of ‘voluntary belief’ (or ‘belief by *fiat*’, or ‘belief at will’) can be understood. As I will be thinking of it, that concept cannot be applied to a belief unless the following conditions are satisfied:

- (A) the belief is acquired independently of any truth-relevant considerations on the part of the subject
- (B) the subject is fully aware of the fact that he attempted to acquire the belief independently of any truth-relevant considerations.¹

¹ Winters 1979: 244-5 provides a somewhat similar account of ‘voluntary belief’ (or as she says ‘believing at will’), but also specifies the condition “that the belief has been acquired directly and as a result of intending to hold it”. As will become clear in the first section, I see no reason to accept this condition. A belief (if such there is) that has been acquired *indirectly* (and as a result of intending to hold it), might still be a case of ‘voluntary belief’.

The argument of this chapter, as indicated, is that belief is involuntary—and this means that its burden is to show that no belief is such that with respect to it conditions (A) and (B) are satisfied.

The claim that belief is involuntary can still be taken in various ways. It can be taken, if true, as contingently true. But it can also be taken, if true, as necessarily true. Taken the first way, it is the claim that *Human belief is involuntary* is true in the actual world but that there are possible worlds in which it is false. Taken the second way it is the claim that *Human belief is involuntary* is true not just in the actual world but in all possible worlds. If *Human belief is involuntary* is only contingently true, then the involuntariness of belief is a contingent property of humans; but if it is necessarily true, then it is an essential property of humans. The claim that belief is involuntary can yet be taken in a third and stronger way, viz. as the claim that for *any* being with a will and capable of belief, belief is involuntary. In the next two sections I will argue for the claim that with respect to us, humans, belief is involuntary in the actual world. Then I will go on, in section 3, to offer an argument in favour of the stronger thesis that human belief is essentially involuntary (so I will argue that *Human belief is involuntary* is necessarily true). I will finally suggest that the strongest claim (the claim that for any being with a will and capable of belief, belief is involuntary) is false.

1. Simple empirical arguments for the involuntariness of belief

The arguments for the involuntariness of human belief that I am now going to discuss are ‘simple empirical’ arguments because they don’t make use of the elaborate empirical research as done at, for instance, psychology labs. In order to see what these arguments amount to, we must first note that the complement of involuntariness is a rather rich set. Something can be voluntary in more than one way, and in order to refer to those ways I am going to use distinctions and labels due to William Alston.² To begin with, something can be voluntary for a person in the way that, for a healthy person, raising his arm is voluntary. All that is needed for that is to carry out the intention to raise one’s arm in one uninterrupted intentional act. The healthy person can raise his arm “just like that”, without having to return to the attempt a number of times. If someone *cannot* raise his arm “just like that” and has to make more than one attempt (perhaps he suffers from a muscular disease, or has a weight of 150 kg tied to his hand), then that act is not voluntary in the sense now intended. This kind of voluntariness

² Alston 1988: 137; also Alston 2005: 58-80.

might be labelled *basic immediate control*. The reason why it is called ‘immediate’ has to do with the “just like that” character of acts that are voluntary in this sense. The reason why it is called ‘basic’ becomes clear upon considering a second way in which an act can be voluntary.

Consider such actions as starting one’s car’s engine and making someone happy by paying him a compliment. One can do these things, just as one can raise one’s arm, in one uninterrupted intentional act, “just like that”, without having to make more than one try. But in order to succeed in doing these things more is needed than just a volition on the part of the agent. In each of these cases the agent’s bodily movements (movements of the hand, and movements of the vocal chords, respectively) must have certain consequences—consequences that are either causally (in the case of starting the car) or conventionally (in the case of making someone happy) connected with his bodily movements. Starting the car’s engine and making someone happy (in the way indicated) are certainly voluntary acts—and they can be done “just like that” (at least, if circumstances cooperate). So these acts are under one’s immediate control. But they are not ‘basic’ in the way that raising one’s arm is, because these actions are causal or conventional *consequences* of acts that *are* within one’s basic immediate control (moving parts of one’s body). The kind of voluntariness involved in cases like these might therefore be labelled *non-basic immediate control*.

There is a compelling argument for the conclusion that for humans in the actual world belief is neither under basic immediate voluntary control, nor under non-basic immediate voluntary control. The argument³ consists in asking yourself the simple question: Can you, at this very moment, bring it about that you believe that the earth has three moons, just by deciding to do so? Can you now believe that the canals on Mars once held water, “just like that”, by *fiat*? In case you are insufficiently motivated to even try, suppose you are offered 500 dollars if you would succeed. Could you then do what it takes to get the reward? All of this seems utterly impossible. The conclusion is that belief is involuntary in the sense that we have no immediate voluntary control over it, neither basic nor non-basic.

This conclusion holds not only for propositions that are false. Even if a proposition is true, we cannot, by a mere act of our will, believe it. Surely, if it is *obvious* to you that a proposition is true, you shall believe it. But this belief is not the result of a voluntary act on your part. The belief passes the will without even being touched by it, so to speak. This also holds for propositions about which we have no clue as to whether they are true or not. Such propositions we cannot decide to believe nor decide to disbelieve. We withhold judgment on

³ Also offered by Alston 1988: 122.

them. But withholding judgement on *p* is not touched by our wills either—we don't *decide* to neither believe nor disbelieve *p*. We simply find ourselves neither believing nor disbelieving.

As the previous paragraph indicates, it isn't only *belief* that is involuntary in the ways specified. *Disbelief* and *withholding* are likewise involuntary. Or as we may say, all doxastic attitudes (belief, disbelief and withholding) are involuntary in the sense of being outside our immediate voluntary control (either basic or non-basic). For ease of exposition I shall in the main speak of belief, but many of the points I will be making hold good for the other doxastic attitudes as well.

In order to elucidate this thesis further, it will be useful to consider a couple of cases that might be thought to controvert it. First there are cases of decision making under uncertainty. An officer in war time, a mountaineer who lost his way, parents whose child escaped from their eyesight have to decide for a certain course of action. They have to decide whether or not to launch an attack, which way to go, and where to look. It surely isn't the case that any decision here would be as good as any other—but still there are, or so we may suppose, genuine options. Whichever option is selected, however, one doesn't thereby *decide to believe* that that option will lead to success. In such cases one decides to proceed in a certain way; one does not *decide to believe* that proceeding in that way will bring victory, lead back to the right track, or guarantee that the child will be found.

Second, scientists may perform experiments on the basis of a certain assumption or working hypothesis just to see what results emerge. Depending on the results, such assumptions and hypotheses can be readily abandoned. Adopting and abandoning an assumption, however, does *not* mean that one first has decided to believe it and then decides to no longer believe it. Surely, in a case like this decisions are made. But what is decided is to proceed on a certain assumption and next to no longer do that. Belief is kept at bay.

Third and final: we often invoke general principles such as “don't believe unless you have evidence”. Now suppose someone who heeds that principle finds a threatening letter on his doormat. Naturally enough, he aims to find out who wrote the letter and looks about for clues and evidence, he meticulously studies the stamp, the handwriting, stylistic features etc. If things are going well, he might wind up believing that, say, Moriarty is the author. Doesn't this indicate that the threatened man has voluntary control over what to believe? Doesn't this controvert the thesis that belief is involuntary? It does not. And this for two connected reasons. First, as was explained in the very first paragraph of this chapter, that thesis doesn't mean to deny that one can do certain things that will make it likely that one will believe *something or other* about a particular topic. What it means to deny is that for any specific

proposition *p*, such as *Moriarty is the author*, we can bring it about *by fiat*, that we believe *p*. Second, upon having made careful inquiries, the man who received the letter did not ever *decide* to believe that Moriarty was the writer. The formation of the belief, subsequent to the inquiries, was involuntary.

The field of the voluntary is not exhausted by immediate voluntary control, however. Consider the kind of voluntariness involved in losing weight through keeping a diet, and helping one's child to persevere in adversity. Losing weight and helping one's son cannot be performed "just like that" in one uninterrupted intentional act, but require many repeated efforts stretched out over longer periods of time. We have voluntary control over these actions, albeit only of an indirect sort. A further point here is that there exist considerable differences in the 'strength' of the control that we have over those things that are within the span of that sort of control. On some things the indirect control may give us a much firmer grip than on some other things. I may have a much firmer control over my weight than over my son's ability to persevere in adversity—and *you* may have a much stronger grip on *your* son's ability to persevere than I have on mine. Still, these cases belong to the same group. They are cases in which the voluntary control is not immediate but only indirect, which means that it is stretched out over a period of time. We can therefore call them cases of *long-range voluntary control*.

As indicated, there are differences in strength of long-range voluntary control. If that sort of control reliably brings about a certain effect, then it is strong. If that sort of control only occasionally brings about a certain effect, then it isn't really control. Control requires that one can, in a high percentage of cases, bring about a certain effect. This holds for immediate as well long-range voluntary control. For the discussion that is now to unfold, this point will prove rather important.

Is our belief under our long-range voluntary control? If it is, there will have to be things that we can *do*, and do them repeatedly, such that doing them will, in all likelihood, and after the elapse of some time, result in the formation of the belief. (The belief involved will still have to be the belief that *p*, with *p* being a specific proposition). What sort of things? The following have been suggested:

- a- pay selective attention to the evidence—only focus on evidence in favour of *p* and turning one's head away from any evidence against *p*,
- b- seek the company of believers in *p*, and avoid the company of believers in non-*p*,

- c- apply some method of self-suggestion, for instance by saying aloud or within yourself “yes, p—and I believe that p”.

These are the most plausible candidates for being modes or methods of long-range voluntary control over belief in p.

Let us look into these possibilities somewhat more deeply, and suppose we start with a-, paying selective attention to evidence. A number of cases, or types of cases, have been described with the aim of establishing that paying selective attention to evidence gives us long-range voluntary control over at least some of our beliefs. One case, due to Roy Edgley⁴, is of a man who is very sick and whose friends persuade him that it would be good for him if he were to believe that he will recover. The sick man is persuaded by them and accordingly tries to believe that he will recover by dwelling on favourable evidence, and by making specific plans about what to do after regaining health. Edgley believes this man may succeed.

In order to assess the plausibility of Edgley’s belief, we need to look into the notion of evidence. That notion is very complex and there are many theories about it.⁵ A first thing to note is that in the expression ‘Paying attention to the evidence’, ‘evidence’, although a singular expression, refers to a plurality. Just like ‘the Dutch’ it is an expression of collective reference. A second thing to note is that ‘evidence’ refers to confirming as well as undermining evidence. Evidence is always evidence *for* or *against* something. This brings me to a third thing we need to note: ‘paying selective attention to the evidence’ can mean two different things, depending on how one understands the notion of evidence. On one understanding, something E is evidence for or against p, provided E either supports or undermines p. On another understanding E is evidence for or against p, provided someone who is aware of E and p, notices that E supports or undermines p. On the first, *objective*, understanding of evidence, the sea’s tides are evidence for gravitation. But on the second, *subjective*, understanding of evidence, the sea’s tides are evidence for gravitation only for someone who is *aware* of the sea’s tides and of the proposition that masses attract one another, and furthermore *notices* that the first supports the second.

‘Paying selective attention to evidence’, with ‘evidence’ understood subjectively, would come to this: one is aware of E and p, and even notices that E supports or undermines p, but either the awareness of E or the noticing of the support/undermining relation (from now

⁴ Edgley 19@@:63-4.

⁵ Some hold that evidence is exclusively propositional in nature, whereas others hold that experiences and even objects can be evidence as well. There are furthermore subjective as well as objective theories about evidence, and different theories of objective evidence at that (Bayesian, hypothetico-deductive, and satisfaction theories).

on I will use these two words in a quasi-technical way to refer to the two different phenomena) is frail, feeble, halting, suppressed, blocked out, or in some other way objectionable. And this not because one lacks the required powers of awareness and noticing, but because of something else: *one doesn't pay sufficient attention* to, *one ignores*, what one *is* aware of and *does* notice—as wise parents occasionally will ignore observed reproachable behaviour of their child. It should be noted that the awareness of E can take the form of being explicitly aware of E, but also the rather common but diffuse form of ‘having heard it say that there is evidence’ for or against a certain proposition, without actually knowing *what* the evidence is.

It is less clear what paying selective attention to evidence, with ‘evidence’ understood objectively, would come to. The only sensible explication in the neighbourhood would be that one pays full attention to objective evidence for p that is, unbeknownst to one, one-sided, i.e. only the evidence *for* or evidence *against* p, not the fusion of them. But paying attention to what is, unbeknownst to one, selected evidence, is something altogether different from paying selective attention to the evidence. I therefore conclude that the only sensible explication of ‘paying selective attention to evidence’ is an explication in which ‘evidence’ is understood subjectively.

The question now is whether paying selective attention to the evidence (understood subjectively) regarding p, can be used as a policy to insure that one comes to believe a specific proposition p? The question is not whether we do sometimes pay selective attention to evidence. “It is as easy to close the eyes of the mind, as those of the body”, said Bishop Butler⁶, and virtually everybody who has pronounced on this issue says the same. There seems to be ample reason for that consensus, whether we like it or not. The question is not whether we sometimes pay selective attention to attention (we do) but whether paying selective attention to the evidence regarding p can be a policy we can adopt when we want to bring it about that we acquire the belief that p? Can we, by adopting such a policy, have the kind of long-range voluntary control over our beliefs that we are currently discussing?

I must say I am very sceptical about this for several reasons. If paying selective attention to evidence is to be a policy that can be adopted in order to attain the specific goal of beginning to believe that p, then adoption of that policy must ensure, or make it highly likely, that the goal will be reached. For if it doesn't, adopting the policy doesn't provide control, hence no voluntary control, not even of the long-range variety. But adoption of the policy

⁶ Butler 1889: 480.

doesn't guarantee, or make it highly likely, that the belief-goal will be reached. The point is that paying selective attention regarding *p* so as to ensure belief in *p* isn't something one can consciously *do*. It is, rather, something that we sometimes seem to be doing, at best half-consciously. My argument is of the simple sort I used above: try it! You won't succeed. Surely, we sometimes pay no attention to claims of the sort "Here is good evidence against your belief that *p*". But that is something different. For sometimes there are very good reasons to discard such claims, e.g. because one knows the person who made the claim and is familiar with his confusions; or because one knows that whatever the contrary evidence will be, one's own supporting evidence will always be stronger.

Applied to Edgley's case of the sick man, this works out as follows. As the description of the case already makes clear, the hoped-for beliefs may *or may not* result from following the advice offered. There is no guarantee that they will. For this reason it seems very dubious that such cases are genuine cases of *control* even if they were to be successful. To have control over a certain result means that one can bring about that result with a rather high likelihood of success. It has been said that by worrying one's hair will get grey. Suppose John likes that colour and aims to get grey through worrying—and accordingly starts worrying a great deal about matters that thus far he took lightly. Then certainly we won't say that he has *control* over that effect. For it is radically unsure that worrying has that un-dyeing effect. Likewise, since it is radically unsure whether dwelling on favourable evidence and planning one's future will propel the belief that one will recover, it can't be said that the sick man has *control* over that belief—even if he were to be successful, or rather 'successful'. For without genuine control there can be as little success as in the case of a drunkard, a novice to the sport even, who shoots an arrow that, helped by an amazing gust of wind, nonetheless hits the target.

The point I am making is not that we never pay selective attention to evidence; the point is that we cannot adopt 'paying selective attention to evidence' as a *policy* that gives long-range voluntary control over believing *p*.

Another thing one could try in order to secure the formation of the specific belief that *p* is to seek the company of *p*-believers and avoid the company of believers in non-*p*. It is a well-known fact that meeting with a company of like-minded persons can be very reassuring and heartening. But that is not the sort of situation that is relevant to the discussion, for this company consists of persons who already share beliefs, whereas our discussion concerns whether someone can *acquire* the belief that *p* by meeting with a company of *p*-believers and shunning the company of non-*p*-believers. Nor is the following very real phenomenon

relevant to the discussion: people sometimes wish, for reasons of status, to become members of a certain social group, and, once entered, acquire particular beliefs shared by its members. This is not relevant, for here the reason to join a certain group is not the wish to believe that *p*, but the wish to have one's status boosted. The sort of situation that *is* relevant to the discussion would be something like this: *S* wants to believe that *p*, and notices that the members of group *G* all believe that *p*; *S* then decides to join *G*, talk to its members, do the things they do; in that manner, *S* thinks, he can ensure (or make it highly likely) that he will acquire the belief that *p*. If what *S* thinks is correct, he would have long-range voluntary control over the belief that *p*; but if it is not, then not.

One might think that whether what *S* thinks (*viz.* that by joining *G* he will ensure that he will acquire the belief that *p*) is correct, depends on what goes on in *G*. If its members, whenever they meet, say to each other again and again such things as 'p', and 'yes, p!', 'hurray for p', and 'three cheers for p, no four!' (thus expressing their belief that *p*), but never really discuss *p*, never consider any evidence for or against *p* (or if there is no evidence, never discuss why the absence of evidence should be no problem), never relate *p* to social practices that might be based upon them, then it seems highly unlikely that *S* will begin to believe *p* *due to joining G*. Joining *G* can't be used as a policy to ensure the formation of belief in *p*. But suppose the members of *G* go about differently, and do discuss evidence relevant to *p* (and do conclude each time that *p*), etc., then, yes, *S* might wind up believing *p*. But there is no assurance that that will happen, for much will depend on further details (well, they are not details) of the discussion among the members of *G*. If the discussion is biased, if relevant counter evidence is suppressed, if counter positions are 'refuted' only after having been mischievously misportrayed, etc., *and if S has some awareness of this*, then again it is surely not to be expected that *S* will wind up believing what he wanted to believe *due to joining G*. But what if the discussion among the *p*-believers *is* fair, open-minded, and due attention is given to undermining evidence, and contrary positions are criticized only after having been given a fair hearing, and *S has some awareness of this*—what then? Can *S* be assured, prior to joining *G*, that he will wind up believing what he wants to believe? Well, perhaps so, perhaps not, that is hard to tell in advance: it all depends on whether *S* will find the transactions of the *p*-believers compelling; perhaps the contrary positions have been described so fair that *S* is unmoved by the criticisms offered, or perhaps the social practices that are based on *p* (or involve *p* in some way or other) are so repelling, that *S* no longer finds belief in *p* attractive. But now suppose that *S*, prior to joining *G*, knows that many others that have joined *G* in the past, have wound up believing that *p* due to arguments and evidence relevant to *p*. Wouldn't

that assure him that the same will happen to him and hence that his joining G will ensure that he will eventually believe what he wants to believe? Well, it seems likely that this will so ensure him. But the thing to see is that the case as described is irrelevant to the discussion. For if S joins G and comes to believe what he wants to believe, then he doesn't come to believe p *by virtue of joining G*, but by virtue of something rather different, viz. *by virtue of the arguments offered by the members of G*, as well as *by virtue of the positively evaluated practices that involve belief that p*. And thus the case is irrelevant to the discussion, because the question was whether someone can ensure that he will believe what he wants to believe, *simply* by joining a certain company—irrespective of truth relevant considerations.

The conclusion should therefore be that whatever effects the joining of a certain company may have, *just* joining a company of p-believers cannot be adopted as a policy to ensure that one winds up believing what one wants to believe. And hence just joining a company of p-believers is not a way to have long-range voluntary control over the belief that p. (None of this implies a denial of the fact that one's joining a particular group might result in picking up beliefs that are indigenous in it).

Before moving on to a discussion of the third strategy I mentioned, self-suggestion, I should like to pause and reflect on a notion that I have been using for some time now—the notion that someone 'wants to believe that p'. Wanting to believe that p is of course not the same as deciding to believe that p. But there is a relation between the two, which is just an instantiation of the general relationship between wanting and deciding. One only *decides* to do what one *wants* to do, and hence one only *decides to believe* that p (assuming that is possible) when one *wants to believe* that p. Etiologically the want precedes the decision. But what is that, exactly, wanting to believe that p? When someone utters the words "I want to believe that p, but I can't", then what is it that he wants, and what is it that he can't?⁷ What he wants is: that p is true; he wants the world to be such that p is true. What he can't is: believe that p is true. Since you can't believe that p, unless you think that p is true, the reason why people sometimes can't believe that p even though they would want to, is that they don't think that p is true (perhaps they know p is false—even though they would greatly prefer the world to be such that p is true.) What I am saying, then, is that when people utter the words

- (i) I want to believe that p,

⁷ The same question could be posed in terms of 'wishing'. What is it that someone wishes when he wishes that p be true?

what they usually mean is:

- (ii) I want p to be true.

And this is only natural, given the explication in chapter 3 of what it is to believe something: to believe that p is to believe p is true.

However, the explication of (i) in (ii), can't be the explication that is needed in the discussion of the idea that seeking the company of p-believers will ensure that one will wind up believing what one wants to believe. There the want is not directed at the proposition that p, but at the belief attitude. What is needed for that discussion to make sense is that what one wants, is to have is a certain propositional attitude towards p, viz. the belief attitude (in contrast to a doubting or rejecting attitude).

Having said this, let us now discuss the third and final suggestion about what one can do in order to try to bring it about that one starts believing p: self-suggestion. It is part and parcel of folk wisdom that when people say something long enough to themselves, they are bound to believe that it is true. And this seems to be the idea about what is going on in self-suggestion: you say to yourself 'p', and 'yes, p', and 'p is true', 'that famous Mr. X believed p—and he was no fool', and the like, and as a result you wind up believing that p. Self-suggestion, then, can be taken to be this: you suggest p to yourself by saying to yourself 'p' and the like over a stretch of time, and as a result you acquire the belief that p. When stated thus bluntly, it would seem that if self-suggestion is to be a policy one can consciously adopt in order to acquire the belief that p, it is severely limited in its application. It seems incredible that simply by repeating to oneself 'the earth has three moons', one can bring it about that one starts believing so. And something similar holds for all sorts of factual propositions about history, geography, cosmology, culture, politics, jurisprudence and so forth. There would seem to be only two classes of propositions where self-suggestion might have a chance of success: a class of propositions about oneself, and a class of propositions about what is good and bad, as well as about what ought to be done, and what ought to be prevented—moral propositions for short. Let me go into these a little more deeply, and start with the class of propositions about oneself.

The propositions I am thinking of are of the following kind: *I am a fine fellow; I have nothing to write home about; I can't do math; I can do the triathlon; I am not going to be nervous when I have to perform in front of the queen; I will get that job; when I am there I am going to make a fool of myself; there was nothing wrong in what I did; I am not going to feel*

embarrassed when I see her. Can we, by saying (or thinking) these things to ourselves, ensure that we will eventually believe any of it?

This is a very difficult issue, but still some relatively uncontroversial points can be made. First, saying something to oneself about oneself, does seem to have genuine effects on the one saying (and hearing) it. The Frenchman Emile Coué (1857-1926) taught his patients to say to themselves every morning and every evening: “Tous les jours, á tous points de vue, je vais mieux et mieux” (Every day in every way I am getting better and better). And he added that these words must be spoken unthinkingly: “don’t think what you are saying, Say it as you say the litany at church!” This simple advice proved extraordinary effective: the Coué movement swept across Europe like a prairie fire.⁸

Second, the genuine effects of saying something to your self, however, may not be belief. Saying to yourself ‘I can do it’ may help you to actually do the triathlon. But even if it so helps you, that doesn’t mean that before crossing the finish you *believed* you could. It might be that you said to yourself ‘I can do it’, but all the while had the gravest doubts about whether you really could, in which case you didn’t have the *belief* that you could do it. But even then, saying to yourself that you can do it, can actually help you doing it. Isn’t this very curious? Well, it is as curious as the effects that exhortations *of others* can have on our performances. There is power in hearing someone else yelling to you “Come on!”, “You can do it”, “Go on!”. In the case of self-suggestion the only difference is that the exhortation (the yelling) comes from your self. So my point is that saying repeatedly to yourself that p may be effective in the sense that it may help you succeed in something you are trying to accomplish—even if it doesn’t make you begin to believe that p.

Third, saying something to your self may have a genuine effect other than that of an exhortation. When you repeat to yourself “I am not going to be nervous when I have to perform in front of the queen” while you are preparing for just such a performance, that may help you in the sense that you when you envisage that prospective situation, try to imagine what it will be like to be in it, you get a feel for its difficulties and challenges, and thus go there forewarned and hence forearmed. So, your self-addressed and oft-repeated speech has a genuine effect—but it isn’t the effect that you come to *believe* that you are not going to be nervous when you have to perform in front of the queen.

There is yet another genuine effect of repeatedly saying ‘p’ to your self: repeatedly saying to your self “I am going to get that job”, might help you to stay focussed on a goal that

⁸ My source here is Weatherhead 1951: 122-128.

you have set for yourself. It might keep your vision alive. But again, this self-addressed speech need not induce in you the belief that you are going to get that job.

So far I have only suggested that self-addressed speeches *p* may have effects *other* than the formation of the belief that *p*. The crucial issue, however, is not *that* but whether such speeches can also result in the formation of the belief that *p*. To be more precise, the issue is whether we can adopt addressing such speeches to ourselves as a policy in order to secure that we come to believe that *p*? If so, we would have long-range voluntary control over beliefs in a certain kind of propositions about ourselves, and if not, then not. I can only report that it seems to me that we have no such control. I have never seen a convincing case of someone doing that; and all the cases that seemed somewhat plausible, are much more naturally understood as cases where the self-addressed speeches are exhortations, or help to be forewarned, or to stay focussed. And what the previous discussion illustrates is that my denial needn't commit me to deny what I take to be genuine effects of self-addressed speeches.

Let us now move to the other class of beliefs over which we, initially at least, may seem to have long-range voluntary control: beliefs in moral propositions. One special problems arises here, however. According emotivism, and moral antirealism in general, there are no moral truths, no moral propositions. On that position, what it means to say "I believe that I ought to help my neighbour" is *not* that I believe *that it is true* that I ought to help my neighbour. It means something like "Hurray for helping my neighbour"—and thus expresses an emotion or other affective state on the part of the speaker. Since the emotivist's understanding of belief is radically different from what I take belief to be (*viz.* that to believe that *p* *is* to believe that *p* is true), I shall pass it by.

So suppose we think of moral propositions in a realist sense, *i.e.* as having a truth value, then do we have long range voluntary control over belief in moral propositions? Now there is a wide variety of moral propositions. Some are highly general ("We should keep our promises"), others very specific ("I ought to help Lizzy with her math"), and yet others somewhere in-between ("and unjust act has more demerit than an ungenerous one"). And there are many more classes into which moral propositions can be put: some prescribe ("one should speak truthfully"), others forbid ("one should not lie"); some state a merit ("whatever goes on four legs is good"), others a demerit ("whatever goes on two legs is bad"); some specify a value ("Bach's music is invaluable"), others a disvalue ("that is worthless rubbish"); some name a virtue ("Lizzy is dependable"), others a vice ("Jack is a cheat"); etc. Now suppose someone wants to believe that *p*, where *p* is a specific moral proposition—can he then bring it about simply by saying to himself 'p', 'yes, p!' and the like, that eventually he

will wind up believing that p? I must confess to not being able to see that this can be done. When I repeat to myself that whatever goes on four legs is good, I don't eventually wind up believing it at all. As a matter of fact, after having tried, it seemed to me as obviously false as it ever did. But what about moral propositions about which I am in doubt as to their truth value? What about the proposition that it is my duty to help my neighbour's son with math, or the proposition that Lizzy is dependable, or the proposition that what F. Weinreb did during world war II, was good—propositions about which I have my suspicions? Can I eventually bring myself to believing any of these, simply by repeating them to myself? Can I, with respect to those propositions, adopt the policy of self-suggestion as explained, and thereby ensure that I will wind up believing them? That too strains credulity. We simple have no such powers.

The upshot of this section, then, is that there are very simple empirical arguments for the conclusion that with respect to a specific proposition p we have neither the immediate voluntary control, nor the long-range voluntary control that would be required for having the ability to believe at will.

So far I have only dealt with the phenomenon of *acquiring* a particular belief. But couldn't a-, b-, and c- prove feasible when it comes to *maintaining* a belief that one has already got?

2. Empirical-psychology, the unconscious, and the involuntariness of belief

The topic of the previous section currently enjoys great interest from the side of empirical psychologists who work in the field of social psychology. In this section I want to discuss two lines of research in that field that are pertinent to the topic of the involuntariness of belief. The first line of research concerns so-called 'motivated reasoning', the second unconscious belief.

From the outset it should be clear that my goal here cannot be to evaluate the empirical adequacy of the research nor of its conclusions. My main goal is to see whether what I have been arguing so far on the basis of simple empirical arguments finds any support in the research of empirical psychologists.

The first line of research I want to go into concerns, as indicated, so-called 'motivated reasoning'.⁹ This research is especially relevant to the claim in the previous section that 'paying selective attention to the evidence', 'seeking the company of p-believers' and 'self-

⁹ I am drawing here primarily on Kunda 1990, and Kunda 1999.

suggestion' aren't policies one can adopt in order to ensure that one winds up believing p (if it is p one wants to believe). Let me first state what I take to be the fundamental conclusion of the research and next provide some illustrations of how this conclusion is reached.

The fundamental conclusion of the research, in a first approximation, is that motivation affects reasoning. But what is 'motivation' supposed to be here, and what 'reasoning'? To start with the latter: 'reasoning' is not what philosophers typically mean by it; it does not, or not exclusively, refer to the process that consists in the thought progression that moves from premis(es) to conclusion(s). 'Reasoning' is an umbrella term for "the cognitive processes we engage in to arrive at a judgment".¹⁰ The literature I am referring to makes no systematic distinction between 'belief' and 'judgment', and the terms are used more or less as synonyms. So in order to obtain at least a verbal connection with the topic of the previous section, we might say that 'reasoning' refers to any process that results in the formation of belief. Such processes include: "forming impressions, determining one's beliefs and attitudes, evaluating evidence, and making decisions"¹¹, as well as modes of dealing with information. As to 'motivation': subjects in whom processes of belief formation are at work (subjects that engage in reasoning) will as a rule be driven by motivations—and the processes operative in the formation of belief (their reasonings) will be driven by those motivations. Or, as Ziva Kunda has it, their reasoning will be motivated. Now there are motives and motivations. Subjects can be driven by the motivation to arrive at an accurate conclusion, but they can also be driven by a motivation to arrive at a particular, 'directional', conclusion. Given these clarifications, the fundamental conclusion of the research I am referring to can be stated as follows: both the motivation to be accurate and the motivation to arrive at a particular conclusion affect the processes operative in belief formation (both motivations affect reasoning).

For present purposes only one half of this conclusion is relevant. The present purpose is, after all, to connect the research on motivated reasoning with the claim made in the previous section that 'paying selective attention to the evidence' and the like aren't policies one can adopt in order to ensure that one winds up believing p (if it is p that one wants to believe). And to that purpose it is only relevant that the motivation to arrive at a particular conclusion can affect belief formation.

How is this fundamental conclusion reached? By empirically finding out that cases like the following are not at all exceptional. The case is of a woman who throughout her three

¹⁰ Kunda 1999: 211.

¹¹ Kunda 1990: 480.

pregnancies continued smoking. It was the days when it wasn't common knowledge that that might not be a good idea. When her children are grown up, she reads a newspaper article on the effects of smoking during pregnancy on babies. She exclaims to her daughter, who had read the article too: "This article is full of nonsense. It says that if you smoke during pregnancy you will have small babies. Well, my babies were huge!" And they were huge indeed. The daughter then points out to her mother that the article also said that smoking during pregnancy could cause lung problems in babies, and reminds her of the fact that both of her huge brothers had suffered from spastic bronchitis throughout their childhood. To which her mother's dismissive response was "that is only a sample of two".¹²

This, it is said, is an example of motivated reasoning. The reasoning here is motivated by the wish to arrive at a particular conclusion—namely that the article is full of nonsense, and need not be taken seriously. This, it is said, is a case where someone comes to believe what she wants to believe *because* she wants to believe it. And there are many more such cases on a hum drum level: lovers may be blind to any blemish in their loved ones, parents may not see the obvious faults in their children, etc. A more dramatic case is discussed by Randy Shilts in his book *And the Band Played On*¹³ in which he describes the reactions of various communities to the ever more disturbing findings about the, at the time, unknown disease among gay men, now known as AIDS. The gay community, the scientific community, the U.S. government, and the blood banks, all initially refused to believe evidence for the devastating conclusion that the disease was infectious and could be transmitted sexually. Each group had its own motivation for believing that this horrific scenario was false. For the gay community a characteristic lifestyle was at stake, the blood banks would have to acknowledge that their blood was unsafe and to modify their procedures at great cost, the government would have to face vast financial and social implications, whereas the scientific community would have to openly declare that many lives, both inside and outside the gay community, were in jeopardy. And all of these groups were to some extent successful. Scientists followed every lead that pointed to alternative explanations, and within the gay community an alternative theory that postulated a right-wing conspiracy to undermine the gay lifestyle became popular.

If these cases are illustrative for directional motivated reasoning, then it is clear that motivated reasoning is a bias that pollutes the beliefs it engenders.

¹² The daughter is Ziva Kunda, who relates the story in Kunda 1999: 212.

¹³ Shilts 1987. This case is also described by Kunda 1999: 223-4. Another case would be the various beliefs people had about what happened to Hitler in May 1945; see Trevor Rope @@

It is important to note that the empirical research also suggests that the power of motivated reasoning is limited. We are not at liberty to believe whatever we want to believe simply because we want to. Research indicates that the desired conclusions are only drawn when subjects can come up with enough evidence to support it, when they can construct a justification for the desired conclusion. In the case of the heavy smoking woman, the justification (evidence) for resisting the belief that her sons' lung problems were caused by her smoking habits, was that two children whose mother smoked during pregnancies suffer from lung problems, doesn't prove that the smoking *caused* the problems. In the AIDS case, the conclusion that the disease is infectious and sexually transmittable, could be resisted so long as there seemed to be reasonable alternatives. But as the evidence mounted, these alternatives no longer seemed reasonable and virtually everybody came to believe the fearsome conclusion.

An important question now arises. In the cases mentioned (and the literature suggests that such cases, in all sorts of variations, abound) bias is at work. But why do psychologists think that the best account of these biases is in terms of 'motivation'—couldn't there be an account of the phenomena (resisting a certain conclusion) purely in cognitive terms—in which case it would be unwarranted to speak of a bias? To return to the AIDS case: couldn't it be said that, prior the heaping up of evidence pointing to the fearsome conclusion, the evidential situation was simply unclear, and that various alternative hypotheses were live options, and that there was simply no compelling reason for thinking that the disease was infectious and sexually transmittable? This account of the case is not in terms of 'motivation' but purely in 'cognitive' terms such as: evidence, hypothesis and reason. And doesn't this account work as well as does the account in terms of motivation? If so, it would seem to be preferable because it does not invoke the extra notion of 'motivation'—and we should always favour the more parsimonious account.

Psychologists claim, however, that there is independent evidence for the thesis that motivation affects belief.¹⁴ Various studies indicate that humans are strongly motivated to believe that other persons, about whom they know next to nothing but on whom they are going to be dependent, are likable—much more likable than when they were *not* going to be dependent on them. Someone who is about to go on a blind date with Mr X, will be strongly motivated to believe that Mr X is likable. And if you know that the success of your team depends on the competence of a new but unknown team mate, you may be strongly motivated

¹⁴ Kunda 1990; Kunda 1999: 215-223.

to believe that this person is competent. That is the idea. The evidence for it is of the following sort. In one study¹⁵ participants take part in a dating experiment. Before meeting their dating partners, the participants watched a videotaped discussion among three people, one of whom was said to be their future partner. There were three groups of participants, each expecting to date a different one of the three discussants, and an additional control group that was not expecting to date any of them. The participants rated their expected date as more likable than the two other discussants, and as having more positive personality.

As Kunda says, results such as these are hard to account for in cognitive terms. After all, the participants knew nothing about their future date, and hence couldn't have formed any specific expectation about him or her. Kunda's conclusion with respect to this study and similar ones is: "The fact that participants nevertheless came to view the person whom they wanted to like and respect as especially likable and competent in these studies therefore provides strong evidence for the role of motivation in judgment".¹⁶

Other studies provide other and additional evidence for motivated reasoning.¹⁷ It has been found that if a certain person finds an event desirable, he will believe that the event is much more likely to occur than if he had found the event undesirable. It has also been found that people tend to take credit for their success but to deny responsibility for failure.

Suppose now that we accept Kunda's conclusion¹⁸, so suppose there is the phenomenon of motivational reasoning, i.e. the phenomenon that the motivation to arrive at a particular conclusion affects the processes operative in belief formation. How does that relate to the thesis that 'paying selective attention to the evidence' and the like are not policies one can adopt in order to secure that one winds up believing the specific proposition p?

The first thing to say is that the phenomenon of motivated reasoning does not show that belief is voluntary after all. It doesn't show that we can decide, by *fiat*, to believe just any proposition p we like. It *does* show that belief formation, as it actually works with us, can be very different from the clean-cut, clear-headed processes that involve evidence weighing and the like. But saying that is not the same as saying that beliefs can be formed simply by an exercise of the will. Second, the phenomenon of motivated reasoning doesn't undermine my claim that that 'paying selective attention to the evidence' cannot be adopted as a policy to ensure a wanted belief. The point to see is that motivated reasoning, just as 'paying selective

¹⁵ Berscheid et al. 1976; this study is discussed in Kunda 1990: 486, and Kunda 1999: 215-6.

¹⁶ Kunda 1999: 216.

¹⁷ All of these are discussed in Kunda 1990.

¹⁸ An important part of Kunda's work is to delineate the mechanisms that underlie motivated reasoning. One such mechanism is 'biased memory search'.

attention to the evidence’, is a highly unconscious phenomenon—at least when it comes to ourselves. (We may observe that others engage in motivated reasoning, and we may note that others pay selective attention to evidence, but we usually don’t observe that in ourselves.) But exactly this feature prevents either of them from being candidates for policies of belief formation that can be adopted at will.

Thus far I have broached empirical psychological studies that might be taken to show that what I have argued in the previous section, viz. that belief is involuntary, is false (but in fact don’t). I now turn to psychological studies that might seem to *support* the thesis that belief is involuntary.

The possible support comes from research in which the crucial concept is the concept of ‘the unconscious’. That research indicates that we are unconscious of much of what we see and feel, but also, more pertinent to the topic of this chapter, unaware of at least many of our beliefs and of the processes that propel it. My argument that this research supports the thesis that belief is involuntary, is this:

1. Belief as well as belief formation is often unconscious.
2. If one is unconscious of X, X cannot be under one’s voluntary control.
3. Therefore, unconscious belief and unconscious belief formation are not under one’s voluntary control.

It is clear that this argument, if successful, at best *partially* supports the claim that I have made in the previous section. After all, my claim was that belief is involuntary (by which I meant that we neither have immediate nor long-range voluntary control over a specific belief), whereas the conclusion of the present argument is only that *unconscious* belief formation is involuntary, and thus doesn’t rule out the possibility that conscious belief formation is voluntary. Furthermore, so far I have been assuming that believing is a conscious affair: if one believes that p, then one is conscious of the fact that one believes p (well, this is not exactly right—given what I have said in the previous chapter). Still, the argument, if successful, supports the conclusion of the previous section. For it might very well be that processes that lead up to belief are themselves oftentimes unconscious. My discussion of this argument consists of a discussion of the premises, from which the conclusion self-evidently follows.

The first premise is a summary statement of findings of empirical psychology. It is, furthermore, an empirical, as opposed to a conceptual, thesis. Now what does ‘unconscious’

in this premise mean? What is it for something to be ‘unconscious’ in the sense intended? Wilson says: “A simple definition of the unconscious is anything that is in your mind that you are not consciously aware of at a particular point in time”.¹⁹ He avows, however, that this definition is too simple because it ‘stretches things’. It would imply, for instance, that someone who isn’t currently thinking about his hometown is unconscious of “Philadelphia”, if that is his hometown—which doesn’t seem right. The problem with the simple definition (which has this unwelcome implication), he says, is that it equates consciousness with attention or short-term memory, while there are things one isn’t currently paying attention to, or aren’t currently in one’s short-term memory that still aren’t ‘unconscious’ in the preferred sense. He therefore says “A better working definition of the unconscious is *mental processes that are inaccessible to consciousness but that influence judgments, feelings, or behaviour*”²⁰, but acknowledges that the unconscious is notoriously hard to define and that this is only one definition out of many.²¹ He even says that he doesn’t want to get bogged down by definitional issues and that it is more interesting to take a look at what humans can accomplish outside the spotlight of consciousness. Still, it merits taking a closer look at the definition and bringing out certain problems. (i) The definition restricts the unconscious to ‘mental processes’, but, as the literature on the unconscious makes abundantly clear, not everything that ‘is part of’ one’s unconscious is a *process*; there are, it is claimed, unconscious feelings, and, more relevant for my topic, unconscious beliefs--but neither feelings nor beliefs are processes. (ii) The definition says that what is unconscious is ‘inaccessible’ to consciousness. But this can’t be correct, for the research of the unconscious is intent on studying it and hence on bringing to consciousness how the unconscious works and what it does (and even how one can get to know one’s unconscious). Perhaps the definition needs to be taken as affirming that subjects can have no access to their *own* unconscious—although they can have it to the unconscious of others (especially when they are empirical psychologists). But this cannot be right either, for in the same book in which he offers the definition, Wilson aims to provide clues as to how one can get to know at least parts of one’s own unconscious.²²

So, if the working definition doesn’t work, there is no alternative to giving examples of the unconscious at work. The argument is limited to unconscious belief and the formation

¹⁹ Wilson 2002: 22. Wilson 2002: ch. 1 makes it clear that the modern notion of the unconscious’ is very different from Freud’s in that it has nothing to do with repressed infantile thoughts that are kept out of consciousness because they are a source of psychic pain.

²⁰ Wilson 2002: 23.

²¹ Miller 1942 offers sixteen different definitions of the unconscious!

²² See especially the later chapters of Wilson 2002.

thereof²³ and so I will give examples of that. One example is due to the French physician Edouard Claparede and concerns a woman suffering from amnesia. Every time she sees Claparede, she wouldn't have recollections of previous consults and accordingly he would have to introduce himself anew each visit. One day Claparede decided to shake hands with the patient (as he always did) but this time while holding a hidden pin in his hand. The woman felt, of course, a painful prick and withdrew her hand immediately. What is interesting is what happened during the next consult. Again the patient showed no sign of recognizing Claparede, but when he reached out his hand (as he always did) the patient refused to shake it!²⁴ One way to interpret this behaviour is to say that although the patient did not consciously recognize Claparede, and accordingly didn't have any conscious beliefs about him, she unconsciously believed that shaking hands with this man (who is in fact Claparede) might give her a painful sensation.

What I have described in terms of an unconscious belief that the patient acquired, can also be (and is in fact by psychologists) described in terms of unconscious learning. These two descriptions, however, aren't at odd with each other in that believing and learning seem to be closely connected with each other. When someone learns that p, he thereby comes to believe that p—whether the learning is conscious, or unconscious.

Another example of implicit learning is described by Pawel Lewicki c.s.²⁵ In a laboratory situation participants were asked to watch a computer screen that was divided in four quadrants. On each trial, the letter X appeared in a quadrant, and the participants' task was to press one of four buttons to indicate in which of the quadrants the letter X appeared. Unbeknownst to the participants, the presentation of the X's followed a complex rule. As time went by, their performance improved, they became faster and faster at pressing the correct button, and hence they appeared to learn the rule. But they were unable to verbalize, or even notice, it. After a certain elapse of time (and still unbeknownst to the participants), the researchers changed the rule with the result that the participants' performance deteriorated dramatically. Although the participants noticed this, none of them knew why. And this shows, or so argue Lewicki c.s., that the participants had learned the original rule *unconsciously*, and hence also knew (and believed) it unconsciously.

²³ The unconscious is also said to be at work in, among others: proprioception; the way that light rays strike the retina and transform these into three-dimensional vision; filtering of information; interpretation of behaviour; evaluation; goal-setting.

²⁴ This case is described by Wilson 2002: 25.

²⁵ Lewicki, Hill & Bizot 1988; this is also described by Wilson 2002: 26-7.

A further, and widely recognized, example of unconscious learning and belief is language learning, which, in young children, goes almost without effort and without conscious awareness of what has been learned. Small children learn the syntax, grammar and semantics of a language without consciously learning syntactical, grammatical and semantical *rules*. They are typically unable to even explicate them. Still, they would be able to identify syntactical, grammatical and semantical mistakes—which suggest that they know (and hence believe) the relevant rules, albeit it unconsciously.

So far, the examples were of unconscious belief—beliefs of which people are supposed to be unconscious. In addition to this, research also indicates that many processes that trigger conscious belief escape the light beam of consciousness. One famous experiment that suggests this is due to Robert Nisbett and Timothy Wilson. Passengers in a shopping mall were asked to examine the quality of four pairs of nylon hoses that were neatly arranged on a table and say which pair was of the best quality. As a matter of fact the pairs were of equal quality. But 40% of the respondents said that the right-most pair was the best, and 31% the pair second from right. Upon being asked why they believed that the pair they selected was best, the respondents would always point to an attribute of the preferred pair, such as its superior knit, sheerness, or elasticity. What is at work here, Nisbett and Wilson argue, is a so-called ‘position effect’—in the case the effect that people have a marked preference for items on the right side of a display. None of the respondents, of course, was conscious of this effect, and they rejected the very suggestion that their belief might be due to such an effect. Nonetheless, the effect is real, or so these studies argue.

Other research indicates that belief formation is often automatic in the sense that we are not aware of the cognitive processes that underlie them. Numerous studies indicate the unconscious effects of priming. A summary statement of such research by a renowned psychologist reads: “Much of our mental life goes on automatically. Unlike controlled processes, automatic processes occur outside awareness, are carried out without intention, are uncontrollable (i.e. one cannot stop them), or are highly efficient (i.e. require few cognitive resources). Our judgments [beliefs, RvW] can be influenced by factors that we have never been aware of and have only been exposed to subliminally ... and by factors that we can still recall but whose influence we are unaware of”.²⁶

Let me now move to the second premise that states that if one is unconscious of X, X cannot be under one’s voluntary control. I believe this premise is nearly self-evident. If I am

²⁶ Kunda 1999: 308.

unaware that a certain process is taking place, I cannot control its execution. Control over a process requires awareness of the process. Such awareness is of course insufficient for control, for one might be aware of a certain process and still have no control over it, as in the case of the process of ebb turning into flow. Still, control over X requires that one is conscious of X. Applied to belief this means that if one is to have control over the formation of belief, it is required that one is aware of the processes that propel it.²⁷

If the second premise is true, this implies that there is no point in talking about ‘unconscious control’. Surely, we can be unconscious of certain effects of our actions. But to say that X is an unconscious effect of what someone does, isn’t at all the same as saying that X has ‘control’ (albeit it unconsciously) over X. For control over X requires having a conscious intention vis-à-vis X. But that is not required for having an unconscious effect.

And thus the conclusion of the argument stands: unconscious belief and unconscious belief formation are not under one’s voluntary control. And this, as I have explained earlier on, in turn supports the claim that belief is involuntary—the claim that I set out to defend in the previous section.

²⁷ This intuition is shared by Kunda; see Kunda 1999: 268.