University of Alberta

The Prescriptivity of Conscious Belief

by

Andrei Buleandra

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Philosophy

Fall 2011
Edmonton, Alberta

Abstract

In my dissertation I explain and defend the claim that conscious beliefs are essentially prescriptive. I argue that norms of conscious belief are explained by the fact that consciously believing *p* involves a commitment to the truth of *p*, a commitment analogous to the one involved in the act of accepting an assertion in public linguistic practice. Having a conscious belief implies being vulnerable to certain questions and criticisms from other agents. For instance, when asked for reasons for her belief, a person should provide a justification which demonstrates her entitlement to accepting the given proposition as true. Moreover, if a certain belief logically follows from the agent's beliefs then she should either accept it as a conclusion or revise her initial beliefs. I argue that both deliberative and non-inferential conscious beliefs can be construed as acceptances of assertions and that they carry the same normative import as public acts of accepting claims put forward by others. The intrinsic relation between conscious belief and language-use shows that conscious belief is irreducible to unconscious or lower-level belief, the type of belief which we attribute to non-human animals or small children. Rather than trying to reduce conscious belief to lower-level belief, I suggest that we should offer an account of the emergence of the linguistic practice of assertion in terms of animal belief and then explain the normative features of conscious belief by reference to the norms implicit in assertional practice. In addition, my work proposes a way of formulating the norms of conscious belief which is consistent with the fact that actual human beings do not have perfect logical

abilities; that they can only dedicate a limited amount of time and cognitive resources to the task of reasoning.

Table of Contents

Introduction

One of the aims of philosophers working in the areas of philosophy of mind and epistemology is to illuminate the nature of cognitive states. When trying to explain the nature of the intentional state of belief theoreticians are confronted with the issue whether beliefs are essentially normative. That is, whether having beliefs essentially involves being subject to certain norms. The norms taken to characterize beliefs are norms of evidence, coherence and consistency. In this dissertation I will defend the thesis that beliefs are intrinsically normative but, before sketching my arguments, the nature of the problem and the concepts involved should be specified.

We must begin by acknowledging the fact that there are many types of belief. One important distinction is between conscious and unconscious beliefs. A preliminary way[1] of drawing the distinction goes as follows. The notion of conscious belief does not only refer to the beliefs which are occurrent, or presently given in introspection. Rather, it is meant to capture the types of beliefs which can be made available to consciousness. For instance, my belief that Plato was Socrates' student is conscious although I do not constantly entertain it. However, if I happen to think about Greek philosophy, I can easily bring it to mind. Unconscious beliefs are beliefs which cannot be brought to consciousness. They can causally influence the agent's behavior and we can become aware of the fact that we have them by reflecting on our behavior. One way of explaining the notion of unconscious belief is by reference to the behavior of animals and small children. Their behaviors are shaped by certain beliefs about their environment but they are unable to consciously articulate and reflect on those beliefs.

Another prominent distinction is between deliberative belief and non-inferential belief. Deliberative beliefs are the outcome of deliberation or reflection. Normally, in the context of doxastic deliberation, we weigh the evidence for and against a specific proposition and, based on reasons which we find convincing, we make up our mind whether that proposition is true or not,

[1] This distinction will be discussed in more detail in Chapter 2, section 6.

whether we should believe it or not. Perceptual beliefs are the main examples of non-inferential beliefs. Our perceptual beliefs are normally the result of specific sensory experiences and they are formed automatically on their basis. They do not involve inference or reflection.

Philosophers also distinguish between full belief and partial belief. Whereas full belief is an all-or-nothing affair, partial belief admits of degrees. Accordingly, the former is expressed by the use of unqualified assertions while the latter is usually articulated as a qualified claim like "It's probable that *p*" or "I'm pretty sure that *p*". Another important difference is that only full belief can also represent knowledge. Fully believing *p* implies taking *p* to be true and taking your evidence for *p* to be conclusive. Although this does not guarantee that the subject has knowledge that *p*, it shows that full belief is a state which 'aims at knowledge' (Owens 2000; Engel 2005). In a similar vein, Jonathan Adler (2002) argues that, *from the first-person perspective*, full-belief shares one essential quality of knowledge: facticity. Knowledge is factive in the sense that if someone knows that *p* then *p* is a fact. Full belief, according to Adler, is factive in the sense that if one fully believes that *p* then one takes *p* to be a fact[2]. By contrast, partial belief explicitly signals that the subject does not have knowledge of the subject-matter; the subject does not take *p* to be true but only probable. Also, being confident that *p* implies acknowledging that your evidence in support of *p* is not conclusive or sufficient for knowledge[3].

The question whether beliefs are intrinsically normative usually concerns conscious, deliberative, full beliefs. Unless otherwise specified, I will use the term 'belief' to refer to this type of belief. Although conscious, deliberative full beliefs will be the focus of this work, I will try to generalize the conclusion regarding their normative character to the cases of conscious non-inferential belief and

[2] Adler's claim that, from the first-person perspective, full belief is factive is connected with the idea that full belief is transparent. He characterizes belief transparecy as follows: if one believes that *p* "one sees through one's attitude to the world without seeing that attitude" (Adler 2002, p. 11). In Chapter 1, I will take a closer look at this feature of full belief.

[3] The distinction between full belief and partial belief will be considered in more detail in Chapter 7.

conscious partial belief. However, I will argue that unconscious beliefs are not normative. This is because they are independent of language and, in my view, the key to understanding the normativity of conscious belief lies in its relation to language use.

Next, we should distinguish two senses in which cognitive states like belief are said to be normative (Hattiangadi 2007). First, such states are subject to evaluations in terms of correctness. We say, for instance, that a true belief is correct and a false belief is incorrect. Second, beliefs are said to imply prescriptions to the effect that the epistemic agents ought to perform certain actions (e.g. revise beliefs which are false). It is a matter of debate whether correctness is an intrinsically prescriptive notion (Wedgwood 2007; Glüer 1999; Rosen 2001; Boghossian 2003). Some philosophers plausibly argue that correctness is not prescriptive in itself, but only in conjunction with a subject's desire to do what is correct (Glüer 1999; Rosen 2001; Dretske 2000). Thus, without assuming that correctness is a prescriptive notion, in this work I will try to answer the question: what are the prescriptions involved in having conscious beliefs?

Before tackling the question whether beliefs are constitutively normative, we should examine what features of belief are taken to play a constitutive role. Many philosophers argue that what differentiates beliefs from other intentional states is that it 'aims at truth' (Velleman 2000; Shah 2003; Wedgwood 2002; Engel 2003). A belief is a cognitive state which aims at representing reality truthfully. Put differently, belief bears a constitutive relation to truth.

Furthermore, rational or inferential relations with other beliefs are also specific of this intentional state (Davidson 2004; Cherniak 1986; Wedgwood 2007; Zangwill 2005). Belief is taken to be subject to norms of logical implication and consistency. The fact that belief is subject to these constrains is a consequence of its intrinsic connection to truth. This point can be illustrated by reference to a preliminary description of these norms[4]. An agent should not

[4] The inferential norms for belief are discussed in Chapter 6.

believe *p* and not-*p* because two inconsistent beliefs cannot both be true. Similarly, if the logical implication between *p* and *q* is truth-preserving then if an agent believes *p*, she should also believe *q*.

In this work I will develop and defend the thesis that truth is prescriptive for belief. In other words, I will argue that the relation between belief and truth is both constitutive and prescriptive. One way of making prominent the prescriptive nature of belief is by reflecting on the fact that belief is *not just* taking a proposition as true. It is taking a proposition to be true in a way which makes one vulnerable to questions regarding his reasons and to criticism. For instance, let's suppose that, as a result of deliberation, John forms the belief that global warming is a myth. His reasoning goes like this: if global warming were a fact then the winters in my town would be less cold but this winter is the coldest on record. However, suppose John expresses his belief in an argument and Lisa, an activist in the green movement, is quick to point out that extreme temperatures are an effect of global warming. She explains that global warming not only means the rising of the worldwide average temperature, but also the increased variation of temperatures. Given this criticism of his reasoning John is expected to revise his belief on the issue.

This example shows that we are in a sense active with respect to our beliefs, they are the result of our making up our minds about some issue or settling a question. Our beliefs are not just things that happen to us, states which we contemplate passively. In this sense they are different from mental states like being in pain or experiencing certain emotions. The fact that what we believe is in this sense up to us explains the element of responsibility or answerability involved in forming a belief. The agent ought to articulate his reasons for adopting a belief and respond to criticism directed at his reasoning. Thus, the prescriptivity of truth for belief is explained by the fact that we are active with respect to our beliefs. This prescriptivity also plays a constitutive role because the element of responsibility or answerability is a feature of belief which distinguishes it from

other mental states (e.g. having certain sensations or feelings). This feature of belief should figure in an explanation of the meaning of the term.

The idea that when one believes *p* one is vulnerable to questions and criticisms is captured by the notion of commitment and the idea that belief implies a *commitment to truth* (Moran 2002; Hieronymi 2006; Brandom 1994). The claim that belief involves such a commitment fits the way the concept of belief is used in ordinary discourse. We normally tend to provide reasons for our beliefs if challenged. In light of this fact, the claim that belief implies a commitment to truth is taken to suggest that the believer has an obligation to demonstrate his entitlement to that commitment. In addition, our practice of expressing our beliefs and supporting them by arguments is shaped by two other prominent inferential norms. When a proposition logically follows from the agent's set of beliefs then he ought to either accept it or revise at least one of his initial beliefs. Furthermore, the agent is not entitled to undertake a commitment to propositions which are incompatible with the ones which he holds-true[5].

Besides formulating the obligations specific to having beliefs, this thesis aims at explaining *why* conscious belief is normative. In other words, how can we explain that belief bears these intrinsic relations with activities like asking for reasons or evidence, articulating criticism, drawing an inference and so on? I think that we can account for these features of conscious belief only by examining its relation to the linguistic practice of assertion. Although both assertions and intentional actions need the support of reasons, only in the former case should those reasons have a relation to truth. Actions are explained by reference to what the agent finds desirable. Moreover, the asserter is criticized when he is not able to prove entitlement to the claim he put forward as true, and he is subject to norms of consistency and logical consequence.

Several philosophers have observed and articulated the parallel between belief and assertion (Adler 2002; Brandom 1994; Williams 2002). However, they failed to recognize that the analogy faces an important difficulty. Whereas we use

[5] The content of these norms will be qualified in Chapter 6, in light of the fact that normal human beings do not have ideal logical abilities.

assertions to transmit information to others, it does not seem that when we form a belief or judgment we aim at informing others. Assertion has an explicit public character whereas belief seems to be a private affair. In response to this issue, I will develop the view that belief should be construed on the model of the public act of accepting an assertion, not of producing one. When we deliberate whether to believe *p* we consider whether to accept the claim that *p*. Making up our mind and accepting *p* means that we settle on a positive answer to the question whether *p* (Hieronymi 2008, 2009). However, given that inner deliberation is the internalization of abilities exercised in a public normative practice, I will maintain that the private act of accepting a claim is subject to the same norms as its public counterpart.

The thesis is organized as follows. In chapter 1 I present and evaluate a recent view according to which belief is governed by a norm of truth of the form: one *ought* to believe *p* only if *p* is true (Boghossian 2003; Shah 2003; Velleman and Shah 2005). The proponents of this view argue that truth is both prescriptive and constitutive of belief. They emphasize the fact that the constitutivity of truth for belief is expressed in the phenomenon of doxastic transparency: the fact that the deliberative question *whether to believe p* immediately gives rise to the question *whether p is true*. Or, in other words, that from the first-person perspective, belief is one's take on how things are; to believe *p* is to take *p* to be a fact. . The defenders of the norm of truth view argue that we can explain the phenomenon of transparency by reference to the norm of truth. This view faces three important challenges:

(1) The account explains transparency by reference to the norm of truth only at the cost of losing the genuine prescriptivity of the truth norm. The truth norm is compatible with transparency only if it cannot be broken. On the other hand, the possibility of the truth norm being violated is in tension with the phenomenon of transparency.

(2) The norm-of-truth account is circular. The norms of truth and evidence are construed as explicit norms which guide our doxastic deliberation. However, understanding the notion of following an explicit norm requires grasping the concept of belief.

(3) The view does not explain how we exercise control with respect to our beliefs. The defenders of the account argue that when forming beliefs we take guidance from doxastic norms but they reject doxastic voluntarism, the view that we can form beliefs at will. But then, how can we follow doxastic norms and decide what to believe on their basis?

Chapter 2 develops an alternative account of the prescriptivity of truth for belief and shows that it can answer objections (1) and (2). The basic insight behind the proposed view is that belief involves a commitment to truth and that the notion of commitment is a prescriptive notion. This commitment is modeled on the obligations one incurs when one publicly accepts a claim. In such a context, one can be challenged to argue why one gave one's assent to a particular statement. In response to this challenge, one can either produce reasons which are considered appropriate or fail to do so. In the latter case one does not succeed in demonstrating one's entitlement to accepting the claim. These ideas point to an answer to the first objection. On this picture, doxastic transparency is compatible with the prescriptivity of truth for belief because to believe *p* means to be committed to the truth of *p*. That is, from the first-person perspective, belief is transparent – it is just equivalent to holding-true. However, the notion of commitment is also prescriptive. The subject has to prove his entitlement to that commitment. Moreover, the agent may or may not be able to demonstrate his entitlement. The reasons he provides for his belief are assessed by reference to epistemic standards which are in force in his respective community. Since *it is possible* that his reasons do not meet those standards, it follows that the notion of doxastic commitment has genuine prescriptive force.

Answering the second objection involves offering an account of the notion of commitment to truth and of the practice of assertion which does not rely on the notion of conscious belief. This is achieved by linking two important ideas. First, Robert Brandom (1994) defines the notion of commitment in terms of the attribution of commitment which, in turn, is defined as a disposition to sanction failure to discharge the obligation. Second, Ruth Millikan (2005) defines assertion as a form of behavior which got selected because it serves a biological proper function: it helps cooperation and coordination between the members of a group. These two insights pave the way for a naturalistic account of the notion of assertional commitment. It is plausible to argue that, in order to ensure that the members of the community do not deceive one another when transmitting information, a system of sanctions is embedded in the practice of assertion. This account of assertional practice has the virtue of not making reference to the explicit beliefs of the participants. Non-human animals, for instance, can engage in cooperative forms of behavior without having the ability to form conscious beliefs. Thus, this view is not vulnerable to the circularity challenge.

In Chapter 3 I qualify the claim that we judge whether an agent has entitlement to holding a claim true by reference to the standards in force in our community. Clearly, there are different epistemic communities with different standards in place. Thus, I am faced with the option of either embracing epistemic relativism or arguing that there are cross-cultural norms of rationality. Following the argument regarding the relation between belief and the practice of assertion developed in Chapter 2, I maintain that there are ways of proving entitlement to a claim which are used in different cultures. For instance, members of various cultures follow epistemic principles like observation, induction and deduction. Put in conjunction with the claim that entitlement has a default and challenge structure, the observation that there are norms which are considered valid across cultures leads to the idea that members of one tradition can formulate *legitimate challenges* to specific ways of discharging entitlement used in other cultures. Thus, the logical space is created for an objective assessment of modes of

argumentation used in different epistemic communities. In Chapter 3 I apply these general ideas to specific examples from the history of science and from anthropology.

Chapter 4 shows how the assertion account of belief outlined in Chapter 2 can answer objection (3). An agent acquires the status of being an asserter or believer by displaying an ability to participate in the practice of giving and asking for reasons. That is, by proving an ability to assess the force of epistemic reasons and evidence which fits the standards in force in his epistemic community. I will submit that having doxastic control over one's belief formation just means being responsive to epistemic reasons. More specifically, although a subject may fail to demonstrate entitlement to his doxastic commitments, he is still taken to be in control of his beliefs if he reacts to criticism by revising them. This shows that he is responsive to reasons. By contrast, if someone holds a belief which is the result of a neurotic disorder like phobia or paranoia then they will not be able to react to criticism in accord with the prevalent practice. Belief responsibility is made possible by this notion of doxastic control. Moreover, exercising this type of control and responsibility does not require having explicit thoughts about the norms which govern the practice of assertion.

In Chapter 5 I present and evaluate the teleological account of belief. The advocates of this account (Velleman 2000; Vahid 2006; Steglich-Petersen 2006) argue that we can make sense of the metaphor that belief 'aims at truth' in non-normative terms. They develop and refine the claim that to believe *p* is to accept *p* with the aim or intention of accepting a truth. In reply, I will articulate two critical points. First, it is not clear that on this account truth does not have any normative import. For instance, given that evidence is a guide to truth, it follows that aiming at truth necessarily involves following evidential norms. One cannot aim the acceptance of *p* at truth unless one takes into account the evidence regarding *p*. Second, it appears that the teleological account faces the objection that it is circular, since it uses the idea of having an explicit intention in the definition of belief. But it is hard to make sense of the concept of intention without that of

belief. Forming an intention is based on having a desire and beliefs about how to realize that desire (Davidson 2001a; Bratman 1987). I will argue that one plausible way in which the proponents of the teleological account can avoid the charge of circularity is by arguing that 'aiming at truth' is not a property of the individual but of the communication practice he participates in. Thus, the intuition behind the teleological account can be incorporated into the view of assertion and belief developed in Chapter 2. Believing *p* is accepting *p* with the aim of accepting a truth, but the acceptance of *p* has this aim because *p* is an assertion and the practice of assertion is shaped by the biological purpose of sharing information.

Chapter 6 is dedicated to articulating the inferential norms which govern belief. One initial suggestion is that if a subject believes *p* then he ought to accept the logical consequences of *p* and reject claims which are incompatible with the truth of *p*. This view is motivated by the fact that belief involves a commitment to truth and logical inferences are truth-preserving. An intuitive objection to this proposal is that human beings do not posses perfect deductive abilities. Since we can invest a limited amount of time and cognitive resources into the task of reasoning, we find ourselves in a *finitary predicament* (Cherniak 1986). Then, we cannot compute *all* the logical consequences of our beliefs and we cannot check our entire belief-set for inconsistencies every time we accept new information. In light of this objection, I will propose that inferential obligations apply only when the agent *becomes aware* of a logical consequence of his beliefs or of an inconsistency between claims he holds true[6]. However, I will maintain, this does not imply that the subject can be massively incoherent or inconsistent as long as he is not aware of his logical mistakes. Following Cherniak, I will argue that having a *minimum deductive ability* is constitutive of being a cognitive agent. This observation fits with the account of assertion outlined in Chapter 2. If someone does not have minimum deductive abilities he cannot successfully

[6] It might be objected that these inferential norms are too weak. It is possible that the agent avoids becoming aware of an inconsistency or he outright denies a conflict. These are epistemically irresponsible actions and an account of the inferential norms of belief should explain this fact. This objection will be considered in Chapter 6, section 5.

participate in the game of giving and asking for reasons and, therefore, cannot acquire the status of being an asserter or believer.

Chapter 7 concerns the relation between full and partial belief. As already mentioned this thesis focuses on the phenomenon of full belief rather than graded belief. However, this approach presupposes that full belief is a distinct phenomenon and can be understood independently of partial belief. This claim has been challenged by several philosophers (Jeffrey 1992; Frankish 2004; Foley 2003). The challenge takes one of two forms. Either it is argued that flat-out belief is reducible to graded belief or that we cannot understand full-belief without first grasping the concept of degrees of belief. The latter claim is based in the idea that flat-out belief is the result of adopting an epistemic policy based on the agent's degrees of belief, his desires, and epistemic goals. On this picture, full belief is just the outcome of a pragmatic decision made on the basis of partial beliefs and desires (Frankish 2004; Foley 2003). In Chapter 7 I reject both of these suggestions. Since full belief is conceptually connected to knowledge by sharing the property of facticity or transparency, it follows that it cannot be reducible to graded belief. Partially believing that *p* immediately implies not knowing *p*. Moreover, the fact that full belief is transparent, that it involves holding a proposition true, entails that one cannot, in full consciousness, form a belief on the basis of pragmatic factors (Shah 2003; Owens 2000). But this possibility animates the idea that full belief is adopting a policy of accepting *p* as true, based on one's confidence in *p*. Therefore, it follows that the notion of full belief characterized by Frankish and Foley is not the ordinary notion but a different epistemic concept. Their analyses refers to the act of accepting a proposition as true when one is pressed by practical factors like time and limited cognitive resources. In this context, the agent takes an epistemic risk by accepting *p* but this acceptance lacks the central properties of full belief: it is not transparent and cannot amount to knowledge.

Chapter I - Belief and the Norm of Truth

1. Belief and the 'aim of truth'

One standard way of defining belief is in terms of its relation to desire and action (Dennett 1978; Stalnaker 1984). On this picture, if we know what a cognitive system desires and we observe its action we can infer what the system holds true. Alternatively, if we know what the system holds true and observe its behaviour we can determine what its goals are. Thus, belief and desires are construed as interconnected concepts.

One difficulty with this view is that there are other cognitive states besides belief which interact with desire to produce action. For instance, although someone might not believe that *p* he may accept *p* as true in the context, guess that *p* or assume *p* and act on it (Bratman 1999). To illustrate, suppose John hears from a few sources that the housing market is going to crash. However, those sources are not completely reliable and John has some evidence that the market will not crash and, consequently, he does not form the belief that home prices will go down. Furthermore, John intends to sell his apartment and if the market were to go down he would not be able to sell it for a good price and will face bankruptcy. Given the cost attached to being wrong, John accepts in that specific context that home prices will go down and decides to sell his apartment. In that context, his acceptance or assumption has the same functional role with belief. John acts as if he believed that the market is going to crash but he does not really believe it and decided, based on pragmatic reasons, to accept that claim. Furthermore, David Velleman argues that, in the context of pretence, imagination may also motivate behaviour. When a child pretends to be an elephant his behaviour is caused by his imagination regarding the way elephants act. The child may, for instance, pretend that a chair is a pail of water and form the desire to drink. In that context, its behaviour is caused by its imagination not its beliefs (Velleman 2000, p. 255-263). But, in general, we want to be able to distinguish

believing that *p* from accepting or assuming that *p* in a context and from imagining that *p*.

Robert Stalnaker points out a different drawback of accounting for belief in terms of its capacity to produce action. On his view, this account gives rise to a relativity or indeterminacy of the content of beliefs. For instance, suppose that Mary wants her neighbor Fred to suffer because he insulted her and she believes that she can achieve her goal by playing the cello badly at three o'clock in the morning. Stalnaker imagines someone arguing that Mary wants Albert to suffer, not Fred. However, the fact is that Mary has never met or heard from Albert and that Albert lives 3.000 miles away. In spite of this fact, the proponent of the alternative hypothesis might argue that Mary believes that Albert is her neighbor, that Albert insulted her, and that Albert's name is 'Fred'. Stalnaker observes that the two hypotheses can explain Mary's behavior equally well. He attempts to solve this puzzle by pointing out that beliefs "represent what they represent not only because of the behaviour they tend to cause, but also because of the events and states that tend to cause them" (Stalnaker 1984, p. 18). Thus, given that Fred, not Albert, played a causal role in producing Mary's attitudes we can infer that her beliefs are directed at Fred. At the end of his discussion of this example Stalnaker draws the following conclusion: "The relativity of content that results forces us to recognize that belief is a backward-looking propositional state. What I want to suggest is that belief is a version of the propositional relation called indication. We believe that *p* just because we are in a state that, under optimal conditions, we are in only if *p*, and under optimal conditions, we are in that state because *p*, or because of something that entails *p*" (Stalnaker 1984, p. 18).

Thus, according to Stalnaker, focusing on the belief's tendency to produce action does not delineate what is specific about this cognitive attitude. We should also analyze the mechanisms which produce beliefs. Stalnaker's characterization of belief in terms of the relation of indication is closely connected with the ideas that belief has a mind-to-world 'direction of fit' and that belief 'aims at truth'. According to Searle, belief has a mind-to-world direction of fit; its function is to

represent how things are in the world. Intention and desire, on the other hand, have a world-to-mind direction of fit; their content is to be made true[7]. He writes: "It is the responsibility of the belief, so to speak, to match the world, and where the match fails I repair the situation by changing the belief." (Searle 1983, p. 8).

The idea that belief is an intentional state which 'aims at truth' intends to capture the specific difference between belief and other intentional states (Williams 1970; Velleman 2000; Shah 2003; Steglich-Petersen 2006; Adler 2002; Engel 2005; Wedgwood 2002). The fact that belief bears this relation to truth is constitutive of belief in the sense that it helps us distinguish belief from other intentional states like guessing, imagining, assuming and so on. Although guessing, imagining and assuming are cognitive states which involve accepting a proposition as true, belief is the only state which 'aims at truth'. Velleman spells out the metaphor by noting that belief involves accepting a proposition as true with the aim of accepting a truth. Accordingly, belief involves more than accepting a proposition as true. For example, in making a supposition or assumption one can accept a proposition as true for the sake of argument without actually believing it. However, when one believes a content one does not accept it as true for the purpose of seeing what follows from it. Moreover, Velleman points out that beliefs are 'regulated for truth'; the way they are formed and revised is truth-conducive (Velleman 2000; Shah and Velleman 2005). One forms a belief with the aim of truthfully representing the world and if one's belief turns out to be false one ought to revise it. It might be argued that the difference between genuine belief and accepting something as true in a context lies in the way these two

[7] Searle introduces the phrase 'direction of fit' in the context of his analysis of speech acts (Searle 1979, p. 4). He distinguishes the two directions of fit, for this application, as the words-to-world (assertion, predictions, etc.) and the world-to-words direction (commands, promises, etc.). Searle indicates that the best illustration of this distinction is provided by Elisabeth Anscombe in her book *Intention* (1957). In her presentation of the distinction, Anscombe describes a man going shopping with a shopping list while being tailed by a private detective listing the man's purchases, and asks what distinguishes the shopping list from the detective's list. She answers the question thus: "It is precisely this: if the list and the things that the man actually buys do not agree, and if this and this alone constitutes a mistake, then the mistake is not in the list but in the man's performance (...) whereas if the detective's record and what the man actually buys do not agree, then the mistake is in the record" (Anscombe, 1957, p. 56). In other words, the purpose of the shopper's list is to get the world to match the list whereas the purpose of the detective's list is to match the world.

attitudes are formed. Whereas one's acceptance of something as true in a context is influenced by pragmatic factors, by the need to make an immediate practical decision, accepting something as true with the aim of accepting a truth is only influenced by evidential considerations[8].

A related feature of belief which differentiates it from other cognitive states is *transparency*. This phenomenon is described in different, but related, ways by different philosophers. According to Nishi Shah[9], transparency is a phenomenon that occurs in deliberation about belief; deliberation framed by the question *whether to believe that p*. The phenomenon is described as follows: when asking ourselves *whether to believe that p* we immediately recognize that this question is settled by, and only by, answering the question *whether p is true*. From the first-person perspective there is no intermediate inferential step between the two questions (Shah 2003, p. 447). Transparency is an essential feature of belief because other attitudes which involve accepting a proposition as true do not exhibit transparency. For instance, the question whether to assume that *p* does not usually give rise to the question whether *p* is true.

Adler characterizes transparency as follows: from the first person perspective, belief is factive in the sense that believing *p* is equivalent with it being a fact for the agent that *p*. The agent can see directly through the belief to the state of affairs it represents. He notes: "When one attends to one's belief that *p*, one takes it to be the case that *p*" (Adler 2002, p. 193). Adler's account of transparency is similar to Shah's in the sense that, from the first-person point of view, believing *p* as a result of theoretical deliberation involves having answered

[8] Velleman's view will be discussed in more detail in Chapter 5.

[9] Shah's characterization is inspired by Richard Moran's account of transparency (Moran 1988). According to Moran, "transparency means that I must treat the practical question 'Shall I believe *p*?' as the impersonal theoretical question about *p*, and this means that the reasons I may have for adopting the belief are restricted to reasons connected with the truth of *p*" (Moran 1988, p. 146). Shah adds the qualification that the question 'Shall I believe *p*?" is a question which frames doxastic deliberation. He draws the difference between the deliberative question about belief and practical questions of the form: 'Would believing *p* be desirable or prudent?' Answering the latter kind of question does not result in the formation of a belief.

the question whether *p* is true in the affirmative. Once the agent accepts that *p*, *p* becomes a fact for him.

Both Shah and Adler are evidentialists about reasons for belief: they defend the idea that only evidence can be a reason for belief. Shah derives this thesis from the description of the phenomenon of transparency. If the only way of making up one's mind about whether to believe *p* is by answering the question whether *p* is true then only considerations relevant for answering the latter question have any bearing on belief. However, Shah and Adler stress, this view is perfectly consistent with the fact that sometimes our beliefs are influenced by non-evidential, pragmatic factors, like in cases of self-deception or wishful thinking. Nonetheless, they argue that the influence of these practical reasons goes unacknowledged; the agent does not explicitly take these factors into consideration when reflecting about belief.

Pragmatists about reasons for belief (Foley 1993; Papineau 2002) agree with the fact that usually citing pragmatic reasons for belief is ineffective. One cannot immediately decide that the earth is flat when one is offered one million dollars for forming that belief. Pragmatists agree that normally we have to go through an elaborate process of self-deception in order to trick ourselves into giving our assent to such statements. However, the essential difference between pragmatists and evidentialists is that pragmatists do not think that these psychological facts about belief should be explained by reference to the idea that there is an internal, constitutive relation between belief and truth. They argue that this relation is contingent and the fact that we cannot directly believe something based only on pragmatic reasons is a contingent psychological fact about our minds.

Although assessing the debate between pragmatists and evidentialists is beyond the scope of this work, I will formulate an objection against the idea that there is only a contingent connection between belief and truth. One challenge pragmatists about reasons for belief face is giving an account of belief without mentioning the relation between belief and truth. If we can believe something we

know is false then what is the difference between belief and pretence, for instance, or between belief and imagining or fantasizing? Although, we may agree with the idea that there is no clear-cut difference between believing that *p* and pretending that *p* (Morton 2003), we may still argue that *the concept* of belief is different from the concept of pretence. We want to say that when someone pretends to be an elephant, just for fun, he does not actually believe that he is an elephant. Similarly, we want to say that imagining or fantasizing about being a well-known rock star, for example, is different from actually believing that you are a rock-star. This distinction can be made even if there are people who, because of a psychological disorder or slight cognitive malfunction, have a disposition to accept as true what they fantasize about or do not have a good grasp on the difference between reality and fantasy.

2. The norm of truth – formulation and role

In addition to acknowledging the fact that there is an intrinsic relation between belief and truth, many philosophers also argue that truth is also prescriptive or normative for belief. One widespread way of arguing for the prescriptivity of truth for belief is by supporting the idea that belief is governed by a 'norm of truth'. The 'norm of truth' for belief is taken to be both constitutive and prescriptive (Shah 2003; Shah and Velleman 2005; Boghossian 2003; Wedgwood 2002; Engel 2005). The advocates of this view cash out the metaphor that belief 'aims at truth' in terms of the idea that beliefs are regulated by the norm of truth. However, they propose different formulations of the truth norm. The initial formulation of the norm goes like this:

(NT1) An agent ought to believe *p* if and only if *p* is true.

One standard argument against this formulation is that it implies that if a proposition is true then an agent ought to believe it. However, the argument goes, it is impossible for a finite human being to believe all truths about the world and,

given that 'ought' implies 'can', it does not make sense to prescribe to a finite human being something that is beyond her capacities (Hattiangadi and Bykvist, 2007).

As a result of this criticism the norm of truth has been reformulated. There are two prominent statements of the norm, one proposed by Boghossian (2003) and Shah (2003) and the other by Wedgwood (2002). Boghossian and Shah formulate the truth-norm as follows:

(NT2) An agent ought to believe that *p* only if *p* is true.

This weaker norm does not state that if *p* is true then one ought to believe it. But rather that: if an agent ought to believe *p* then *p* is true.

Wedgwood states the norm of truth as follows:

(NT3) if an agent considers *p*, then he ought to believe *p* if and only if *p* is true.

On this formulation of the norm, one ought to believe *p* when one thinks that *p* is true. This formulation also avoids the objection from 'ought' implies 'can' because the agent is not obligated to believe every true statement there is but only those statements that he considers and turn out to be true. The rule captured by (NT3) is thus based on the plausible idea that any proposition you can consider, you can believe.

As indicated, (NT2) and (NT3) are said to be constitutive of belief in the sense that they are part of *understanding* the concept of belief[10]; the fact that it is

[10] "When philosophers use the idea of some fact B being constitutive of some other fact A, they typically mean, I think, not that A analytically implies B, but rather a stricter condition along the following lines: that it's a condition on understanding what it is for A to obtain that one understands what it is for B to obtain. In other words, B is constitutive of A means: grasping the concept of an A-fact requires grasp of the concept of a B-fact." (Boghossian 2003).

regulated by truth is what differentiates belief from other intentional states (Boghossian 2003; Shah and Velleman 2005; Wedgwood 2002).

Nishi Shah tries to *explain* the phenomenon of doxastic transparency by reference to the fact that the norm of truth plays a *prescriptive* role in deliberation about belief. When engaged in theoretical deliberation, the epistemic agent tries to answer the question *whether to believe that p or not*. Thus, Shah notes, the question that frames doxastic deliberation and sets it apart from other types of reflection is a question about belief; the agent makes use of the concept of belief. But belief is defined in terms of the norm of truth. Shah writes: "To say that it is a conceptual rather than merely a metaphysical matter that truth is the standard of correctness for belief is to say that a competent user of the concept of belief must accept the prescription to believe that *p* only if *p* is true for any activity that he conceives as belief-formation." (Shah 2003, p. 470) It follows that doxastic deliberation is regulated by the norm of truth because it is an activity framed by the question *whether to believe p or not*.

3. The Objection from Normative Force

One important objection against the idea that the norm of truth has a prescriptive role in doxastic deliberation is that the norm of truth does not have genuine normative force (Katrin Glüer and Åsa Wikforss 2009; Steglich-Petersen 2006; Hattiangadi and Bykvist 2007). The proponents of this objection make two related points. First, they argue, it is unclear what the norm requires one to do, how this norm is supposed to guide one in the context of deliberating about what to believe. For instance, it seems that (NT2) and (NT3) entail that if *p* is false than it is not the case that one ought to believe *p*. However, the consequent of the conditional is not equivalent to the claim that one *ought not* to believe that *p* (Hattiangadi and Bykvist 2007). The claim that it is not the case that one ought to believe that *p* is compatible with the claim that one is permitted or entitled to believe that *p*. Moreover, it looks like the only positive guidance we can get from (NT3) is that if we consider whether *p* and *p* turns out to be true then we should

believe it. But it is not clear how this norm can offer us guidance in deliberation about belief because we engage in such deliberation precisely when we *do not know* whether *p* is true or false.

The second, closely related, critical point is that it is difficult to see how the norm can be broken in doxastic deliberation given that theoretical deliberation exhibits transparency. In other words, it is difficult to see how the possibility of breaking the norm of truth is compatible with the phenomenon of transparency. Steglich-Petersen (2006) develops the second critical point against Shah's attempt to explain doxastic transparency by reference to the norm of truth. He argues that Shah's explanation is flawed because it is impossible for an agent *to break the norm of truth* while engaged in doxastic deliberation. Since a norm or a prescription is by definition something that can be broken it follows that the norm of truth is not actually a norm; doxastic deliberation is not regulated by any norm. Steglich-Petersen writes:

> According to Shah and Velleman's explanation it is a necessary condition for an instance of deliberation to count as deliberation about belief that the deliberation exhibits transparency. This means that if transparency is produced by the norm of belief, this norm motivates one necessarily and inescapably to act in accordance with it. The transparency is immediate and does not involve an intermediary question about whether to conform to the norm of belief; the norm is thus unlike the norms governing promising. It is thus doubtful whether a consideration which necessitates motivation should be considered a normative consideration at all. This argument does not undermine the claim that deliberation about belief necessarily exhibits transparency, but only challenges the thesis that a norm or normative considerations can explain the necessity (Steglich-Petersen 2006, p. 507).

Steglich-Petersen draws a contrast between the norms governing promising and the norm of truth. In the case of promising, a speaker can think about whether he wants to keep his promise or not. The speaker has the freedom

to decide whether to keep his promise or not even if he acknowledges the fact that he ought to keep his promise. This shows that the norm of keeping one's promise does not necessarily motivate the speaker to keep his promise. By contrast, Steglich-Petersen notes, when an epistemic agent is deliberating about what to believe, the norm of truth necessarily and inescapably motivates him. The possibility that the norm of truth does not immediately motivate an agent is in conflict with the phenomenon of doxastic transparency. For instance, if the agent considers the question whether to believe that *p* and then, instead of thinking whether *p*, he wonders whether *p* is useful for him to believe then, according to Shah's definition, the agent's reflection does not count as deliberation about belief. The subject considers whether believing *p* is useful, not whether *p* is true. Thus, his deliberation does not exhibit transparency and, since transparency is an essential feature of doxastic deliberation, the agent's reflection is not deliberation about what to believe. However, Steglich-Petersen argues, following a norm involves being able to break it; that is, the claim that the norm of truth governs doxastic deliberation implies that we are able to break it in the context of deliberation about belief. He concludes that doxastic transparency cannot be explained by reference to following doxastic norms in the way proposed by Shah because *the possibility* of breaking the norms in the context of deliberation about belief is in conflict with the claim that transparency is a necessary condition for deliberation about belief.

4. Objective and Subjective Norms

Both Boghossian and Shah are aware of the fact that (NT2) does not offer clear guidance as it stands. Boghossian claims that the norm of truth is an 'objective norm' in the sense that it is a norm 'whose satisfaction is not transparent', but this does not mean, he maintains, that it is not a real norm. To illustrate his point he draws the analogy between the norm of truth and the norm: 'buy low and sell high!' Although it is not clear how to follow this norm in particular cases (one does not know when the market is at a low) it is still a genuine norm. Shah and

Boghossian argue that the fact that the norm of truth is an objective norm implies that it does not offer guidance to an epistemic agent by itself, but only in conjunction to subjective epistemic norms. We can attempt to follow the objective norm of truth by following the subjective norm of evidence:

(NE) An agent ought to believe that for which he has evidence.

Shah writes:

> The objective norm of belief prohibits believing that *p* unless *p* is true. But since we cannot directly check our representations against the truth, we need subjective norms by which to guide our doxastic activities. The role of the objective norm, whose acceptance is expressed in the phenomenon of transparency, is to provide *a standard of success* for subjective norms of good evidence that an agent can directly apply to his deliberation. It thus is not a brute fact that subjective norms for rational belief must be evidential in character. Rather, this constraint falls out of transparency, which is just an agent's recognition of the authority of the objective norm of truth (Shah 2003, p. 471).

Thus, the fact that an agent applies the norm of truth to his doxastic deliberation explains the requirement of evidence for belief. The agent ought to believe that proposition for which he has evidence. Evidence points to truth and therefore, in searching for the truth, one has to be sensitive only to evidential factors. Only evidential factors can be *reasons* for belief.

Critics of the idea that truth is normative for belief attack both the suggestion that (NT2) and (NT3) are objective norms analogous to the norm 'buy low and sell high!' and the idea that the norm of truth has prescriptive power only in conjunction with the subjective norm of evidence. First, Katrin Glüer and Åsa Wikforss (2009) agree with the idea that the norm 'buy low and sell high!' is not transparent because in most cases we do not know whether the circumstances in which the norm applies obtain. In consequence, if we knew that the market is at a low we would know what we should do to follow the norm. However, they ask, what are the relevant circumstances in connection with the norm of truth? As indicated, (NT3) states that one should apply the norm when one believes that *p* is

true; it implies that if *p* is true then one should believe that *p*. In other words, if the subject thinks that *p* is true then she should form the belief that *p*. Glüer and Wikforss point out that no guidance can be had from this norm. On the other hand, if we consider (NT2), it appears that this rule should be applied when *p* turns out to be false. But, as Hattiangadi and Bykvist observe, the only thing that follows from this is that it is not the case that one ought to believe that *p*. But the agent is under no obligation to *reject p* since he is permitted to believe it, but not required to. Thus, the objection goes, the fact that the norms (NT2) and (NT3) are construed as objective norms does not show that they are genuinely prescriptive.

In his reply to Glüer and Wikforss, Steglich-Petersen (2009) rightly notes that the condition for the application of (NT3) is not, strictly speaking, that the subject believes that *p* is true. The norm states that if an agent considers p then he should believe *p* if and only if *p* is true. In other words, if it is a fact that p is true then the agent should believe p. But the agent can track that fact without necessarily tracking his belief that *p* is true. The norm does not say anything about what the agent believes. Similarly, an agent can accurately represent the fact that the market is at a low, without having explicit, conscious beliefs about the state of the market. Although Steglich-Petersen's point is well-taken, it is still hard to see how the norm of truth, as described by Shah and Wedgwood, can offer us guidance in deliberation about belief, since it is obvious that when we engage in such deliberation we do not know whether *p* is true or false. The truth norms tell us what to do in case *p* is either true or false, but this offers no help in our reflection *whether p is true or false*.

Let us now turn to the idea that the prescriptive power of the norm of truth consists in the requirement for evidence. There are three related objections against the suggestion. First, Katrin Glüer and Åsa Wikforss argue that the two norms involve different notions of correctness: "I may believe that which is false, but on excellent grounds. Conversely, I may believe what is true on flimsy grounds" (Glüer, Katrin and Wikforss, Åsa 2009). Thus, the fact that someone follows the norm of evidence does not guarantee that he will succeed in following the norm of

truth. Therefore, the normative force of the norm of truth is not exhausted by the prescriptive force of the norm of evidence. Therefore, the onus is still on 'normativists' like Shah and Boghossian to show what is the specific normative force of the norm of truth over and above the norm of evidence.

Second, Shah's and Boghossian's suggestion implies that following the norm of truth is beyond an agent's capabilities. If an epistemic agent wants to follow the norm of truth all he can do is follow the norm of evidence. However, this does not guarantee that he will *succeed* in following the norm of truth; the norm of evidence is different from the norm of truth. A responsible agent may gather an impressive amount of evidence for a belief that turns out to be false. In consequence, it seems that it is not within one's power to follow the norm of truth. But 'ought' implies 'can': it makes sense to impose a prescription on an agent only if he can follow it. However, in this case, if it is impossible for the agent to follow the prescription then we are not dealing with a genuine prescription.

The third objection against the suggestion that the positive normative force of the norm of truth can be captured in terms of the norm of evidence goes as follows. The norm of truth is supposed to guide agents who are considering what to believe, who are engaged in doxastic deliberation. But doxastic deliberation is defined as a mental process in which the subject forms a belief on the basis of appreciating reasons. Shah writes: "the function of deliberation as opposed to mere reflection is to come to a decision or belief on the basis of one's appreciation of reasons." (Shah 2003, p. 472). Therefore, it seems that there is no room for breaking the norm of evidence in the context of doxastic deliberation. If someone forms a belief on the basis of no evidence then that belief is not the result of deliberation about what to believe.

It is noteworthy that neither Shah nor Boghossian state explicitly how much evidence is required for belief (Steglich-Petersen, 2006). One option would be that the agent ought to have conclusive evidence for his belief; evidence that establishes the truth of the belief. But again, this norm is in conflict with the

intuitive principle that 'ought' implies 'can' and it is clear that agents have limited amounts of time and cognitive resources to invest in gathering evidence. At this point 'normativists' are again faced with the problem of specifying exactly how the requirement of truth for belief translates into a requirement for evidence.

5. Doxastic Transparency and Prescriptivity

As mentioned, an important objection to the idea that the norm of truth is prescriptive is that it is not clear how this norm can be broken in the context of doxastic deliberation, given that this kind of deliberation exhibits transparency. As we have seen, Boghossian argues that the norm of truth is like the norm 'buy low and sell high!', in the sense that it is not always transparent what one ought to do in order to follow the norm. But in the case of the norm 'buy low and sell high!' we have a distinct idea of when the norm is broken; someone can *decide* not to follow the norm and buy high and sell low, for instance. However, can someone decide to break the norm of truth? Norms like (NT2) and (NT3) are said to forbid two things. First, they are supposed to stop agents from forming beliefs based on pragmatic, non-evidential reasons (Steglich-Petersen 2009). Second, although it does not strictly follow from either of the formulations, the norm of truth is meant to forbid the agent from believing something he knows is false.

Let us suppose, for the sake of argument, that an agent can break the truth norm either by believing something directly based on pragmatic reasons or believing something he knows is false. I think Shah and Wedgwood should reject this possibility because it is in tension with the idea that belief is defined as a state which 'aims at truth'. Recall that 'normativists' like Shah argue that belief is transparent: when asking ourselves *whether to believe that p* we immediately recognize that this question is settled by, and only by, answering the question *whether p is true*. From the first-person perspective there is no intermediate inferential step between the two questions (Shah 2003, p. 447). In other words, from the first-person perspective, belief is *factive*, the agent 'sees' through the belief the state of affairs that it represents (Adler 2002).

It might be argued that in *cases of self-deception* an agent ends up believing something he knows is false. David Papineau describes two such cases:

> Consider people who aim deliberately to mislead themselves. Suppose an elderly man realizes that he is likely to be upset if he learns about the real probability of his developing cancer, and so arranges to avoid any evidence that might undermine his sanguine belief that this probability is low. Or suppose an adolescent youth learns that people with an inflated view of their own worth are generally happier and more successful, and so deliberately seeks out evidence which will make him think overly well of himself. Of course, there are familiar psychological difficulties about *deliberately arranging to have false beliefs*, but examples like these suggest they are not insuperable.
>
> Note that these examples involve the manipulation of belief by the deliberate avoidance or pursuit of evidence, not any deliberate refusal to tailor your beliefs to the evidence you already have. I agree that cases of this latter kind are of doubtful psychological possibility. It is intriguing to consider exactly why this should be so (Papineau 1999).

It is noteworthy that the cases described by Papineau are not examples where an agent deliberates whether *p* is true, finds out that it is false and then still decides to believe *p*. This sort of case would be in violation of the norms (NT2) or (NT3), if they are taken to imply that an agent ought not believe something he knows is false as a result of deliberation. Thus, the elderly man does not find out as a result of careful deliberation that the probability of his getting cancer is high and then decides to believe that it is low. Rather, Papineau argues that the elderly man wants to avoid forming the belief that the probability of his getting cancer is high and to this end he does not consider evidence which might support that belief and only seeks evidence which shows that he is not in real danger of getting cancer. There are two possibilities: either the self-deceiver adopts this policy towards evidence in a conscious way or he does it unconsciously. In the first case,

when the agent deliberates whether the probability of his getting cancer is low he decides to focus on specific types of evidence and ignore other relevant evidence (e.g. what the doctor says) because he is *afraid* that considering *all* available evidence will show that he is in real danger of having cancer. However, Shah would deny that this type of reflection is an instance of doxastic deliberation because non-evidential factors explicitly influence the outcome of deliberation. In other words, in his reasoning the subject does not accept the authority of truth for belief which gives rise to the requirement of evidence. Since only evidence is a reliable guide to truth, the agent involved in genuine doxastic deliberation would only explicitly consider evidential factors and no pragmatic factors. Thus, conscious self-deception is not an example of breaking the truth norm in the context of deliberation about belief. On the other hand, if the elderly man has an unconscious policy to avoid evidence about the real probability of his developing cancer, then he does not break the norm of truth at a conscious level. Shah agrees with the possibility that self-deception might influence doxastic deliberation but he claims that such influence must go unacknowledged[11] (Shah and Velleman 2005, p. 8, n. 16).

It is noteworthy that Papineau and Dretske take these kinds of examples to show that truth is not an intrinsic goal of belief but rather it is a goal we sometimes have. In some cases, falsity is what serves us best (Papineau 1999; Dretske 2000) Thus, they argue, the norms of truth and evidence are hypothetical rather than categorical; they depend on the agent's desire to have true beliefs. But if the norm of truth is a hypothetical norm then we can have a grasp of the concept of belief independently of this norm, and, therefore, truth is not constitutive of belief. However, as I have argued in the first section, philosophers

[11] Frankish (2004, chapter 8) offers an account of self-deception along these lines. He argues that the self-deceiver has the unconscious belief that *non-p* and a strong unconscious desire that *p*. Moreover, consciously entertaining the thought that not-*p* fills the agent with anxiety and they are unable to accept not-*p* and consider its consequences. As a result, the agent pursues a *shielding strategy* of avoiding any evidence in favor of not-*p* and actively seeking evidence which supports *p*. I think that Frankish's account offers a conceptual framework in which we can offer a coherent and non-paradoxical description of self-deception. In particular, it avoids the problematic claim that the agent consciously believes both *p* and not-*p* at the same time.

who argue that there is only a contingent connection between truth and belief are faced with the challenge of drawing distinctions between the concept of belief and concepts like imagining or pretending – distinctions which are made in folk-psychology.

To sum up, I think that 'normativists' like Shah and Boghossian cannot accept the possibility that an agent can believe what he knows is false because this possibility is in conflict with the idea that the norm of truth is constitutive of belief. But then, given that they are committed to the truth norm being genuinely prescriptive, they have to show how it can be broken.

6. The Norm of Truth and Rule-Following

Another objection, formulated by Katrin Glüer and Åsa Wikforss, against the idea that belief is governed by the norm of truth involves the idea of rule-following. Glüer and Wikforss start by observing that there is an important distinction between regularity of behavior and rule-following. In his book *Rationality* (1964) Jonathan Bennett describes this distinction by reference to the behaviour of honey-bees. When a honey-bee finds a source of sugar it returns to the hive and performs a certain dance. Other bees observe the dance and fly straight to the food. Apiologists have found rules which correlate certain aspects of each dance with the distance between the hive and the discovered food, the direction from the hive to the food and the concentration of sugar in the food. Although we may say that the behaviour of the bees is regular and it is covered by rules, we would not claim that they are guided by rules. Bennett observes that the fact that, from time to time, a bee's dance misrepresents the location of food is not sufficient to make their behaviour genuinely guided by rules. He argues that rule-following involves an awareness of the rules as rules and of breaches of rules as breaches of rules. Thus, he imagines that if the bees which were misinformed about the location of

food were to have a dance which we could interpret as a denial[12] of the dance performed by the discoverer of food then we would be more justified in saying that the bees are aware of rules (Bennett 1964, p. 8-21).

Like Bennett, Glüer and Wikforss argue that the mark of rule-following behavior is the fact that it is intentional: the agent is aware of the rules and intends to bring his behaviour in accord with them. Both the advocates and critics of the norm of truth view seem to be in agreement regarding the idea that following the truth norm is an intentional action. Shah and Wedgwood argue that the norm of truth is an explicit norm which guides the agent's deliberation. Shah claims that the agent *accepts* the norm of truth and *aims* at following it. Katrin Glüer, Åsa Wikforss (2009) and Steglich-Petersen (2006) claim that the norm of truth is not a genuine norm because we cannot *decide* to break it like we can decide, for instance, to break the norm that forbids one to smoke in public places. According to this picture, rule-following is *an intentional activity and the rule-follower has an explicit intention to follow the rule and has beliefs about what the rule requires.* The fact that they adopt this picture of rule-following makes Shah's and Wedgwood's accounts vulnerable to the objection that their view leads to an infinite regress. Glüer and Wikforss make this point by observing that rule-following involves practical reasoning on the part of the subject (Glüer and Wikforss 2009, p. 55-57). Where R is an epistemic norm like the norm of truth, the reasoning of the subject motivated by this norm is of the following form:

(P1) I want to believe what is in accordance with *R*
(P2) To believe that *p* is in accordance with *R*
(C) I want to believe that *p*.

[12] Bennett's concept of a denial plays a similar role to Brandom's concept of sanction. Both kinds of bahaviour show that the participants in a practice have an awareness of the rules governing the practice. This idea will be developed and defended in Chapter 2.

Thus, it is clear on this construal that the agent has to have beliefs about what the norm requires. But then these beliefs also should be regulated by the norm of truth. In consequence, an infinite regress ensues.

7. Responsibility and Doxastic Control

A well-known objection against the idea that belief is governed by epistemic norms is that we are warranted in holding agents responsible for acting in accord with these norms - asses them as blameworthy or praiseworthy - only if they can exercise control over their cognitive lives. However, the most common form of control is voluntary control. Thus, it seems, the thesis that belief is governed by the norm of truth and norms of evidence leads to defending doxastic voluntarism, the controversial view that one can believe at will. Although normativists like Shah and Velleman reject doxastic voluntarism, they have to give an account of the control we have over our belief-formation. One potential problem with any such account of control is that it might make the account of the normativity of belief vulnerable to the circularity objection described in the previous section. If, for instance, the notion of doxastic control requires that the agent has explicit beliefs about what the epistemic norms requires of him and he wants to bring his reflection in line with these requirements then it follows that making sense of the notion of doxastic control requires having a grasp of the concept of belief and then the claim that belief is essentially normative is circular and leads to infinite regress.

This objection may be formulated in the form of a dilemma. On the one hand, if the normativist accepts that epistemic responsibility requires doxastic control then he needs to give an account of doxastic control without using the notion of belief. Besides doxastic voluntarism, the only prominent account of control is reflective control, but this account employs the notion of belief. On the other hand, if the normativist concedes that belief-formation is not within our control, then his claim that belief-formation is subject to prescriptions is called

into doubt. Something is a prescription if it can guide action or a certain activity. If a prescription does not offer any real guidance then it is not a genuine prescription.

8. Conclusion

In this chapter I presented and evaluated one way of developing the idea that truth is both constitutive and prescriptive for belief. The main thesis proposed was that belief is a cognitive state governed by the norm of truth. There are three main challenges such an account should meet. First, such an account should either explain doxastic transparency or show that this feature of belief is compatible with the prescriptivity of truth for belief. In so doing, the theory should make reference to prescriptions which can be broken, norms which have genuine prescriptive force. Second, an account of the prescriptivity of belief should avoid circularity and infinite regress: it should include an account of rule-following or norm-governed behavior which does not make reference to explicitly intentional notions. Third, the proponents of belief prescriptivity should consider the intuitive thesis that epistemic responsibility requires doxastic control. If normativists accept this claim then they should offer an account of doxastic control that does not make use of the notion of belief and does not commit them to the controversial view that one can believe at will. The next chapter will propose an alternative way of developing the idea that truth is constitutive and prescriptive for belief and show how this proposal can respond to the first two objections. Chapter 4 will discuss whether the account proposed in Chapter 2 has the resources to answer the third objection.

Chapter II - Belief, Assertion and Commitment to Truth

1. Introduction

In the first chapter I presented and evaluated one way of capturing the normative relation between belief and truth: the view that belief is subject to the norm of truth. This chapter will explore a different way of articulating the thesis that truth is prescriptive for belief: rather than claiming that truth is a norm for belief it might be argued that belief involves a commitment to truth; to believe that *p* means to be committed to the truth of *p*. The idea of commitment captures the following intuition: believing *p* is not a fact that happens to us, a fact with regards to which we are essentially passive. Rather, our beliefs represent our take on how the world is; our stance on how things are. The fact that we are active with respect to our beliefs explains why, if someone asks for reasons for our opinions, we feel the obligation to offer reasons. It is because our beliefs represent our take on what is true that we are answerable for them (Moran 2002; Hieronymi 2006). If beliefs were just facts about us, mental states with regards to which we can only take a spectators' stance and report on them passively, like we report on our sensations of pain or hunger, then they would not be mental states for which reasons could be asked.

In addition to having a normative import, the idea that belief involves a commitment to truth also captures the constitutive relation between belief and truth. What differentiates belief from other intentional states is that belief involves a commitment to the truth of the proposition believed whereas other intentional states do not involve such commitment. As Velleman notes, we need a way of distinguishing belief from other states which involve accepting a proposition as true, like supposing, assuming and so on. The idea proposed here is that assuming or supposing are not beliefs because they do not involve a commitment to the truth of the respective proposition.

The task of sections 2 and 3 is twofold: to explain why beliefs involve such a commitment to truth and to specify what are the obligations involved in undertaking such a commitment. Many defenders of this view (Brandom 1994, 2008; Baldwin 2008; Williams 2002) argue that there is a significant parallel between belief and the speech act of assertion. Like assertion, belief involves a commitment to the truth of the proposition. In the following I will present Robert Brandom's account of the relation between belief and assertion. I focus on Brandom's theory because he offers a detailed account of what is involved in undertaking a doxastic commitment. I will argue that construing belief along the ways suggested by Brandom enables us to show that the phenomenon of doxastic transparency is compatible with the genuine prescriptivity of belief. In addition, I will argue that Brandom's idea that the norms of assertion are implicit in practice gives as a starting point in answering the circularity objection. It is only when an agent follows an explicit norm that he needs to have beliefs and intentions with regards to what the norm requires. However, if a norm is implicit in a social practice, an agent can follow it without explicitly thinking about it. Sections 4-6 will build on Brandom's insight in an attempt to answer the circularity objection. For the sake of analysis, in my presentation of Brandom's ideas, I will start by assuming some of the controversial theses of his theory, and then I will discuss them explicitly.

2. Brandom on Belief and Assertion

2.1. Brandom's General Approach

Brandom distinguishes two ways of approaching the general problem of the relation between intentional states and speech acts. First, a theorist can start with an account of intentional states and of action and then construe speech acts as specific kinds of action. Second, linguistic theories of intentionality account for the propositional content of intentional states in terms of their relation to linguistic

practice. Philosophers like Paul Grice and John Searle are defenders of the first order of explanation. They explain speech acts by reference to the intentions and beliefs of the speakers. Brandom calls this kind of approach 'agent semantics' because it presupposes that the agent has the capacity for practical reason independently of having a language (Brandom 1994, 147). Also, this kind of approach presupposes that the propositional content of the intentional states of an agent can be accounted for independently of speech acts. Brandom opposes this order of explanation and defends the linguistic account of intentionality. He is committed to the idea that we can understand the propositional content of a sentence in terms of its inferential role and the notion of inference is interconnected with the notion of assertion. Assertion and inference are closely connected because assertion is an act for which reasons can be asked and which can be used in offering reasons. In offering a reason for an assertion one implicitly endorses an inference. Thus, Brandom uses the notions of inference and assertion in order to account for intentional conceptual content.

Nonetheless, Brandom does not attempt to offer an account of speech acts which makes no reference to intentional states and then explain intentional states in terms of speech acts. Like Davidson (2001b), he believes that thought and language are essentially interconnected and we cannot make sense of one without the other. In other words, Brandom does not attempt to give an account of belief in terms of assertion but construes them as intimately related.

Like Dennett (1981, 1989) and Davidson (2001b), Brandom defends an interpretivist account of intentionality. The basic idea of such an account is that having intentional states is to be explained in terms of attributing intentional states. According to Dennett, there is nothing more to a system's having intentional states than its being interpretable in terms of intentional states; that is, explaining and predicting its behaviour by reference to beliefs and desires. We can adopt the intentional stance towards animals but also towards chess-playing computers. In addition to ascribing beliefs and desires to the systems (or analogues of these states) we need to use the rationality assumption: the

assumption that the system will do what it ought to do given its intentional states. For instance, if we observe a mouse in a situation when he can see a cat waiting at one mouse hole and cheese at the other, we will predict that the mouse will go for the cheese. But in our calculation we rely on the assumption that the mouse will do what it ought to do (Dennett 1981, p. 10). The intentional stance is not to be contrasted with offering a physical description of a system because such a description is the product of adopting a different stance towards it: the physical stance. We adopt the physical stance towards objects whose behaviours can be explained and predicted based on their actual physical states and an application of our knowledge of the laws of nature. Assuming this stance, we can predict, for instance, that the branch of a tree will break under someone's weight. Both the adoption of the intentional and the physical stance involve a degree of simplification or idealization aimed at maximizing predictive success. A useful analogy here is with the way physicists use the notion of a center of gravity. When the trajectory of an object is calculated, we can represent the object by a single point, thus abstracting from its shape. We can make this simplification because the total weight of the object is concentrated in its center of gravity. In a similar fashion, we are entitled in attributing beliefs and desires to a cognitive system as long a these ascriptions allow us to better calculate or predict its future behaviour.

Both Dennett and Brandom argue that the cognitive systems that attribute intentionality are linguistic creatures. They have to possess intentional concepts and engage in linguistic activities like explanation, prediction, theory construction and so on. Thus, on their account non-linguistic creatures have derivative intentionality in the sense that their status as intentional systems depends on the intentionality of linguistic creatures who posses original intentionality. Brandom offers two connected reasons that explain this dependency. First, the intentional states of non-human animals have propositional content and we can understand propositional content only in terms of human rational beings participating in linguistic practices. Second, we explain the behavior of non-human animals by

reference to reasons but we can make sense of the concept of reason only in the context of the practice of giving and asking for reasons. In other words, we cannot understand what a reason is unless we understand what it means to ask for reasons and give reasons. In addition, we give reasons for what we do by making assertions. In consequence, grasping the notion of reason and of intentional action requires understanding the notion of assertion. For instance, when we see an animal running from a fire we say that it runs because it thinks that the fire is dangerous and wants to avoid danger. That is, "we seek to make its behaviour intelligible by treating it *as if* it could act according to reasons it offers itself" (Brandom 1994, p. 171). But the act of offering reasons is intelligible only in the context of linguistic the practice of giving and asking for reasons. In this sense, the animal has *derivative* intentionality.

The question that arises naturally at this point is: how is original intentionality explained? According to Brandom, the intentionality of linguistic creatures is also dependent on them being interpretable by other linguistic creatures, on their ability to participate in social linguistic practices. On this picture, only linguistic communities, not individuals, can be interpreted as having original intentionality (Brandom 1994, p. 61). A member of a community counts as an intentional agent only if he can successfully participate in the game of giving and asking for reasons; that is, if he can make assertions and draw inferences from them. This game confers genuine propositional contents to the claims and cognitive states of the participants. Moreover, the functional roles of belief and desire (and the notion of intentional action) can be explained once we have a grasp of the concepts of assertion and inference (Brandom 1994, p. 55-62). Brandom's view of assertion and its relation with belief is the focus of the next section.

2.2. Brandom on Assertion

As indicated, in Brandom's view, assertion plays a dual role: it is a speech act for which reasons can be asked and it can be used in offering reasons. In other words, assertion is defined by reference to the practice of giving and asking for reasons. The notion of assertion is inextricably linked to the notion of inference because when one gives one's reasons for an assertion one implicitly endorses the inference from the premise or reason to the conclusion. The semantic content of a claim is constituted by its inferential connections with other claims, inferential connections that arise out of the practice of giving and asking for reasons.

The deontic notions of commitment and entitlement play a central role in Brandom's account of assertion. These notions correspond roughly[13] to the traditional deontic notions of obligation and permission. Brandom uses these concepts to characterize both assertion and the related notion of inference. In order to understand the commitments and entitlements specific of assertion we first need to look at the way Brandom describes the inferential relations between claims. He draws the difference between commitment-preserving inferences and entitlement-preserving inferences. Deductive inferences and materially good inferences are commitment-preserving inferences in the sense that commitment to the premises of these inferences involves a commitment to their conclusion. An example of the latter kind of inference is: if A is to the west of B then B is to the east of A. Material inferences are content-determining inferences in the sense that it is part of the semantic content of the concepts involved in these inferences that they have this inferential-role. Inductive inferences are examples of entitlement-preserving inferences. Being committed to the premises of an inductive argument gives one the permission to endorse its conclusion but one is not committed to endorsing it. In addition to characterizing these inferential relations Brandom uses the concepts of entitlement and commitment to define the relation of material

[13] Brandom avoids using the notions of obligation and permission because they are associated with the idea that it is a superior who lays an obligation or offers permission to a subordinate (Brandom 1994, p. 160). The deontic statuses of commitment and entitlement are instituted by all the participants in linguistic practices.

incompatibility between claims: two claims are incompatible if commitment to one claim precludes entitlement to the other (Brandom 1994, p. 168-169).

The notions of commitment and entitlement are normative or prescriptive notions in the sense that they concern what an agent ought to do or is permitted to do as opposed to what he actually does. Thus, for Brandom, the practice or game of giving and asking for reasons is a normative practice. However, given his endorsement of a linguistic account of intentionality, Brandom's account of the practice of giving and asking for reasons cannot make use of explicit intentional notions; the participants in the practice do not have explicit knowledge of the norms that govern the practice and they do not consciously intend to follow those norms. Rather, the norms that govern the game of giving and asking for reasons are implicit in the practice of playing the game. Participating in this practice involves knowing-how to play the game rather than having explicit knowledge of the rules of the game.

The notions of commitment and entitlement are essential to understanding speech acts and, therefore, intentional states. In accord with his interpretivist view of intentionality, Brandom advocates phenomenalism about norms. On this view, we can explain having a normative or deontic status (being committed or entitled) in terms of attributing normative status. Phenomenalism explains undertaking commitments in terms of attributing commitments. An agent undertakes a commitment when he does something that entitles others to attribute that commitment to him. For instance, a speaker undertakes a commitment to do what he promised when his interlocutors decide that his utterance counted as a promise and attribute that commitment to him. In addition, failing to discharge a commitment makes one a candidate for negative sanctions. Thus, attributing commitments is closely related with the possibility of sanctioning failure to discharge the commitment. Moreover, given his aim of explaining normative practices in non-intentional terms, Brandom construes attributing commitment as adopting a practical attitude (i.e. having a disposition to sanction). Adopting a practical attitude is characterized as a doing and not an action. The concept of

doing is equivalent with that of behavior whereas the concept of action is connected with intentional notions like belief and desire.

Besides the difference between commitment and entitlement, there are two more distinctions essential to Brandom's characterization of the practice of assertion. First, given that assertion is used in communication, Brandom draws the difference between the intrapersonal and interpersonal aspects of assertion. The former notion concerns the inferential consequences of making an assertion for the speaker and the latter notion refers to the inferential consequences of the assertion for his interlocutors. Second, Brandom distinguishes between the authority and responsibility specific of assertion (Brandom 1994, p. 169-173).

In Brandom's analysis, when a speaker makes a claim he puts forward a sentence in the public arena as true, he undertakes a commitment to the truth of that content. The speaker authorizes his interlocutors to rely on the truth of the sentence asserted. However, the interlocutors rely on the speaker's assertion because they expect that the speaker was entitled to make it. In consequence, the asserter is held responsible for the truth of the assertion; if challenged, he ought to demonstrate that he was entitled to make the assertion. In addition, in undertaking a commitment to the truth of a sentence the speaker is committed, on the intrapersonal level, to what follows from that claim as a conclusion of a commitment-preserving inference. Furthermore, the speaker is not entitled to make claims that are incompatible with his original claim. On the interpersonal level, the speaker makes available to his interlocutors a premise which they can use in their own inferences (Brandom 1994, p. 173).

As mentioned, in Brandom's account, assertion involves a task-responsibility, a responsibility to do something; it implies that one ought to demonstrate that one was entitled to make the claim. In this sense, assertion is similar to promising, they both involve commitments to undertake certain doings. Brandom distinguishes three ways in which a speaker may demonstrate that he was entitled to make a claim (Brandom 1994, p. 174-175). First, the speaker may justify the claim by making another claim. In this case, he is entitled to the

commitment involved in his initial assertion if the members of his linguistic community take the inference between the two claims to be correct. Brandom argues that in justifying his claim a speaker can appeal to both deductive and inductive inferences. He writes: "Justificatory practices depend on entitlement-preserving inferences. But commitment preserving inferences are also entitlement-preserving." (Brandom 1994, p. 673) Second, a speaker may discharge his commitment by making reference to the authority of other speakers. Third, the speaker may appeal to his own authority as a reliable observer. Agents are entitled by default, Brandom argues, to make perceptual, non-inferential claims. In this case, the interlocutor who doubts the truth of the perceptual claims should demonstrate that he is entitled to challenge the speaker's authority as a reliable observer.

For Brandom, the success of an assertion is equivalent to transmitting knowledge: "The complex hybrid deontic status of knowledge defines the success of assertion. Treating an assertion as expressing knowledge – attributing to the asserter entitlement to the commitment undertaken thereby and endorsing the commitment oneself – is the response that constitutes the practical recognition of the authority implicitly claimed by the assertion" (Brandom 1994, p. 203).

Brandom submits that the concept of assertion defined in terms of the deontic statuses of entitlement and commitment plays the same theoretical role as the concept of belief. In other words, linguistic belief has the same conceptual features as assertion. He writes: "It has been suggested that the doxastic commitments undertaken by speech acts having the significance of assertions can serve as analogs of belief – that such deontic statuses can do much of the explanatory theoretical work usually done by the paradigmatic sort of intentional state" (Brandom 1994, p. 194). In other places, Brandom suggests that belief is inner assertion and doxastic deliberation is a private conversation regarding what is appropriate to assert.

3. Assertion, Transparency and Normative Force

3.1. Refining the Assertion Account of Belief

Let us have a closer look at the relation between belief and assertion and how exactly the analogy proposed by Brandom is supposed to work. One initial problem with the analogy is that assertion is described as a public act by which the speaker informs others of the truth of a sentence and, in consequence, undertakes a certain responsibility. However, it is not clear that when we believe something we inform others about anything; belief is not an act of communication.

One suggestion is that maybe belief is inner assertion, an act of informing ourselves about something – in this case we are our own audience. Another proposal is that belief is what we would assert in contexts we care about (Williams 2002). I will consider these two ideas in turn. First, to clarify, the deliberative belief that *p* is normally the result of an act of judgment, and this act of judgment is the result of deliberation about whether *p*. When we judge that *p* we do it because we were convinced by the reasons for *p*. Now, one difficulty with the inner assertion model is that it is hard to see what the point of asserting *p* to yourself is. In other words, once you made up your mind about *p* and accepted that *p* is the case, why would you have to assert *p*? This assertion seems redundant since you already know that *p*. The second suggestion is that belief is what we are disposed to assert in contexts we care about. But why are we disposed to assert *p* in such contexts? Because in our minds we considered the question whether *p* and were convinced by the arguments in support of *p*. So, the assertion is the expression of the judgment that p; the expression of our private acceptance of *p*. A different way of interpreting the second suggestion is that deliberation itself should be understood as putting ourselves in a context in which we really care about the truth of *p*. I agree with this idea, but once we are in this private context, what is the goal of actually making an assertion? Deliberation is triggered by in

interest in whether *p*, but once we made up our minds that *p* is the case, why do we still have to assert it to ourselves?

Then, the question remains, how can we define judgment and belief based on the model of the practice of assertion described in section 2? I think it's useful to start by offering an account of what is to deliberate about whether *p* is the case. As indicated, on Brandom's account of assertion, when an asserter is challenged and has to demonstrate entitlement to his claim, his audience must assess the inference from the claims offered as reasons to the initial assertion. Brandom suggests that "first-person deliberation is the internalization of such third-person assessment" (1994, p. 231). In other words, when we deliberate about belief we consider a sentence *p* as if it was asserted by someone else, we consider various reasons for it and, if we take the inference from one of these reasons to the claim under consideration to be correct then we accept the claim. In the public setting, if the asserter is able to vindicate the claim that p then his assertion is successful and his interlocutor accepts p as information or knowledge and commits himself to its truth.

In the case of inner deliberation, the agent is his own interlocutor and reflection consists in the appreciation of the reasons he is able to produce. This phenomenon is sometimes described as self-interrogation (Frankish 2004; Dennett 1991) and characterized as follows: "Since we routinely rely on verbal interrogation as a way of acquiring information from others, it would be natural for us to develop habits of verbal self-interrogation, instinctively questioning ourselves and supplying answers" (Frankish 2004, p.105). In deliberation, we not only question ourselves and supply answers but we also evaluate the correctness of those answers.

What are the prescriptions involved in this act of acceptance? In analogy with the third-person case, in the context of private theoretical deliberation, if the agent accepts *p* he becomes committed to the truth of *p*. One consequence of this commitment is that the agent ought to use *p* as a premise in future theoretical

deliberation[14]. This inner commitment parallels the third-person case, where accepting an assertion obliges one to use it as information if needed. For instance, if one accepts the claim that global warming is produced by humans and he is then asked, in a context of inquiry, whether global warming is caused by human activities then he ought to make public what he already accepted as information.

One apparent disanalogy between the public and the private commitment to truth is that it is not clear that the latter sort of commitment involves an obligation to prove that one is entitled to holding a sentence true. However, I think there are cases in which such entitlement is brought into question in first-person deliberation. Suppose someone accepts *p* as a result of deliberation and then uses *p* as a premise. However, it turns out that if he puts *p* in conjunction with a few other premises he gets a conclusion which is obviously false. Then, suppose that he is quite confident in the truth of the other premises. In this case, it makes sense for him to ask himself again why he holds *p* true. He needs to demonstrate entitlement to that claim. When the agent reevaluates his justification for *p* he might not find it as convincing as he did the first time. This is usually what happens when we change our minds about something. Sometimes, in light of new information, we reassess the justification we had for holding certain beliefs only to find it incorrect. In these cases, we press ourselves, as someone else might, to demonstrate our entitlement to the claims we hold true[15].

The definition of belief in terms of accepting an assertion also illuminates why we can ask for reasons for belief and why we have to prove entitlement to our beliefs in a public setting. Saying 'I believe *p*', on the present account, is synonymous to uttering 'I accept *p* as information'. When making this utterance the speaker commits himself to the truth of *p* and to having entitlement to that

[14] Keith Frankish (2004) proposes a similar idea. In his view, believing *p* is adopting a premising policy to the effect that one should use *p* as a premise in deliberations which are truth-critical with respect to premises. Frankish's view will be considered in more detail in Chapters 4 and 7.

[15] Jonathan Adler, a defender of the idea that belief parallels assertion, makes a similar point: "If we maintain a belief on becoming aware of it, we enter a claim to truth on ourselves. Effectively, we assert the belief to ourselves, and incur responsibilities that parallel those we incur when we assert our beliefs to others" (Adler, 2002). On the account defended here, the responsibilities we incur in the private sphere parallel those we incur when we *accept* an assertion in the public domain.

commitment. Thus, saying 'I accept that *p*' means that you take it that *p* is worthy of being asserted, that that information can be disseminated. If, after attending a pro-life convention and finding the arguments presented irresistible, you make up your mind that abortion is murder then you accept this sentence and you take it to be worthy of being further asserted and disseminated. Believing is accepting the assertion of *p* and such an acceptance implies that *p* is a truth which can be spread to others. This is in line with the observation made by a few philosophers (Wittgenstein 1953; Heal 1994) that the statement 'I believe that *p*' has roughly the same use, and thus similar meaning, to the simple assertion of *p*.

Clearly, there are cases in which an agent accepts a claim and does not intend to further share that information. But even in such a case, his belief is still subject to norms. In accepting a claim the subject exercises his authority and implicitly trusts his own epistemic competence in a specific subject matter. Once his doxastic commitment is challenged, either by his own further observations or reflections or by claims made by others, he has to reassess his own judgment. Thus, an agent can lose his entitlement to holding a proposition true even when he does not publicly assert it. Private belief is still subject to the force of reasons.

Thus, on this account, we can ask for reasons for belief, just as we ask for reasons for assertions, reasons that would prove entitlement, because both belief and assertion involve a commitment to truth.

3.2. Answering the First Challenge: Transparency and Normative Force

The first chapter presented three main challenges that an account which states that truth is prescriptive for belief should face. The first challenge is to show that the prescriptivity of truth for belief is compatible with doxastic transparency while at the same time preserving the genuine prescriptivity of the norms of belief. In other words, the challenge is to show that truth can be both constitutive for belief (in the way manifest in transparency) and prescriptive (involving prescriptions which can be broken). So far, the present chapter developed the view that belief

involves a commitment to truth, a commitment which has prescriptive force. The rest of section 3 will aim at showing that this commitment is compatible with, and can illuminate, the phenomenon of doxastic transparency and that the prescriptions which it involves are genuine.

Recall that, on Shah's and Velleman's account, belief transparency means that the deliberative question whether to believe that p gives rise to the question whether to judge that p which in turn gives rise to the question whether to accept the assertion that p. Given that assertion involves a claim to truth, we can only answer this last question by determining whether p is true. This is in analogy with the public case, when we ask the asserter to show us that his claim is true by offering reasons. Trying to answer the question whether p is true is what motivates the game of giving and asking for reasons played both in the public and the private arenas.

In consequence, given that in his inner deliberation the agent considers accepting an assertion, only considerations that bear on the truth of the sentence have any direct influence on his reflection. As a result, we can derive the requirement, emphasized by Shah, that only evidential considerations are relevant in deliberation about belief, from the fact that an assertion involves a commitment to truth, without relying on the controversial idea that belief is governed by on objective norm of truth.

Since belief is defined as the acceptance of an assertion and assertion is meant to transmit information or knowledge, it follows that, once the subject judges that p as a result of deliberation, he takes himself to know that p. This is why, as Adler (2002) rightly points out, from the first-person perspective, belief and knowledge have the same essential features: they are both factic and transparent. To believe p, on this picture, is to take p to be a fact. From the subject's perspective, the belief represents how things are. However, there is more to belief than taking a proposition as true. Since deliberative belief is the result of accepting a claim, the subject is vulnerable to a request for reasons and potential criticism from others. This normative aspect is captured by the idea that belief

involves a commitment to truth. It becomes clear now why, on this perspective, there can be no conflict between transparency and the prescriptivity of truth for belief. Transparency just means that the subject settled the question whether p is true. Prescriptivity means that the act of settling the question of p's truth has normative consequences: if challenged, the agent has to show why he found his reasons for p convincing, he must be sensitive to warranted criticism of his arguments, and he must take p as a premise in future deliberation.

As indicated in connection with the norm of truth view, Steglich-Petersen makes the critical point that genuine prescriptions are prescriptions which can be broken. If there are no circumstances in which we would say that a norm has been broken then it is not clear why we should even call it a norm. Now, my present suggestion is that there are at least two central prescriptions involved in accepting an assertion[16]: using it as a premise in future deliberation and demonstrating entitlement to it if appropriately challenged. I think both these prescriptions have genuine normative force. Someone who accepts the assertion that p but then claims that not-p breaks the obligation involved in the assertion of *p*. Similarly, an agent fails in discharging his commitment to the truth of *p* if he judges that *p* as a result of reflection but then does not use p as a premise in his deliberation. For instance, after accepting a new conclusion and rejecting a long held belief, some people have a tendency to fall back on old habits of thinking and acting. The new item of information, although accepted, fails to be integrated into their cognitive lives. A sexist English teacher, for example, may theoretically accept that women are just as intellectually gifted as men, but still find himself bewildered when one of his female students submits a very good essay. His first thought might be that it must have been plagiarized and he might invest a lot of time and energy into trying to prove it. This is why we need to keep track of our commitments both in the public arena and in private deliberation.

Similarly, as already mentioned, an agent may fail in demonstrating his entitlement to holding a sentence true, if pressed for reasons. In a social setting,

[16] Accepting an assertion also involves certain logical commitments of coherence and consistency. These inferential doxastic obligations will be discussed in Chapter 6.

there are at least four ways in which such a failure might occur: (1) the justification for the initial claim may be incorrect, (2) the claim that is used in justification may be a claim the speaker is not entitled to, (3) the informant who communicated the information may be unreliable (to the speaker's knowledge) or (4) the asserter himself may wrongly think that his perceptual mechanisms are reliable (despite evidence to the contrary). As was pointed out, Brandom explains the deontic status of being entitled in terms of the deontic attitude of attributing entitlement. However, an important qualification of this thesis is that the attribution of entitlement might itself be correct or incorrect. Thus, to be precise, having entitlement means being correctly taken to be entitled. We can judge whether someone is correctly or incorrectly taken to be entitled to a commitment by reference to the intersubjective epistemic standards which are adopted in a community. Therefore, it is always possible that an agent might think that he has good reason or evidence for his belief but fall short of the standard that is in place in his epistemic community. For instance, while deliberating about the quality of life in his city, one may commit the fallacy of hasty generalization and base his conclusion only on information about the people he knows, without realizing that all of them are members of the upper middle class. In this case, the agent has some evidence for his conclusion but the evidence is not sufficient to entitle him to draw the conclusion that the quality of life in his city is high. He considers an inference to be correct in his inner reflection, but his assessment of the inference is itself incorrect.

A different case of lack of entitlement is when the inference to the conclusion is correct but the premise of that inference is one the agent is not entitled to. For instance, someone may correctly infer from the claim that all fiction writers are rich the conclusion that John Grisham is rich. However, he would have a hard time proving entitlement to his premise, since there are so many counterexamples. The idea that we have to prove entitlement to the premises of our reasoning may give rise to the concern that proving entitlement is an activity that does not have an end. However, in answering this difficulty,

Brandom emphasizes that entitlement has a default and challenge structure. There are many claims to which we are entitled by default. As Wittgenstein puts it, 'justifications must come to an end'. There are Moorean sentences like 'I have two hands', 'There are other people', 'The Earth existed before I was born' which are not usually called into question. Moreover, if someone wants to question these kinds of basic truths then he has to prove entitlement to his challenge. A challenge to a claim, Brandom notes, can be seen as making an incompatible claim (Brandom 1994, p. 178). Thus, since the challenge is itself a claim, the challenger may have to prove entitlement to it. The speaker has the burden to prove entitlement only if the challenge is considered appropriate, by the standards which are in force in the linguistic practice of the community[17].

So far, the focus of Chapters 1 and 2 has been the normativity of deliberative belief, belief that is the result of explicit refection. However, since normativists want to establish the more ambitious claim that beliefs in general are constitutively prescriptive, it is worth analyzing the case of beliefs that are not the result of explicit deliberation. Sometimes we just accept the information that we get from newspapers, the internet or TV without thinking about it. Similarly, we normally accept what is given in our external perception and the information we get from introspection. In what sense is truth prescriptive for these kinds of beliefs? The proponents of the norm of truth account do not have a straightforward answer to this question. Remember that, on their view, deliberative belief is normative because deliberation is shaped by the question "Should I believe *p*?" and in formulating this question we make explicit use of the concept of belief which is defined by the norm of truth. However, once outside the context of deliberation, once we stop using the concept of belief, it is not clear

[17] It might be argued that skeptical challenges are examples of challenges which are not considered warranted in the ordinary practice of giving and asking for reasons. For instance, Descartes' dreaming argument and the Evil Demon argument do not carry too much weight in normal conversations. People resist the idea that they cannot distinguish dreaming from being awake and that there is such a thing as an Evil Demon. This shows that modern philosophers like Descartes have a sense of responsibility for grounding all their beliefs which is not shared by ordinary people and is not embedded in their practice. Arguably, if this sense of skepticism and epistemic duty were widespread, the linguistic practice could not have gotten off the ground.

what kind of guidance the norm of truth has to offer. Of course, Velleman might argue that perceptual belief 'aims at truth' in the sense that it is the result of perceptual mechanisms designed to track the truth[18]. Even if this is the case, it is not obvious why truth plays a prescriptive role in this context. The fact that a mechanism inside an agent was designed to function in a certain way does not have any implication regarding what the agent has to do.

I suggest that one virtue of construing the normative relation between belief and truth in terms of commitment is that there is a clear sense in which commitment is involved in both non-inferential and inferential beliefs. Recall that, on Brandom's account, justifying a claim by reference to another claim is only one way of demonstrating entitlement. If pressed for reasons, a speaker might defer to another speaker the responsibility of grounding a claim. Thus, in the case of information we get from mass-media, we can point to the authority of the source of information in order to justify our claim. In addition, Brandom indicates that we can support our perceptual judgment by reference to our authority as reliable observers. Thus, non-deliberative beliefs are also caught in the game of giving and asking for reasons. They are not just the output of cognitive mechanisms that are independent of us, but they are also items of information for which we are answerable.

Let us take a closer look at the way we can fail to demonstrate entitlement to our non-deliberative beliefs. In the case of testimony, when we defer the responsibility of grounding a claim, it might turn out that our source of information was, after all, unreliable. In that case, our interlocutor might pressure us to reevaluate our belief. One difficulty here is whether the speaker succeeds in demonstrating entitlement if it turns out that his source was unreliable. There are two possibilities: the speaker knew, or had good reason to suspect, that his informant did not have the required authority or he was ignorant of that fact. If we argue that, in the latter case, the asserter did not demonstrate entitlement then it follows that the asserter is responsible, in order to make sure that she follows the

[18] A critical discussion of the teleological account will be the focus of Chapter 5.

norm of entitlement, to find out whether her sources are reliable before accepting their testimony. But this is in conflict with the principle that 'ought' implies 'can'. It implies that we need to make extensive inquiries into the reliability of our sources of information – news channels, newspapers, friends, and so on – before forming any beliefs. I think Brandom's remark that entitlement has a default and challenge structure might be useful in this context. The idea is that the practice of assertion is not based on the principle of being guilty until proven innocent but rather that one is innocent until proven guilty. Thus, one is entitled to rely on the authority of one's interlocutor, as long as one does not have reason to doubt that authority. An agent is not, in consequence, blameworthy for trusting an unreliable source of information, as long as he was not aware that the authority of that source has been appropriately challenged. Then, the asserter can still count as having proven entitlement to a claim in this context, by deferring to his source, but it is the original speaker who might fail in discharging his responsibility. However, if the speaker knew that his source of information was unreliable then he is blameworthy and fails to prove entitlement. Such a failure occurs very often. For instance, there are newspapers and TV channels which are proven to be biased; the fact that they are not worthy of trust is common knowledge. In spite of this fact, many people uncritically accept their information.

Analogously, in the case of perceptual judgments, an agent cannot prove entitlement to his claims about colors if he knows that he is not a reliable observer of colors; that he is color blind, for instance. Similarly, the judgments concerning distant objects made by a person suffering from myopia do not carry much authority if she does not have her glasses on.

There is another aspect of our responsibility involved in our commitment to the truth of non-deliberative claims. Even if we are not blameworthy for not knowing that our interlocutors were not to be trusted or that our perceptual mechanisms are not reliable, we are still responsible to reassess our beliefs and change our epistemic habits once that unreliability is proven to us. This is a different way in which we are essentially active in connection with our non-

inferential beliefs and it shows that they are not only products of mechanisms which function automatically and independently of us. Instead, perceptual beliefs and beliefs based on testimony are caught up in the public game of giving and asking for reasons.

3.3. The Objection that 'Ought' Implies 'Can'

It might be objected that the doxastic obligations involved in undertaking a commitment to the truth of a claim are in conflict with the principle that 'ought' implies 'can'. For instance, we are not able to remember all the evidence we had for our beliefs and also, sometimes we may forget that we made up our mind about something and fail to use that claim in our deliberation. As Harman (1986) persuasively argued, we do not have to give up a belief just because we forget our reasons for holding it. Harman observes that, given the limited storage capacity of our long-term memory and our finite processing power, we need to avoid cluttering our minds with information that is not important. Thus, in his view, when one reaches a significant conclusion, one needs to remember the conclusion but not all the intermediate steps involved in reaching the conclusion (Harman, 1986, p. 42).

Accordingly, Harman criticizes the foundationalist Principle of Negative Undermining which states that one should stop believing p whenever one does not associate the belief with an adequate justification. Given that we normally forget our reasons for holding certain beliefs, this principle implies that we should reject many of our beliefs, which is unintuitive. Instead of following this disastrous epistemic strategy, Harman suggests that we should guide our cognitive lives by the Principle of Positive Undermining which claims that one should stop believing p whenever one positively believes one's reasons for believing p are no good. This implies that one has to be reminded of his reasons for holding p in the first place and shown that those premises are false. Later on, Harman defines belief as full acceptance, an acceptance which ends inquiry and claims that "full

acceptance ends inquiry into *p* in the sense that, having accepted *p*, one is justified in continuing to accept *p* in the absence of a special reason to doubt *p* or at least a special reason to reopen one's inquiry" (Harman, 1986, p. 48). A reason to reopen one's inquiry is fining new evidence which casts doubt on the truth of one's conclusion.

It is noteworthy that finding a special reason to doubt *p* is different than finding out that one's initial reasons were no good. In the former case, one is warranted in reopening the investigation regarding *p* whereas in the later case one should stop believing *p*.

In the following I will evaluate Harman's remarks and show that they pose no challenge to the idea that belief is prescriptive. We may start by noticing that an agent has to have certain minimum cognitive abilities in order to be able to participate in the practice of assertion. Someone cannot be a participant in such a practice if he has very poor short-term memory, for instance. In order to be a player in the game of giving and asking for reasons someone has to keep track of his and his interlocutor's commitments and entitlements in a normal conversation. To this end, the speaker has to remember what his interlocutor has said and to follow his inferences. If after following a short chain of reasoning, he is not able to remember what the premises were and what the initial claim was, then he's not fit for the practice. More importantly, we should observe that the norms of assertion and belief specified in the previous section do not imply the Principle of Negative Undermining. The norm of entitlement says that one should prove entitlement to a belief or acceptance if one is appropriately challenged. As long as one is not challenged one does not have to do anything. Recall that a challenge, on Brandom's construal, can take the form of an assertion that is incompatible with the speaker's belief that *p* and the speaker can ask for the hearer's entitlement to that assertion. In the event in which the challenger is successful in proving that his challenge is warranted, the speaker has to face it. But this is close to what Harman calls 'a special reason to reopen one's inquiry'.

Let us take a concrete example. Consider someone who believes that free will is compatible with determinism. Suppose that he accepted compatibilism a long time ago and does not remember the details of the reasoning that led him to this conclusion. Now, in conversation, his interlocutor presents a compelling argument for incompatibilism. According to the view presented here the compatibilist now has to defend his position. Saying that he does not remember his reasons for defending compatibilism is not a move in the game. Claiming that is of no consequence. The opponent is perfectly entitled to ask the compatibilist to revise his position. The compatibilist might be forced, in this case, to reopen the inquiry. It is noteworthy that he can still prove entitlement to his claim. The fact that one is not able to bring to mind his reasons for p does not mean that one cannot do some research that would help him retrace the steps of his reasoning. The status of being entitled or being able to prove entitlement has no direct connection with memory or memory capacity. Demonstrating entitlement means having the ability to produce good reasons. Whether the reasons come from your memory or from some other source is not important. We can suppose, for instance, that someone has very poor long-term memory, or thinks that his long-term memory is unreliable and writes down his reasoning every time he reaches a significant conclusion. Now, although he cannot retrieve much from his memory, he can always check his notebooks. If the arguments in the notebooks are good then he has entitlement. Similarly, one can reopen the inquiry into the issue of free will and remind himself of the grounds for which he is a compatibilist and succeed in proving his entitlement.

A different possibility is that one does not really have an interest in that issue anymore and does not want to reopen the inquiry. In that case he is not considered to be entitled to his position and breaks a norm of assertion and belief. However, this should not be taken as a troubling result and he himself, as an expression of self-trust, might still hold his belief, but it is a belief for which he lacks support. Analogously, a promiser may realize that keeping his promise requires too much effort and decide to break his promise.

In addition, I think that Harman's claim that we should avoid cluttering our minds with our justifications for the significant conclusions that we reach is more controversial than it may first appear. In fact, an agent who does not keep track of his justifications, at least for his significant beliefs, may end up using his cognitive resources in a very uneconomical way. To adapt a previous example, consider the case of an atheist who lives in a very religious community. As a result of carefully considering the problem of evil, he makes up his mind that God, as a supreme, perfect being, does not exist. I suggest that it would not be wise for him to try to remember his conclusion and forget about the argument which supports it. The conclusion and the argument are equally important because his belief will be constantly challenged. If he forgets his justification and his belief is challenged (let's say, someone comes up with a new argument for the existence of God), he might be forced into reconsidering his conclusion, reconsidering the problem. But this is a waste of cognitive resources. Remembering his argument enables him both to demonstrate his entitlement and to expose the gaps in the reasoning of his interlocutors. For instance, if his opponent argues 'Don't you think that there is a supreme intelligence that is the source of the order in the world?' he might reply 'Maybe, if you think that that intelligence is God, the supreme being, all powerful and all good, then how do you explain that there is evil in the world that is not man-made?'.

This observation also holds true in more concrete situations. Scientific concepts are reflective concepts in the sense that they have a technical meaning which is determined only in the context of a specific theory. As Sperber puts it "these are concepts that scientists can indeed *think with*, in most cases, only by *thinking about them*" (Sperber 1997, p. 77). Typically, scientists discharge the entitlement they have to beliefs involving theoretical concepts by providing arguments or demonstrations. For instance, in Newtonian physics, there is a difference between the weight of an object and its mass. The weight of an object is defined by reference to its mass and the acceleration due to gravity. Thus, in order to justify his belief about an object's weight, a scientist should present a

short demonstration. When using such concepts scientists keep these ways of discharging entitlement in front of their minds.

To sum up, I think the basic premise of Harman's argument is not plausible as it stands and should be qualified. It is very important for people to keep track of their reasoning so that they are not forced to reopen issues that they consider settled and waste cognitive resources. Also, if an agent has a good grasp of the argument which supports his conclusion then he can show that his interlocutors' attempts to challenge it are not warranted; he is in a position to point to the gaps in their reasoning and to the premises they use that he does not agree with. Moreover, in context of scientific inquiry, theorists should have the contexts which validate their beliefs in front of their minds. These demonstrations determine the meanings of the terms they use and helps them work out the implications of their theories and describe the outcome and relevance of specific experiments.

Finally, it is noteworthy that occasional failings to demonstrate entitlement do not lead to losing one's status as an asserter. In the same way, if we do not remember why we held a sentence true, we may continue to do so because we implicitly rely on the responsibility we exercised in accepting claims. We can trust ourselves in the same way we may trust others.

Summary

The goal of this section has been to show that construing the normativity of truth for belief in terms of commitment to truth avoids an important challenge faced by the norm of truth account of belief. The defenders of the idea that belief is constituted by a norm of truth have a hard time showing that this norm is compatible with the phenomenon of doxastic transparency and has genuine prescriptive force. This section demonstrated that construing belief as a cognitive state which essentially involves a commitment to truth can explain the phenomenon of doxastic transparency. In addition, the prescriptions which result

from undertaking such a commitment to the truth of a sentence can be broken and thus, have genuine normative force: one may fail in proving entitlement to those commitments and one may find that his own theoretical and practical reasoning is divorced from his own doxastic commitments. The rest of the chapter will attempt to answer the circularity challenge – the criticism that construing belief as essentially normative leads to circularity and regress.

4. Norms Implicit in Practice

4.1. The regress and circularity objections

In the first chapter it was pointed out that the view that belief is governed by the norm of truth is vulnerable to the objection that it leads to an infinite regress. The crucial premise used in the objection is that rule-following is an intentional activity; that the rule-follower has an explicit intention to bring his behavior in line with what the rule requires. Glüer and Wikforss point out that the rule-follower engages in a bit of practical reasoning in order to determine what the rule requires of him in a certain circumstance. Now, Brandom paints a different picture of rule-following. In his view, it is not a necessary condition of rule-following that an agent has an explicit conception of the rules. An agent can follow norms just by participating in a practice governed by implicit norms. The participants in this kind of practice do not have an explicit knowledge of the rules governing it. The correctness of a bit of behavior is explained in terms of the attitude of treating it in practice as being correct. However, Glüer and Wikforss maintain that going implicit does not solve the problem of regress. In their view, the idea of explaining correct behavior in terms of taking it to be correct leads to a different kind of infinite regress. Treating a behavior as correct is itself subject to normative assessment. One may be wrong in one's practical assessment of correctness. But one may also be wrong in assessing another person's attribution of correctness. And so on. In this section I will describe Brandom's account of

implicit norms in more detail and assess his reply to the regress objection. I will argue that, although Brandom's account has the resources to deal with the regress objection, it is vulnerable to a circularity objection. On his view, we can make sense of the idea that a practice is governed by implicit norms only by interpreting that practice by using our own concepts. However, this leads to circularity because practices governed by implicit norms were supposed to explain what it means to be a concept user and an intentional agent. Then, I will propose a way of avoiding the circularity objection.

The problem of regress briefly described above bears directly on the proposed view that belief is prescriptive because it involves a commitment to truth. In order to avoid circularity and regress, we need to construe the notion of commitment such that it does not involve any intentional notions. For instance, we cannot explain the fact that an agent commits himself to doing something in terms of his decision to do it. Making a decision already presupposes having intentional states. As already indicated, Brandom endorses phenomenalism about normative statuses; he explains the notion of normative status in terms of the attribution of normative status. Commitments and entitlements are normative statuses and they are explained in terms of the attribution of commitment and entitlement. Similarly, the correctness and incorrectness of a performance is explained in terms of attributing correctness and incorrectness. However, in order to avoid the regress, the attributions of deontic status should not be construed as explicitly intentional but as practical attitudes. The basic kind of practical attitude mentioned by Brandom is the disposition to sanction. Thus, for example, someone is committed to do what he promised because someone else has the disposition to sanction him in case he does not keep his promise.

It is essential to emphasize that Brandom's phenomenalism is a normative phenomenalism. He does not want to explain the normative in non-normative terms and he takes norms to be fundamental. Therefore, although Brandom explains attributing a commitment in terms of having a disposition to sanction someone who does not discharge the commitment, he does not imply that we can

ultimately reduce commitments to dispositions to sanction. Rather, his strategy is to explain commitments in terms of correct dispositions to sanction. Let me illustrate why this further condition is needed. Recall that Brandom wants to explain what we need to be able to do in practice in order to count as speaking. In other words, he wants to determine how conceptual norms are instituted. Now, in order for a community to be interpreted as a linguistic community their deontic practice should have a certain structure. But sometimes our actual dispositions to sanction are themselves wrong. For example, someone may be wrongly sanctioned for allegedly not keeping their promise while in fact they kept their promise or they may not be sanctioned when they do not keep their promise. But promising is defined as a speech act which involves a practical commitment. Then, we may not be able to capture the deontic structure which constitutes promising in terms of the actual dispositions of the participants in the practice but rather in terms of the dispositions they ought to have, or their correct dispositions.

Brandom is committed to the possibility that creatures which do not yet have a language or concepts can participate in a practice governed by implicit norms. These pre-linguistic beings can attribute commitments and entitlements to each other by being ready to sanction each other's behavior. Brandom offers various examples. If, for instance, some members of a tribe go hunting and the hunting is successful but they do not bring back any meat, then they may be sanctioned by being beaten with sticks. Also, you might have the permission to enter a hut only by displaying a leaf from a specific kind of tree. However, as indicated above, the practitioners' actual dispositions to sanction may sometimes be wrong. According to Brandom, it is a discursive practitioner who, while assessing the practice, may determine that some practical attitudes of the participants are misplaced. The interpreter uses her own concepts to project from the behavioral dispositions of the participants in the practice and determine what is correct and incorrect according to their practice.

This appeal to an external interpreter also solves the problem of indeterminacy or of the gerrymandering of interpretations. The problem, noticed

by Wittgenstein (1953) and Kripke (1982), is that, at any point in time, there are various ways of interpreting the behavior of the members of a community; interpretation is underdetermined with respect to behavior. To use a well-known example, if the speakers use a certain word in the presence of rabbits, it is natural for us to infer that they mean rabbit by that word. However, it may mean, 'undetached rabbit part' or 'instantiation of the idea of a rabbit' and so on. Brandom, like Davidson (2001b), suggests that we can keep the degree of indeterminacy to a minimum by appealing to the idea that interpreting a foreign language is not essentially different from interpreting someone speaking your own language. In both kinds of interpretations we use our own concepts to make sense of how others act and talk.

Brandom writes:

> But the one who attributes such norms, who takes them to be instituted by the conduct of those taken to be bound by them, must be a discursive practitioner. The interpreter will use her own concepts to project from their dispositions what she takes them to be taking to be permitted or forbidden. And here the fit is bound to be loose. Under a variety of circumstances in which related norms interact with one another the interpreter may take some subset of the actual dispositions to sanction to be mistaken by the hominids' own lights, as she interprets those lights. This need not be made out in terms of full information except in an innocent practical sense that does not require intentionality: the practitioners have committed themselves (according to the interpreter) to deferring to the behavioral verdict of some 'experts' (e.g. the elders) who in fact (according to the interpreter) would be disposed not to sanction in the case in question (Brandom 1997).

It is here that the regress objection is to be launched. It seems that the practical attitudes which were supposed to explain normative statuses are themselves subject to normative assessment by some external interpreter. But, it is also the case that the assessment of the external interpreter of the practice may be wrong.

This appears to lead to an infinite regress. Glüer and Wikforss make the critical point as follows: "According to Brandomian pragmatic phenomenalism, normative statuses are explained by normative assessments. This is a view about the metaphysics of norms. Metaphysically, assessment is prior to normative status. For any particular assessment, a further norm determining the assessment's own normative status therefore already needs to be instituted. Which, of course, can only be done by means of further assessments. And so on. And so a regress ensues: the metaphysics of pragmatic phenomenalism makes it the case that for any norm to be instituted at all, an infinity of prior norms already needs to be instituted" (Glüer and Wikforss 2009).

Gideon Rosen also emphasizes the point that there being some interpreter who evaluates the normative behavior of the practitioners makes it hard to understand how the practitioners themselves institute the norms (Rosen 1997). It seems that the participants in a practice cannot institute the norms unless we know what is correct in that practice. But, given the regress of normative assessments, we may never know what is correct. Then, how can implicit norms be instituted in the first place?

> It is true that for any given normative fact, the regress permits us to cite another fact in virtue of which it obtains. But at no stage is this further fact one that is in any clear sense of our making. The regress provides no insight into how anything we do determines what is correct according to the norms implicit in our practices (Rosen 1997).

Let us evaluate this objection by focusing on an example. Consider a tribe in which adultery is sanctioned by death. Suppose one woman commits adultery and her husband kills her in some ritual way. Now, let us say that one external interpreter of this practice considers the sanction to be correct, given the norms that he takes to be instituted in the practice. However, a second interpreter considers that the sanction was unwarranted. In his view, given that the husband committed adultery first, he should have been killed, and what the woman did

would not have been considered adultery anymore. The second interpreter also argues that if this case had been judged by an expert, the expert would have declared the sanction unwarranted and punished the husband. Now, it seems that both interpreters agree that the participants in the practice instituted the norm that adultery is to be punished by death. Their disagreement makes sense only given their agreement that the practice is governed by implicit norms. Their dispute is sparked by some indeterminacy in the content of the norm. It is not clear, given the observed behavior of the practitioners, whether the norm applies only to women or both to men and women. However, this indeterminacy is perfectly consistent with the fact that the members of the community have instituted one of these norms by their behavior. Thus, it seems that potential disagreement regarding the correctness of sanctioning is not in tension with the fact that there are norms instituted in a society.

Similarly, it is not clear that the regress of assessments must be an infinite regress. Brandom's remarks that interpretation begins at home may be construed as an answer to this challenge. When we interpret an alien community we use our own concepts and are guided by the norms that are implicit in our own practice. However, it is clear that our own linguistic practice is governed by norms, some of which are already made explicit in rules. Norms are surely instituted into our linguistic practices. Also, the fact that there may be disagreement about how to understand the behavior of an alien community, how to fit it into our conceptual scheme, does not imply that interpretation leads to infinite regress. On Brandom's view, the regress stops at home.

Nonetheless, even if we accept that Brandom may be able to reply to the regress objection, it seems that his account is circular. This critical point is developed by Anandi Hattiangadi (Hattiangadi 2007). She observes that, according to Brandom, understanding what it means for norms to be implicit in practice requires conceptual or explicit norms; there is a need for a linguistic being to interpret that practice in her language. But, Hattiangadi wonders, how is this compatible with Brandom's attempt to explain the emergence of conceptual

norms from practices governed by implicit norms? In the following passage Brandom explains why his account of implicit norms may not be circular.

> Preconceptual normativity in one sense accordingly derives from the conceptual norms the interpreter brings to bear. And in this sense, conceptual norms are prior in the order of intelligibility to preconceptual norms. Indeed, in a related though distinct sense, conceptual norms underlie the facts (the sense, elaborated in the concluding chapter, in which the realm of the normative includes the factual, since facts are intelligible only as claimable contents). This sort of priority is compatible with the possibility of a world in which there are preconceptual norms, but no conceptual ones (and indeed, a world in which there are facts, but no normative ones) (Brandom 1997).

In this passage Brandom draws an analogy between the existence of preconceptual normativity and the existence of objective facts. In one sense, he suggests, both implicit norms and objective facts depend on conceptual norms. More specifically, we need to have concepts in order to talk about a pre-linguistic community or make claims about facts. But the fact that we have conceptual abilities, which enable us to talk about norms implicit in practice and objective facts, does not imply that the objective existence of pre-linguistic communities and facts in the world depends on us. Rather, these entities exist independently of us, but we need to be conceptual beings in order to meaningfully talk about them.

However, this is problematic for his account of norms, of how commitments and entitlements occur in the first place and how the practice of assertion gets off the ground, because, ultimately, it seems that we need a discursive practitioner to make sense of the norms instituted in basic, pre-conceptual practices – "conceptual norms are prior in the order of intelligibility". But appealing to the idea of a discursive practitioner makes the account circular because the account is supposed to explain how the ability to participate in linguistic practices emerges from more primitive abilities. However, now it

appears that in order to understand those primitive practices we need to appeal to the idea of an interpreter who is also a language user. But then the things that need to be explained are appealed to in the explanation – the idea of a language-user already involves the idea of a normative practice, of commitments and entitlements. Indeed, as Hattiangadi notices, this order of explanation seems to go against Brandom's manifest strategy because it explains norms implicit in practice in terms of the explicit conceptual norms followed by the interpreter (Hattiangadi 2007). Moreover, the primitive normative practices which were supposed to explain our attributions of intentional states seem to depend on an interpreter who is already construed as an intentional agent. Thus, intentionality seems to be presupposed by the account, and not explained by it.

4.2. A solution: Millikan on Conventions

In order to escape circularity we need an account of norms implicit in practice that does not make reference to a possible discursive interpreter of that practice. We need to explain how certain standards of behavior became adopted in a community. I think that Millikan's account of coordinating natural conventions provides a good starting point (Millikan 2005). Millikan proposes an account of natural conventions which does not rely on intentional notions. She distinguishes between natural conventions and stipulated conventions. In contrast to stipulated conventions, natural conventions are not the outcome of explicit decisions made by the members of a community. On her view, there are two features which characterize a form of behavior based on a natural convention. First, these patterns of behavior are 'reproduced' in the sense that they are based on imitating or copying someone else's behavior. Second, the fact that these forms of behavior proliferate is due, in part, to weight of precedent, not only to the fact that they are good for performing certain functions. This account applies to coordinating conventions. There is a need for coordination when the members of a group have a purpose in common and achieving that goal requires actions by each of the

partners. Furthermore, more than one combination of actions can achieve the purpose. Millikan offers a variety of examples: hand-shaking, keeping a required distance while speaking to someone, the widespread use of certain units of measurement (e.g. money) in a community, waiting for the other person to call when a phone connection is broken and so on. Once a form of behavior solves a coordination problem, it may be reproduced by the members of the community. This initial solution can count as a positive precedent and become stabilized in a given population (Millikan 2005, p. 9-16).

On Millikan's view, speech acts are language forms which were selected because they serve coordination functions. For instance, when a group of people is building a house, some of them may give orders to others to bring certain materials as necessary. Giving an order is the start of a conventional pattern which ends in the order being satisfied. Similarly, a man who is supposed to be on the lookout for potential attackers may signal danger by making a claim. The other part of the conventional pattern is realized when the other members of the tribe accept that they are in danger.

Millikan emphasizes that these conventional patterns of behavior are just standards to which actual behavior may conform or not and they are not intrinsically prescriptive. Usually, they serve their function even if they are performed from time to time, not necessarily regularly. For instance, it may be worthwhile for a beggar to ask each person passing by for money, even if only a small number of them answer his request (Millikan 2005, p. 14). On Millikan's construal the conventional pattern specific of assertion, its biological proper function, is that the speaker makes a claim and the hearer accepts it (Millikan 2005, p. 153). If the speakers prove to be unreliable most of the time then their utterances will cease having any effect and the conventional pattern will be selected out. Nonetheless, Millikan suggests, sometimes these patterns are reinforced by sanctions and they become prescriptive. For instance, although many traffic norms were first coordination conventions, they gradually became written into law, given their importance in preventing accidents. Now, given the

account of assertion developed in this chapter, we should give an argument why the practice of assertion is a prescriptive practice.

We can start by noticing that Millikan's view does not capture the complexity of our ordinary practice of assertion and, in particular, the prescriptive dimension of assertional practice. In particular, it does not account for the promissory aspect of assertion, the fact that the speaker is responsible for the truth of his claim. Millikan's account of assertion does not make reference to the elements of deception and trust which characterize our assertional practice. Sometimes, the hearer does not respond to a claim in the conventional manner because he does not trust the speaker. This mistrust is normally the result of the speaker's past deception or lack of competence. Millikan is right that the sanctions associated with making claims are not aimed at making the hearer believe a message, but she is wrong to infer that there are no such sanctions. The penalties are aimed at eliminating deceptive linguistic behavior. The possibility of deception can be made vivid by reflecting on the structure of assertional conventions. This conventional pattern stabilized because the producers and consumers of assertions had a common purpose and needed to coordinate. For instance, in the context of hunting, the shared purpose is catching the prey. The conventional pattern which defines making claims in this circumstance involves three important elements: the proper context, the message and the acceptance of the message. All these elements should figure in a normal explanation of the purpose of this form of cooperative behavior. For instance, as a result of seeing the prey run in a certain direction, the speaker may inform the others and they will accept his information and act accordingly. The fact that the communicator produces the message in reaction to a specific state of affairs is part of the explanation why this form of behavior got selected. However, once this behavioral pattern becomes stabilized, the stage is set for deceivers to take advantage of it. Someone, for instance, may kill the prey and announce to others that it ran away in a certain direction. He will then have the food to himself. This

kind of cheating is liable to sanctions. The sanctions are aimed at enforcing the connection between the production of the signal and its appropriate context.

Withholding information is a different kind of cheating and can be appropriately punished. Someone may, for instance, withhold information about a good source of food in an attempt to have it all for himself. He may do that while knowing that, standardly, members of the group who discover food should alert the others. Here again, sanctioning is aimed at reinforcing the connection between the appropriate state of affairs and the production of the relevant claim.

Sperber (2006, 2010) argues persuasively that the public practice of asking for reasons for claims and assessing the arguments produced is partly the result of our need to guard against deceivers and misinformation. He writes: "A problem well known to anybody who has ever tried to lie and stick to one's lie over time is that it is increasingly hard to keep it coherent with what is otherwise known to the audience without embellishing it, and it is increasingly difficult to embellish it without introducing internal inconsistencies" (Sperber 2006, p. 186). The capacity to check for the external coherence and internal consistency of a message is a useful defense against deception as well as honest but false information.

The sanctioning behavior may take different forms. In situations when the stakes are high cheaters may be physically punished. However, more common forms of sanctioning are skepticism and criticism. A skeptical reaction implies that the speakers incur the cost of emitting a signal without gaining a benefit. Furthermore, an unreliable asserter may be excluded from group practices which require coordination and lose his right to their benefits. We normally ask someone for reasons for their claims when we cannot accept something simply on trust. When someone's epistemic authority is at a low, they will have to put a lot of effort into proving entitlement to their claims because their interlocutors will not be inclined to simply accept them. As a result, the speaker will have to make an effort to make his claim convincing, normally by connecting it with what his interlocutor already accepts as true. In consequence, the cost of transmitting a

message and persuading his audience will be a lot higher. This fact can be seen as a form of sanctioning[19].

Alongside deception, epistemic incompetence or even momentary lapses in theoretical judgment are also subject to sanctions. In spite of being honest, an epistemic agent may prove to be unreliable as a result of careless reasoning for instance. When scrutinized by his audience, his reasons for making various claims may prove either too weak or irrelevant. In this case he may become subject to criticism. Thus, we can understand the notion of doxastic commitment in terms of epistemic vigilance. Undertaking a doxastic commitment makes the agent vulnerable to questions and criticisms even if the possibility of deception is excluded. In this case, epistemic vigilance will aim at evaluating the quality of the information the agent takes to be true: the questions are aimed at uncovering the reasoning behind the speaker's claims and the criticism is meant as a sanction of the fact that that bit of reasoning is not entitlement-preserving. The criticism is meant to reinforce the inferential patterns which are taken to be entitlement-preserving in the epistemic practice of the community.

To sum up, I think that Millikan's account of natural conventions is an account of norms implicit in practice which avoids the circularity and regress objections because the fact that these norms are instituted is not explained by reference to an external interpreter, but by reference to the biological purpose that form of behavior is meant to achieve. This strategy avoids the indeterminacy problem, the problem of favoring one description of the behavior over others. The conventions adopted in a community are identified by the proper functions they play and these functions are identified by reference to the history of that form of behavior. The reference to the history of a form of behavior is meant to clarify

[19] In this respect, asserting is similar to promising. An agent who does not keep his promises may find himself unable to make promises. He will thus have to put more effort into convincing his interlocutors that he will do what he promised. Therefore, the sanctions associated with promising are not aimed at making the hearer accept a promise, but at obligating the promiser to do what he promised. Thus, the agent who makes false promises or false claims loses her authority and it becomes harder for her to be part of group practices which require coordination and receive a share of their benefits. This kind of exclusion is a form of sanctioning.

why that behavioral pattern got selected and also what biological aim it achieves in normal circumstances. In the way that we know that the proper function of the heart is to pump blood and not to make a certain sound when it beats, we also know that the proper function of a form of behavior is to solve a coordination problem. This explains why that form of behavior proliferated and determines the way in which we should describe that behavior.

Millikan's account also shows how we can construe the notion of doxastic or assertional commitment in non-intentional terms. Brandom defines commitment in terms of the practical attitude of applying a sanction correctly. For instance, an agent is committed to do what he promised because his interlocutor is entitled to sanction him if he does not keep his promise. In this case, we may say that the coordination convention governing a speech act is enforced by sanctions. The sanctioning behavior also has a proper function which can be discerned by reference to the normal explanation of that behavior. In the case of assertion, sanctioning behavior, skepticism and criticism, aims at eliminating deception and lack of competence and reinforcing the connection between the production of the assertion and its normal context. In consequence, we can appeal to the paradigmatic form of the convention and that of the normal punishment in order to determine which sanctions are correct and which are not. Thus, we have an account of the commitment involved in making assertions and promises in terms of a teleological account of natural convention and of enforcing conventional patterns.

5. Prescriptive Norms vs. Norms 'in Force'

Like Millikan, Katrin Glüer and Peter Pagin (1999) develop a non-prescriptive view of language norms or conventions. On their view, the fact that a set of norms is constitutive of a certain practice does not immediately imply that those norms are prescriptive. Rather, Glüer and Pagin argue, constitutive rules have the role of standards of correctness; they help us classify performances into correct and

incorrect ones. In this section I will argue that this picture of norms as descriptive standards of correctness does not apply to our practice of assertion and the way we play the game of giving and asking for reasons.

In the descriptive picture of norms painted by Glüer and Pagin, it is not necessary that the participants in the practice actually intend to perform correctly and follow the rules. There is nothing intrinsically wrong with not obeying an order, not answering a question or not keeping one's promise. The norms which constitute a practice do not have to be prescriptive; they just have to be 'in force'. They write:

> This, we believe, is the key to understanding constitutivity: a practice is constituted by a set of rules if it is possible to engage in that practice only insofar as the rules of that set are in force for the agent. I can make a checkmate or castle only to the extent the rules of chess are in force for me as I make the move. If they are not in force, then I am doing something else. Once this is realized it is obvious that there is no general problem about complying with or violating constitutive rules.
> Accordingly, it is essential, as regards constitutive rules as well as others, to distinguish between being in force and being complied with. This is not to deny that there is a substantial borderline area, where it is difficult to say whether a rule is in force at all because the violations are too frequent. Nonetheless, the distinction remains essential even in such cases (Katrin Glüer and Peter Pagin, 1999).

Glüer's and Pagin's point is that if we construe constitutive rules as rules that are in force then they are not action-guiding. For instance, rules of meaning like ''green' applies correctly to green objects' are meaning-determining but they are not prescriptive. A speaker uses the word 'green' meaningfully if the rule is in force for him, but that does not mean that the speaker cannot make mistakes and misapply the word 'green'. A speaker does not have to follow the rules of meaning in order to speak meaningfully; it is sufficient that the rules are in force

for him. Similarly, a player does not have to follow the rules of a game in order to count as playing the game, it is enough that the rules are in force for him[20].

How does this view apply to the rules of assertion? The claim would be that rules like 'one ought to be entitled to one's claims' are constitutive of assertion but they do not have prescriptive force; one still counts as making an assertion even if one breaks the rules. These norms are just descriptive standards of linguistic behavior. It should be clear that this view of norms applies only to norms which are the result of explicit conventions. The suggestion is that a speaker and his linguistic community might accept the rule of entitlement for assertion as a rule which is in force but qualify their acceptance by saying that the rule does not actually govern their linguistic behavior. However, I think this suggestion is in conflict with our intuition that if someone does not believe that he ought to be entitled to make an assertion he does not understand the concept of assertion. Maybe the picture described by Glüer and Pagin applies to some hypothetical linguistic community, but it is not an accurate depiction of how assertions are used in our linguistic practices.

As it was pointed out, some philosophers emphasize the similarities between the speech act of assertion and that of promising (Brandom 1994, Watson 2004). In the case of promising, we can plausibly argue that someone who claims that the rule that he ought to keep his promises is not action-guiding for him does not understand the notion of promising. Noone would trust the promises made by this speaker and there would be no reason for him to make promises. The whole point of the speech-act of promising is the coordination between speakers, and this coordination is achieved only if it is based on trust and responsibility. By saying that the norm of promising is not action-guiding for him, a speaker undermines his own authority and others would not attribute any actual commitments to his attempts at promising, thus making him unable to make promises. The consumer of promises expects the producer to actually do what they promise. The same point, I maintain, can be made in connection with the

[20] I criticize this view of linguistic meaning in Buleandra (2008).

speech act of assertion: being able to distinguish between correct and incorrect assertion is not enough for understanding the concept of assertion. If someone does not grasp the idea that he ought to be entitled to the assertions he makes then he is not counted as an asserter by the members of his linguistic community. Given that the speaker does not understand the element of responsibility involved in making assertions, his interlocutors are not warranted in relying on the claims he puts forward as true and, in consequence, they do not attribute to him the authority or license necessary for making claims.

My suggestions about the status of being an asserter or a promiser are corroborated by our ordinary linguistic practices. If Glüer's and Pagin's model were correct then someone could hold his status as an asserter or a promiser without ever discharging their commitments. The fact that the rules specific of promising and asserting are in force for a speaker is enough to make his utterances into assertions and promises. However, as Brandom points out, in practice someone can lose their status as an asserter or promiser if they repeatedly fail to discharge their commitments. This is what happened to the boy who cried 'Wolf.' He repeatedly committed himself to the claim that there is a wolf present and therefore he authorized others to draw theoretical and practical conclusions from this claim. However, as it turned out, he was not entitled to commit himself to that claim. In consequence, as a sanction, the members of his linguistic community took away his license to make claims. The point is that, on Glüer's and Pagin's model, given that the rules of assertion are not prescriptive or action-guiding, it is hard to see why the boy was sanctioned. Given that someone does not have the obligation to make correct assertions and that the norms of assertion have only a classificatory use, it follows that a speaker can make as many unwarranted assertions as he wants. But this consequence of their thesis is in conflict with the way we attribute commitments and entitlements to claims in our ordinary linguistic practices.

6. Two Types of Belief

The following objection might be formulated against the account of the notion of commitment put forward in the previous two sections. Although the account shows how commitments emerge from social behaviors which are selected for their success in solving coordination problems, it is still unsuccessful in explaining commitments in non-intentional terms. After all, we have a strong intuition that the behavior of the pre-conceptual creatures in the examples described by Brandom is intentional behavior. Let us take the example of the pre-linguistic community of hunters and gatherers. Is not it obvious that the members of this community want to hunt in order to have food and they know that successful hunting leads to eating meat? Do not they have beliefs that help them navigate their environment and intentions or rudimentary plans that they need to realize?

These strong intuitions generate problems for some central theses of Brandom's theory. Brandom is committed to the idea that the intentionality of non-human animals is derivative and that the home of the attribution of intentional states is in the linguistic practices of rational agents. One reason for the primary role played by the community of rational agents is that the game of giving and asking for reasons is what explains conceptual content. As indicated, Brandom is an advocate of inferentialism about mental content: the view that the content of a concept is determined by the inferential relations in which the propositions containing that concept participate. In addition, Brandom maintains that the home of the concept of reason which is at the heart of intentional explanation of action is the linguistic practice of giving and asking for reasons.

However, the goal of this work is not to defend Brandom's project but rather to use and develop his account of the relation between belief and assertion. A detailed discussion of theories of mental content and intentional explanation is beyond the scope of this project. Our current purpose is to answer the circularity objection: the claim that we cannot make sense of the concept of commitment and

obligation without invoking the intentional state of belief. The line of response developed in the previous sections explained commitments as norms implicit in practices but now we are faced with the plausible observation that the participants in these practices have intentional states.

I agree with these intuitions but I think that the charge of circularity can get off the ground only if it is based on the unity of belief assumption: the assumption that there is only one type of belief. A number of philosophers have recently questioned the unity of belief assumption (Dennett 1991; Frankish 2004; Mercier and Sperber 2009) and offered a two-strand theory of belief. The main idea is that there are two types of belief, corresponding to two levels of cognition. In the following I will sketch the basic differences between these kinds of belief and then I will formulate two arguments in support of drawing this distinction. The lower level belief is the product of specialized cognitive modules which work automatically. We have no control over the output of these mechanisms and we are thus essentially passive with regards to these beliefs. In addition, these kinds of beliefs are not available to consciousness and are not essentially connected with the use of natural language. On the other hand, higher level beliefs are available to consciousness. Usually, they are formed or revised by being involved in conscious reasoning. This is why we do not think that we are passive in relation to these beliefs, that they are just things that happen to us. Rather, it is argued that we are active[21] in relation to them. Moreover, the higher-level beliefs involve language. And, as it was argued in this chapter, reasoning involving these beliefs is mostly internalized speech. As indicated in this chapter, there is an intrinsic connection between the fact that we are active with respect to our conscious beliefs and the fact that they are linked to language use. In doxastic deliberation, we settle for *an answer* to *the question* whether *p*. The claim that we are answerable with respect to our conscious beliefs is explained by the fact that doxastic deliberation takes place in the medium of natural language; that it involves asking yourself a question and settling for an answer.

[21] I will consider more closely the idea that we are active and have control over deliberative belief in Chapter 4.

I think that the distinction between conscious and unconscious beliefs is important and although, apparently, the two kinds of cognitive states play a similar functional role (e.g. they both interact with desires to produce action), they are essentially different and irreducible to each other. First, we cannot construe high-level belief as a kind of low-level belief. Most philosophers agree that an important feature of linguistic belief is that we can ask for reasons for it. Many naturalist philosophers argue that the difference between the intentional states of rational agents and those of non-human animals is just that rational agents are capable of higher-order cognition and they have beliefs about their own beliefs. However, if these higher-order beliefs are essentially of the same kind with the low-level beliefs then it is not clear why we should be responsible for them, why our conscious beliefs are responsive to reasons. This picture explains why we can be conscious of our beliefs by reference to the fact that we have a cognitive mechanism designed to detect and monitor our first-order beliefs, but the fact that we are conscious of them does not make us responsible for them. Similarly, that fact that we are conscious of feeling pain does not make us responsible for the pain. The feeling of pain does not stand in need of reasons (Millar 2004; Boyle 2009).

Second, as it was emphasized, an essential feature of conscious belief is transparency. Transparency shows that only evidential factors can be reasons for belief because believing p is just giving a positive answer to the question whether p. However, as will become clear in chapter 5, we cannot explain transparency if we construe higher-level belief as a kind of lower-level belief. Animal belief, as it might be called, is obviously influenced by non-evidential factors and it is very plausible that our cognitive mechanisms have been designed so that they can sometimes err on the side of caution (in perceiving predators, for instance). Thus, lower-level belief is influenced by non-evidential factors. But then, if rational belief is just a type of lower-level belief, why cannot we make up our minds about something just by reflecting on how useful that belief is for us? Why cannot we

believe that we are very popular and well-liked just by reflecting on how useful that belief would be for us and how much new meaning it would give to our lives?

It might be objected that, on this account, if a belief moves from being unconscious to being conscious then it shifts from being passive and non-prescriptive to being caught up in the space of reasons. This consequence is implausible. However, I do not agree with the idea that the same belief can shift from being conscious to being unconscious. On the view proposed here, unconscious and conscious beliefs are produced by different cognitive systems. Unconscious beliefs are formed by cognitive modules which work at a subpersonal level of cognition whereas conscious beliefs are produced by consciously accepting a statement or making a judgment. Both of these cognitive systems are operating in a normal human being. For instance, suppose that I have to run an errand, find myself in an ill-famed neighborhood and notice the fact that I'm walking faster than usual. Based on the observation of my behaviour I can consciously conclude that I have the unconscious belief that I am in danger. Nonetheless, consciously thinking that I have the unconscious belief that I am in danger does not imply consciously believing that I am in danger.

However, it might be argued that normally our conscious mind takes its input from the deliverances of specialized modules. For instance, we form non-inferential perceptual beliefs based on the output of our specialized sensory mechanisms. Although I agree with this observation I do not think this connection between the two systems implies that unconscious beliefs become conscious. Rather, we should say that based on the information coming from our perceptual mechanisms we form conscious beliefs with the same content. This distinction is important because it allows us to explain the fact that sometimes, when our subpersonal perceptual mechanisms deliver a certain input, we may decide not to accept it. For instance, in the case of the Müller-Lyer illusion, our perception tells us that one line is longer than the other, but we do not form the corresponding conscious belief because we know that they are the same length. Similarly, suppose that someone watches a horror movie in 3D. In one of the scenes a

monster emerges from a mass of green slime and comes towards the viewers. It is possible that the person watching will scream in terror while clutching desperately at their chair. This behavior might be interpreted as the activation of a fight or flight automatic response. However, at the same time, the subject consciously knows that he is in a theater and is not in physical peril[22].

With this distinction between the two types of belief at hand, the emerging account of the normativity of linguistic, conscious belief has the following structure. There are two levels of belief, and more generally, intentionality. We share the first strand of beliefs with non-human animals. Guided and moved by lower level beliefs and desires we participate in social practices that benefit each other. Some of these kinds of behavior are reproduced and flourish. Linguistic behavior is one of these kinds of conventional behavior that proved fruitful in solving a wide range of coordination problems. The main purpose of the practice of assertion is transmitting information; when as asserter makes a claim he is committed to the truth of the sentence uttered. We can decide whether someone is a reliable, trustworthy informant by asking for reasons for his assertions and evaluate the quality of those reasons. By offering good reasons for his assertion the speaker shows that he was entitled in undertaking a commitment to the truth of the respective sentence. Similarly, when we deliberate about whether *p*, we consider various reasons for and against *p*. If we are convinced by the reasons we find in support of *p* then we accept *p* or we judge that *p* is the case. Such a judgment results in the higher-level belief that *p*. If the first-strand belief is the product of specialized cognitive mechanisms, second-strand belief is a product of participating in the practice of assertion. The connection between higher-level belief and the practice of assertion explains why we ask for reasons for belief and why belief is transparent.

[22] This example is used by Tamar Szabó Gendler (2008), to illustrate the notion of *alief*. The notion of *alief* is related with the notion of unconscious belief described in this section. Gendler writes: "I argue for the importance of recognizing the existence of *a*lief – so-called because *a*lief is *a*ssociative, *a*utomatic, and *a*rational. As a class, aliefs are states that we share with non-human animals; they are developmentally and conceptually antecedent to other cognitive attitudes that the creature may go one to develop. And they are typically also *a*ffect-laden and *a*ction-generating." (Gendler 2008, p. 641).

It is noteworthy that the concept of assertional commitment to truth is not to be accounted for by reference to the notions of inference and argumentation, but rather by appealing to the notion of correct sanctioning. It is plausible that speakers were able to attribute and undertake commitments before the practice of giving and asking for reasons was instituted. We can clearly imagine the rudimentary practice of assertion before this filter was invented. In this simpler language-game there was no argumentation, only testimony. However, even in the context of this basic game, there are still speakers who are deceptive or lack competence and they may be sanctioned. But these agents' beliefs are not higher-level beliefs because they are not answerable for them: they do not have to give reasons for them. Consequently, they are incapable of linguistic deliberation and of forming active, responsible beliefs. Thus, their beliefs are lower-level beliefs.

According to Gricean theories of communication even participants in simple linguistic practices should have beliefs about their interlocutors' communicative intentions and beliefs. Also, the possibility of deception involves the capacity to attribute intentional states and to predict and manipulate other people's behavior by influencing their cognitive states. However, although I agree that the ability to engage successfully in communication requires the capacity for mindreading, I don't think that these second-order cognitive states are conscious, higher-order states in the relevant sense. These states lack the central feature of deliberative beliefs: they are not active; the bearers of these states are not answerable for them. The capacity to offer reasons for claims and to evaluate arguments is a late addition to the assertion game, an additional tool for persuasion and calibrating trust. Then, if we can make perfect sense of the capacity for intentional deception in the context of a rudimentary assertional practice which does not involve argumentation, it follows that deception does not require the ability to form deliberative beliefs.

This point can be made more clearly by reference to the simple cooperative practices non-human animals engage in. Obviously, these cognitive systems have only lower-level beliefs. It is reasonable to argue that assertion has

evolved from more simple forms of signaling used by non-human animals. Alarm calls and food calls are paradigmatic examples. It is noteworthy that although the animal behavioral patterns cannot be classified as conventional in Millikan's sense, it may still be argued that they serve specific proper functions. Like bodily organs, these forms of behavior got selected because of their adaptive value (Lorenz 1966). Thus, we are warranted in focusing on a specific description of these behaviors – the one which makes prominent their survival value – without appealing to a linguistic interpreter.

If the notion of commitment is defined in terms of correct dispositions to sanction, then we can legitimately talk about commitments in connection with these forms of animal behavior. Empirical studies show that rhesus monkeys, for instance, use both active and passive deception. Individuals who discover food but do not make the required food call are usually subject to aggression when detected by other members of the group (Hauser 1997). In addition, rhesus monkeys and other animals are able to meet some signals with skepticism. For instance, if one member of the group proves to be an unreliable signal caller, then her signals stop producing an effect on the audience. Interestingly, rhesus monkeys show that they trust some individuals more than others. For instance, if a high-ranking female makes the alarm call then it is more probable that the other members of the group will run and hide, even if that caller made false calls in the past. However, the adults in the group are less responsive to calls coming from low-ranking group-members. This difference "may be the consequence of the adults' more extensive interaction with the group members that are generally reliable signalers" (Gouzales et al. 1996).

Hauser (1997, p. 586-594) discusses the possibility that the deceptive behavior in animals is not intentional. Some experiments suggest that physiological changes may largely dictate the conditions for call production and suppression. More specifically, changes in cortisol level are associated with changes in stress and the production of alarm calls. However, if the cortisol level remains low even in a stressful situation then the animal will not produce an

alarm-call. Thus, rather than being an intentional act of deception, withholding an alarm call might be determined by low cortisol levels in the organism. In addition, experiments regarding monkeys' capacity for mindreading are inconclusive and throw into doubt the idea that they may be capable of intentional deception.

Nonetheless, it is important to point out that these empirical debates are relevant to the argument developed in this section only to a small degree. My point is mainly conceptual. The fact is that we can easily conceive that the deceptive behavior of a group of non-human animals has such a level of complexity that it would be impossible for us not to regard it as intentional. Moreover, the members of this group may display a robust capacity for mindreading. Since they are capable of genuine deception, they can also invent the tools to guard against it and devise a system of sanctions. My suggestion is that this is all they need to be able to do in order for their communicational practice to be a normative practice; for them to undertake and attribute commitments.

7. Conclusion

In this chapter I have argued in favor of the idea that beliefs necessarily involve a commitment to truth and that this commitment is prescriptive: an agent who undertakes such a commitment ought to demonstrate entitlement to that commitment, if appropriately challenged. I suggested that this account avoids the major problems faced by the norm of truth view of belief: it explains doxastic transparency, it involves genuine prescriptions (norms which can be broken) and it is not circular. In connection with the latter issue, I followed Brandom's strategy of explaining commitments in terms of the practical attitude of attributing commitments and understanding practical attitudes in terms of correct sanctioning. However, Brandom's suggestion that we can understand appropriate sanctioning by reference to a discursive interpreter has been rejected. In turn, I sketched an account of correct sanctioning based on Millikan's biological account

of conventions. In this fashion, the way has been paved for an account of the notion of commitment which does not make reference to explicit intentional notions. In consequence, the definition of conscious belief in terms of undertaking a commitment to truth is not a circular definition.

Chapter III - Strange Beliefs

1. Introduction

In Chapter 2 I have defended the view that we decide whether someone succeeds in proving entitlement to a claim by reference to the epistemic standards which are in force in our community. However, the idea that there is only one epistemic community and one relevant practice is simplifying and idealized. It appears that there are many epistemic practices with different epistemic norms in place. This becomes clear if we look at the history of science or across different cultures. For example, before the scientific revolution in the seventeenth-century, the Bible was considered relevant in answering questions about the physical nature of the universe. That is, appealing to the authority of the Bible was a form of demonstrating entitlement to certain views – that the earth is the centre of the planetary system, that it has a certain age, that all species of animals were created by God, and so on. Similarly, in African tribes like the Azande or the Kalabari, one can appeal to oracles in order to find out the causes of a certain disease and to find a treatment for it. Thus, appealing to the authority of the diviner provides a way of discharging their doxastic commitments. These observations provide two challenges to the account of belief and assertion developed in the previous chapter. First, it needs to be shown whether, according the view proposed here, participants in different epistemic practices have epistemic entitlement or not. Second, and more generally, cases of epistemic practices which are widely dissimilar from our own rise the worry that the account of belief outlined in this work is ethnocentric; that it only an expression of the values *we* attribute to certain epistemic virtues and practices. Thus, the objection goes, given that other cultures do not share into these values, it seems to follow that they do not count as having beliefs in the relevant sense – a very implausible conclusion.

2. Significant Cases

(a) Bellarmine vs. Galileo

According to the Ptolemaic view, the earth is at the centre of the universe and the planets and other celestial bodies revolve around it. This picture fits some passages in the Bible, which, being the revealed word of God, was usually interpreted literally. In 1543 Copernicus proposed a different theory which aimed at better explaining various astronomical observations. On this account, the earth revolves around the sun, other planets resemble earth, Venus exhibits phases and so on. In early seventeenth century, Galileo was able to use a new invention, the telescope, and verify the hypotheses advanced by Copernicus. Now, given that the new discoveries were in tension with a literal established interpretation of the Bible, Galileo was summoned to Rome to defend his views. Cardinal Bellarmine was in charge of deciding whether Galileo's views were heretical. Although Bellarmine looked through the telescope, he dismissed the evidence and reasserted the authority of the Bible.

Now, the question before us is: who has entitlement? Can we say that both Bellarmine and Galileo are entitled to their conflicting views? It appears that Bellarmine's deference to the authority of the Bible is an acceptable way of demonstrating entitlement in his respective epistemic practice. On the other hand, Galileo's stance seems more in line with what will become the dominating, scientific paradigm.

(b) Oracles

The Azande do not believe in bad luck, accidents or chance. In their view, every unfortunate event is a result of witchcraft. To take an example made famous by Evans-Pritchard (1937), when a granary falls on someone sheltering under it the Azande immediately infer that the injured person was the victim of witchcraft.

Similarly, diseases, bad crops, and other calamities are explained by reference to magic and witchcraft.

In their view, witches are usually male members of the tribe who have a physical substance in their belly, which lives a fairly autonomous life. Very often, the witch is not conscious that this substance is active and affects the lives of those around him.

When someone's family member has a disease which cannot be cured by usual means, they immediately suspect witchcraft, list the names of people who might be responsible and go to an oracle. The oracle determines whether the disease was caused by bad magic, who is responsible for it, and prescribes ways of protecting the victim. One of the ways in which the diviner comes to this knowledge is by administering poison to a chicken and observing whether it lives or dies. For instance, the diviner might ask whether the sick person has been bewitched or not. If the chicken dies, let's say, the answer is yes, if it lives the answer is no. Then the victim's family member contacts a witch-doctor who is able to perform rituals which will protect her. Also, if one of the community members is identified as a witch then a formal apology ceremony must take place. The individual responsible should participate in rituals aimed at stopping the activity of his witchcraft.

In general, these practices aim at regulating the conduct of the members of the tribe and maintain social order. In some cases, the fact that an individual falls suddenly ill is interpreted as a punishment or sanction of the spirits. In his studies about the Kalabari, Robin Horton (1967) remarks that in their view disease may be caused by the anger of several categories of spirits: heroes, ancestors, water-spirits, and medicine spirits. Heroes, for instance, are spirits which protect the community and they are activated by offences against the laws or customs of the tribe. Ancestors are spirits which protect lineages and bring misfortune to those who betray lineage values. Water-spirits are the patrons of human individualism and can take both creative and destructive forms. Finally, medicine-spirits are activated by someone's enemies when they are plotting against him.

Horton observes that this belief-system allows for many alternative diagnoses of a specific disease. This is part of the reason why the victim's family should consult a diviner. Given his special ability to come into contact with the world of spirits and interpret their signs, the diviner can identify the specific cause of a certain disease. If the remedial prescriptions suggested by him do not work, the failure is usually explained in two ways. First, the procedures used by the diviner are fallible and easily disrupted. For instance, they can be easily affected by those who have a grudge against the client. Second, failure may also be explained by reference to the oracle's dishonesty or incompetence.

One important point emphasized by Horton is that in all African traditional cultures there is reluctance to keep track of the failures of prediction and to act by attacking the beliefs involved. Rather than revising their beliefs, the members of these cultures are quick to come up with *ad hoc* excuses, technically called "secondary elaborations", to explain why the advice of the diviner did not work in a certain case. Horton argues that these traditional cultures are 'closed' epistemic systems in the sense that they do not have the notion of alternative systems of belief. Any challenge to their belief-system triggers deep anxiety because these beliefs also constitute their identity.

Another important feature of 'closed' epistemic systems is that their members are not interested in formulating general epistemic principles. Some anthropologists went as far as claiming that the Azande, for instance, do not follow the norms of deductive logic (Evans-Pritchard 1937; Winch 1964). For example, Evans-Pritchard has noticed that their view of witchcraft entails an obvious contradiction. The Azande hold that the witchcraft substance is inherited patrilineally. It follows that if a man is identified as a witch then all his male descendents are witches. However, the Azande reject this inference:

> To our minds it appears evident that if a man is proven to be a witch the whole of his clan are ipso-facto witches, since the Zande clan is a group of persons related biologically to one another through the male line. Azande see the sense of this argument but they do not accept its

> conclusions, and it would involve the whole notion of witchcraft in contradiction were they to do so (Evans-Pritchard 1937, p. 24).

It appears that the Azande reject the logical conclusion that the whole of a witch's clan are witches and consider only the paternal kinsmen of a known witch to be witches. Their reasoning can be interpreted either as a rejection of the principle of non-contradiction or a rejection of *modus ponens*. For instance, they seem to reject the inference from '*If* Abu is a witch *then* the whole of his clan are witches' and 'Abu is a witch' to the conclusion that 'The whole of Abu's clan are witches'.

3. Objectivity and the Structure of Entitlement

We should start by arguing that neither the European society before the scientific revolution nor the epistemic practices of African tribes count as alternative epistemic systems to our own. There are core epistemic norms of entitlement which are shared by participants in all of these practices. In Chapter 2, three ways of demonstrating entitlement were distinguished: appealing to one's authority as a reliable observer, appealing to the authority of another asserter and formulating another claim in justification for the initial one. I think that the first two methods of demonstrating entitlement are widespread across different cultures and different time periods. First, claims which are the result of induction or deduction can usually be challenged by reference to observation. This can be explained by reference to the survival value of the deliverances of our perceptual mechanisms. Even if an agent has a strong expectation, as a result of induction from past experience, that a certain area of the forest is devoid of predators, he will automatically form the belief that there is a dangerous animal in front of him once he sees it. Accordingly, in public conversation, if someone argues, based on induction or deduction that *p*, his interlocutor may challenge him by saying that he observed that not-*p*. If the interlocutor is an honest and reliable observer then *p* will normally be rejected or qualified. Moreover, as indicated in the previous

chapter, a degree of trust in the testimony of others is a necessary condition for the evolution of communication since communication evolved as part of cooperative activities and its main function is the transmission of information. If all claims were met with skepticism then the practice of making claims would be selected out.

In addition, some forms of justification also have a cross-cultural status. Induction from past experience is a very useful form of reasoning and it is used in different cultures. Any sort of planning is based on expectations or predictions and these expectations are based on induction. Some forms of deduction were widespread and codified in our Western culture since Aristotle and existed in other cultures sometimes in the form of mathematical reasoning. Even our understanding of the Azande and other traditional African cultures presupposes that they are able to follow *some* norms of deduction. Horton, for instance, argues that the Azande views about witchcraft have the structure of a theory which is applied to particular cases. To illustrate, if a member of the tribe is badly injured as a result of an unfortunate accident then the Azande will apply their general knowledge of witchcraft to that specific case of bad magic. They will *deduce* that someone's witchcraft substance is active and will *induce* the identity of the responsible witch.

Now, according to the view of assertion and belief proposed here, entitlement to claims has a *default and challenge structure*. In other words, one has entitlement up until the moment when someone else articulates a legitimate challenge to one's truth-claim. I think that, given that some basic epistemic norms are shared among members of different cultures or theoretical paradigms, it might be argued that members of one culture or paradigm can formulate legitimate challenges to members of different cultures. A legitimate challenge means a challenge which has force given the epistemic practice of the community the accused is a part of.

This idea can be applied to the case of the disagreement between Bellarmine and Galileo as follows. First, Bellarmine should recognize that

observation normally has priority over the teachings of the Bible. In other words, the Bible has epistemic authority because we can discern in reality signs and patterns which fit a literal interpretation of the text. The reliability of the Holy Scripture is established by reference to observed facts. In addition, historically, there were cases in which a literal interpretation of the biblical text has come under pressure and a figurative interpretation was chosen instead[23]. For instance, Psalm 104 was very often cited in support of geocentrism: "Oh Lord my God … who laid the foundations of the earth, that it should not be removed for ever". Although this was taken to imply that the earth does not move, it was not interpreted as meaning that the earth literally has a foundation or that the earth is flat. The Bible scholars were aware that the earth is suspended in space and spherical and, thus, inclined to offer a non-literal interpretation of the expression 'foundations of the earth'.

Now, the telescope is an instrument of observation. Therefore, it might be argued that Bellarmine cannot dismiss it as a source of evidence, based on the standards implicit in his own epistemic practice. In reply to the suggestion that revelation applies to statements about the heavens and observation to claims about immediate facts, Boghossian makes a similar point:

> But this would only make sense if he [Bellarmine] also believed that propositions about the heavens are different in kind from propositions about earthly matters, so that vision might be thought to be an inappropriate means for fixing beliefs about them. But does not he use his eyes to note that the sun is shining, or that the moon is half-full, or that the clear night-time Roman sky is littered with stars? And does not he think that heavens are in a physical space that is above us, only some distance away? If all this is true, how could he think that observation is

[23] This possibility is mentioned by Bellarmine himself in a letter. He writes: "If there were a real proof that the Sun is in the center of the universe, that the Earth is in the third heaven, and that the Sun does not go around the Earth but the Earth around the Sun, then we should have to proceed with great circumspection in explaining the passages in the Scripture which appear to teach to the contrary, and rather admit that we did not understand them than declare an opinion to be false which is proved to be true" (de Santillana, p. 99-100).

> not relevant to what we should believe about the heavens, given that he relies on it in everyday life? (Boghossian 2006, p. 104).

However, Bellarmine has ways of answering Galileo's challenge and save the authority of the Bible. First, the Cardinal can challenge the reliability of the new instrument. Thus, given that all sides agree that observation is relevant, it is up to Bellarmine to prove that the telescope is not accurate. This way of looking at the disagreement is not an example of ethnocentrism because the standards used in our assessment are shared by both arguers.

Bellarmine can also embrace an alternative geocentric account which accounts for the observations, like the planetary system proposed by Tycho Brache. On Tycho's view, the earth is motionless, the sun orbits the earth and all other planets orbit the sun. However, these new versions of geocentrism can be challenged given that there are alternative explanations which are more simple and elegant. These are standards for adjudication between alternative theories which are slowly emerging.

In the case of the Azande, the pronouncements of the oracle can also be challenged by appeal to observation. If the diviner offers diagnoses and prescriptions which prove to be ineffective then his authority decreases. He is considered either dishonest or incompetent. However, one puzzling feature of traditional African culture is that they do not regard these failures of prediction as evidence against the theory. They explain them away by secondary elaboration. Now, the question is: are they entitled to this sort of reaction to evidence?

Before approaching the question of entitlement it is relevant to note that the view of disease which is the core of their belief-system is not without cognitive value. Horton points out that the views of disease of the members of traditional African culture have the structure of a theory. They postulate a number of entities which explain the irregularity and diversity of experience. On his view, spirits can be interpreted as theoretical entities analogous to molecules, atoms or electrons in Western science. Furthermore, the relation between disease and disturbances in one's social life is real. Disturbances in one's social life often lead

to mental stress, which, in turn, can cause a multitude of physical ailments. So it is not the case that their view of disease has no explanatory and predictive power.

Now, we can either argue that the members of these cultures are not entitled to hold on to their theories at any price or that they have epistemic entitlement. There are difficulties with both these options. Claiming that they do not have entitlement may be seen as an expression of arrogant ethnocentrism. We should recall that construing *ad hoc* explanations for recalcitrant data is not always forbidden in normal scientific practice. Scientists usually stop making additions to a theory once they have an alternative theory which better explains the data. In other words, the epistemic standards which prescribe when theoreticians should stop trying to save a theory are established between the participants in an 'open' epistemic practice, a practice in which a certain theory does not have absolute validity and alternatives to it are formulated and assessed objectively. But these standards came to define post-Galilean science and cannot be applied to a 'closed' epistemic system like that of traditional African cultures. For the members of these cultures, the views that they inherited from their ancestors have absolute value and they cannot conceive rejecting them in light of recalcitrant experience. Thus, every challenge to their entitlement to those beliefs will be answered by secondary elaboration. This is also an effect of the fact that they do not have alternative explanations of the relevant phenomena.

On the other hand, one difficulty associated with arguing that they have epistemic entitlement because they cannot imagine that their views are wrong is that then we may be forced to grant entitlement to every agent, even members of our own culture, who cannot even conceive that they are wrong. For instance, there are still people who use the Bible to determine the age of the Earth and deduce that it must be around three or four thousands years old. They regard with anxiety every attack on their belief-system and are reluctant to revise any of their beliefs. However, I do not think that they should be considered entitled to their beliefs and to ignoring all the evidence to the contrary, just because they cannot cope with it emotionally. The difference between these members of our own

culture and the Azande is that they live in an 'open' epistemic system and an alternative theory which manages to better explain all the data is fully developed within their culture. In conclusion, I maintain that the Azande have epistemic entitlement to their beliefs about witchcraft by default (since we cannot construe a legitimate challenge of their use of secondary elaboration), but add the qualification that we cannot generalize this conclusion to epistemic agents who operate in 'open' epistemic systems.

Let us now turn to the claim that the Azande do not follow some basic norms of deductive reasoning like *modus ponens* or the principle of non-contradiction. First, it should be pointed out that, if communication evolved as a tool for the transmission of information, language users should have some minimum deductive abilities. If a speaker usually makes claims which have the form of explicit contradictions then he will soon lose his status as a speaker. Being an asserter implies being able to acknowledge that the claim that *p* is incompatible with the claim that not-*p*. In the case of the Azande, given the importance of identifying whether someone is a witch, it is not at all plausible to argue that, in many cases they believe both that someone is a witch and that they are not. Similarly, our understanding of their culture is made possible by the fact that they follow some basic deductive norms. For instance, anthropologists surely attribute to them inferences of the form:

(1) If a person is hurt in an accident then that person has been bewitched
(2) Abu was hurt in an accident
(3) Therefore, Abu has been bewitched.

It looks that their entire practice regarding witchcraft is based on a multiple uses of *modus ponens*, specific cases being explained by reference to a general principle. Thus, although it is plausible that the Azande, like members of other cultures, are prone to some mistakes in reasoning, it is not intuitive to claim that, in practice, they reject some basic rules of deductive reasoning, since our

understanding of their practice is based on the fact that they are capable of following some basic logical norms.

How about their claims about how the witchcraft substance is inherited? Are they inconsistent? I think there are more charitable interpretations of their view. For instance, one plausible line is that they reject the conclusion that all the members of a clan are witches because they detect an equivocation in the use of the term 'witch'. When an oracle says that someone is a witch that means that their substance is active and has dangerous effects. In this sense, it is not the case that someone's descendent is also a witch because it does not follow that their witch substance is also active. Their substance might be 'cool' (Bloor 1991, p. 141).

Furthermore, instead of jumping to the implausible conclusion that the Azande do not follow basic principles of logic, we can also consider that our translation of the words or phrases in their language may be deficient (Boghossian 2006, p. 106). Maybe, 'transmitted patrilineally' is not the right translation of their claims about the inheritance of witchcraft. Or, the word 'if' is not the right translation of the sentential connective they are using. This last point is supported by the idea that logical inferences like *modus ponens* are used in determining the meaning of the truth-functional connective 'if' and the fact that the Azande reject some of the inferences of this form might imply that they use the sentential connective to express a different concept.

To sum up, since some basic epistemic norms are shared by different cultures, some of the ways of proving entitlement which are widespread in a culture can be legitimately challenged. Against the background of these shared forms of argumentation, we can make judgments regarding entitlement which are not expressions of ethnocentrism and have a claim to objectivity. However, some cultures are more concerned with making explicit the epistemic norms which shape their practice and with assessing their validity. For instance, in the Western tradition, since Aristotle, philosophers became interested in formulating and evaluating the rules behind certain types of deductive and inductive reasoning.

Furthermore, in the post-Galilean science philosophers and scientists were concerned with differentiating scientific knowledge form other types of knowledge and, later on, formulated and debated norms which should govern a choice between two competing scientific theories. This progress was made possible by the fact that the Western culture has always been, more or less, an 'open' epistemic system, one in which people with different views were able to engage in debates and the notion of an alternative view was not completely alien. However, it would be a mistake to judge participants in 'closed' or unreflective epistemic practices by reference to epistemic norms which emerge in 'open' practices. Members of traditional African tribes are not interested in codifying the norms which shape their epistemic practices and they do not even have the concept of a norm of reasoning. This is a consequence of the fact that they were not faced with alternative views and consider the beliefs learned from their elders to have absolute validity.

4. Strange Beliefs and 'the Aim of Truth'

Reflection on the cases of religious or cultural belief raises again the issue of the relation between belief and truth. It might be argued that the idea that belief involves a commitment to truth or following a norm of truth is a philosopher's construct and does not really apply to concrete cases. It seems rather that people go to great lengths to avoid the truth. In this section I will develop two connected arguments in reply to this charge.

First, the fact that belief is defined by its relation to truth means that to believe *p* is to hold *p* to be true. People who have religious beliefs hold onto them because they think that they represent how things are. Scholars in the Middle Ages actually believed that the Bible is the revealed word of God. These beliefs were accompanied by deep-seated ideas about the world or about the nature of reality. On their view, reality had to be perfectly ordered. This is one reason why arguments based on correspondences or finding analogies or similarities between

different kinds of things were considered particularly powerful. For instance, from the fact that there are seven 'windows' to the 'domicile of the head' (two nostrils, two eyes, two ears and a mouth) and the premise that the microcosm corresponds to the macrocosm, theologians were able to infer that there are necessarily seven planets in the heavens (Warhaft 1965). From our perspective, these arguments seem weak and ludicrous and we are tempted to say that they are means for the propagation of false beliefs. However, this reaction is ethnocentric. Those forms of reasoning made sense to their minds because of the deep presuppositions they shared regarding the nature of reality – that the universe is perfectly ordered (Taylor 1982). In addition, these cognitive states should be regarded as beliefs because they are sensitive to how things are. For instance, if it was determined by analogical reasoning that the number of planets is seven but it turns out to be nine then that belief is revised in light of how things are. But this does not immediately involve scrapping the entire theoretical framework or style of argumentation which produced that belief. Maybe our mind did not grasp the right pattern of similarities, but that pattern is still there in reality, waiting to be discovered. Similarly, the Azande do not keep on believing what the oracle has told them when it is obvious that the remedy did not work. But this diminished trust in the honesty or competency of the diviner does not immediately give rise to skepticism regarding the existence of witchcraft. So the fact that these cognitive states are sensitive to how things turn out in the world shows that they are beliefs and they 'aim at truth'.

The second point that needs to be stressed is that although belief is a state which 'aims at truth', it can also be influenced by pragmatic factors like emotions, hope and so on. In the first chapter I argued, following Shah and Velleman, that these pragmatic factors can influence doxastic deliberation only at an unconscious level. To use a standard example, we cannot just decide to believe in God just based on our need for meaning and purpose. Rather, our desire and hope that life has a final meaning is expressed in our reluctance to give up the belief in God in light of evidence. This is, I think, what accounts, on a larger scale, for the

dogmatism of religious belief. Religious beliefs are and were so central to our lives that we would do anything but revise them.

In addition to the desires for meaning and order, religious or cultural beliefs and assertions may be accepted and propagated non-critically because they help define group-identities. As Sperber puts it: "It may be that the content of the ideas matters more to you than who you share them with" (Sperber et. al. 2010). In most of these cases people accept claims which do not have a clear meaning for them: "Mary was and remained a virgin when she gave birth" or "The Father, the Son and the Holy Spirit are one". Although they do not fully grasp these claims and they do not cohere with their other beliefs, people do not challenge the authority of the source. Religious leaders, gurus and some philosophers achieve such inflated reputations that people have a tendency to defer to them all the time and grant them the status of epistemic authorities. However, these reputations, in turn, are the result of lack of epistemic vigilance and a tendency to accept and transmit an idea just because it is widely accepted and transmitted.

The question which arises immediately is whether these cultural beliefs with indeterminate content are really beliefs. Theorists like Dan Sperber argue that they are not genuine beliefs because their transmission does not serve an epistemic function. Sperber draws a difference between factual beliefs and representational beliefs. Factual beliefs are those we take to constitute knowledge and represent established facts. Representational beliefs are beliefs taken together with the context which validates them. For instance, the belief that my parents have told me that Santa Claus will give me presents and they are truthful is a representational belief. The belief that Santa Claus will give me presents is a factual belief. On Sperber's view, there are two kinds of representational beliefs: propositional representational beliefs and semi-propositional ones. In the first case, the agent has a good grasp of the content of the belief but in the second case he has only an incomplete understanding of the concepts which form the belief. For instance, one does not take statements about the Holy Trinity as facts but one takes as a fact the metarepresentation that these statements are accepted on good

authority and a proper interpretation of them must be true. Sperber argues that representational beliefs of semi-propositional contents should not be conflated with genuine, factual beliefs and that the norms of consistency and coherence should be relaxed in their case because that the subjects do not have a grasp of their contents. They do not know what the statements imply and what is incompatible with their truth.

However, I think Sperber's proposal is in tension with the fact that agents take these propositions to be true and are prepared to either defend them by providing justifications or defer to the source. Both of these moves are ways of proving entitlement to holding a claim true. For instance, one standard way of defending the claim about the Holy Trinity is by analogy with water. Many religious people are quick to say that water itself can take different forms and still be the same substance: it can be liquid, solid (ice) and a gas (vapor). Furthermore, when agents defer to epistemic authorities then we should scrutinize the claims and explanations of those authorities. Many religious leaders, for instance, have tried to vindicate claims made in the Bible, offering arguments for them and interpretations which seem more plausible in light of the facts. For instance, in order to make Genesis consistent with the idea that the age of the earth is much greater than 10.000 years theologians in the nineteenth century advanced the thesis that one day of creation represents millions of years in real time. If they did not believe these claims why would they put so much effort into providing arguments for them and making them consistent with scientific knowledge? It is clear, based on their practice, that these assertions are caught up in the game of giving and asking for reasons and that they feel that they should provide a plausible interpretation of these statements to ensure that people will continue to accept them.

Furthermore, I think Sperber's suggestion has the implausible consequence that there are types of belief which should be exempt from critical discussion just because the respective agents do not have a clear grasp of what they believe. But, in everyday argumentation, we feel that our interlocutors have

an obligation to explain and clarify the vague terms that occur in their statements. We often react to their arguments by saying 'Well, that depends on what you mean by…' Noticing that a term or statement in an argument is too vague is a standard form of challenging the argument and the arguer feels the obligation to clarify it and make his statements more precise. If we are able to prove that an assertion does not make any sense we expect the assenter to withdraw that claim and reconsider it.

I think that, rather than claiming that religious or cultural beliefs are not genuine beliefs, it is more plausible to argue that they are beliefs heavily influenced by pragmatic factors. This influence is discernable in the way people accept them uncritically, are reluctant to revise them in light of experience, and in the weak arguments they offer in support of them. It just seems that their epistemic vigilance is at a minimum because they need to accept these truths.

Nonetheless, there are cases when agents make these sorts of claims but argue that even if we do not understand them we have to have faith in their truth. The religious authorities who propagate these ideas say that understanding them is beyond that capacity of human minds and that they are mysteries for us, but we have to have faith in God and accept his word. In these cases, I think that we are dealing with acts of faith rather than acts of judgment. Faith is holding a proposition to be true without reason or entitlement or by explicitly rejecting the idea that that proposition can be vindicated in a standard, rational way. In this context, the proposition is not true in the sense that it represents something in reality but rather it is true of something which transcends reality and should be accepted on faith.

To sum up, religious or cultural beliefs are beliefs which 'aim at truth' to the extent that they are caught up in the game of giving and asking for reasons. The fact that agents usually offer poor reasons in support of these beliefs and are very attached to them is to be explained by the fact that these judgments are influenced by non-evidential factors. If they are statements accepted on faith then they are not genuine beliefs but rather expressions of faith and hope.

5. Conclusion

This chapter aimed at clarifying the view of conscious belief articulated in Chapter 2 and answering the objection that the account put forward is ethnocentric, in the sense that it links the concept of belief with standards of rationality which are specific to our culture. I argued that different cultures share some ways of proving entitlement to doxastic commitments. Against this common background we can make sense of the notion of a legitimate challenge to specific epistemic norms accepted in other cultures. Thus, we are sometimes entitled in making cross-cultural judgments regarding epistemic entitlement without expressing ethnocentric arrogance. Furthermore, I argued that the acceptance of many religious or cultural beliefs is heavily influenced by non-epistemic reasons but that this fact, in and of itself, does not make them exempt from criticism. To the degree that these cognitive states are beliefs, they express a claim to truth and this claim should be supported by epistemic reasons.

Chapter IV - Doxastic Prescriptivity and Control

The idea that belief is prescriptive is usually taken to imply the thesis that beliefs are within our control; that we have either voluntary or reflective control over *belief formation*. However, the thesis that we have voluntary control over our beliefs has been subject to intense criticism (Williams 1970; Hieronymi 2006) and this criticism motivated skepticism in connection with the view that belief is governed by epistemic norms. One important reason for this attitude is the intuitive claim that it does not make sense for us to be held responsible, be assessed as praiseworthy or blameworthy, for something we cannot control.[24] In the following two sections of this chapter I will consider and criticize recent attempts to argue that we in fact have either voluntary or reflective control over our belief formation. In section 3 I will argue against the premise that responsibility requires these types of control. Following Pamela Hieronymi, I will articulate a robust notion of responsibility which applies to a wide range of cognitive attitudes over which we have neither voluntary nor reflective control. Then I will build on Hieronymi's suggestion that we have evaluative control over our beliefs, intentions and other propositional attitudes and that this kind of control is a genuine form of agency. I will connect this proposal with the main thesis that belief is to be modeled on the public act of accepting an assertion and articulate an alterative notion of control, control as responsiveness to epistemic reasons, which will explain our attributions of belief responsibility.

1. Frankish's defense of weak voluntarism

One recent defender of doxastic voluntarism is Keith Frankish (2004, 2007). Frankish distinguishes between weak voluntarism and strong voluntarism. On the latter view, one can form a belief only based on practical reasons, without regard

[24] David Owens (2000) rejects this intuitive thesis. While I sympathize with Owens' criticism, I think that the claim that we are responsible for our beliefs becomes more plausible if supplemented with a notion of doxastic freedom or control. Moreover, I think he would agree with the account of doxastic control as sensitivity to epistemic reasons proposed here.

to considerations related to its truth. Weak voluntarism, the view defended by Frankish, states that although we cannot disregard epistemic reasons in forming a belief, there are still pragmatic reasons involved in our theoretical deliberation. Before stating Frankish's argument in support of weak voluntarism, I will sketch his general view of belief.

As mentioned, Frankish is a defender of the thesis that there are two kinds of belief which should be clearly distinguished and that they correspond to two levels of cognition. At the lower level, the type of belief that humans share with animals, has the following features: it is non-conscious, passive, partial and non-verbal. At the higher level, belief is conscious, active, flat-out and verbal (Frankish 2004, p. 23-24).

Frankish's central idea is that flat-out beliefs are acceptances. When we accept a proposition as true we decide to adopt it as a premise in future deliberations. This is why, for Frankish, having a flat-out belief involves adopting *a premising policy*; the policy is that we should take that proposition as a premise in our reasoning about what to think and what we should do. Now, not every acceptance issues in a full belief. The distinguishing mark of acceptances which result in beliefs is that the agent is prepared to use those propositions as premises in deliberations which are *truth critical with respect to premises* (TCP, for short). These are the kinds of deliberations in which truth is the only criterion for the selection of premises. Thus, for example, a lawyer may accept the proposition that his client is not guilty for professional purposes but that acceptance constitutes a belief only if the lawyer is prepared to use it as a premise in his TCP deliberations.

The policy adoption which results in full belief may take place either at a conscious level or it may be unconscious. In the first case, as a result of theoretical reasoning we decide to accept something as true and use it in our TCP deliberations. However, more often, this decision takes place at an unconscious level, it is unthinking and automatic. For example, most of us routinely accept what we are told unless we have reason to suspect that the speaker is deceitful and

uninformed. Another difference between the two levels concerns the *motivation* we have for the policy. Normally, the motivation for accepting a proposition as a premise in TCP deliberations is based on the high degree of confidence we have in that proposition being true. However, at the sub-personal unconscious level, we may accept propositions for non-epistemic reasons. This level of cognition is not within our control and can be influenced by mechanisms which are not truth-conducive.

Frankish's suggestion is that although our high-confidence in the truth of a proposition is an enabling condition for its unrestricted acceptance, pragmatic factors can also influence the outcome of our deliberation. He writes:

> We can be highly confident that something is true without making up our minds that it is – that is, without having committed ourselves to a policy of taking it as a premise. […] In such a case it may take pragmatic considerations to induce us to come off the fence and commit ourselves to a flat-out belief on the matter. These motives might be of various kinds – prudential, moral, financial, and so on. But most often, I suspect they will be connected with the pragmatics of inquiry and practical deliberation. By forming superbeliefs we equip ourselves with premises which we can feed into conscious reasoning, and the desire to engage in such reasoning and to secure the benefits of doing so will thus be a powerful motive for forming superbeliefs (Frankish 2004: 151).

Let us consider a concrete example. John is quite confident, based on his evidence, that the housing market in his area is at a low and that this would be a good time for him to buy. However, he did not form a belief about that yet because he does not find the evidence fully convincing and he does not think he has a good grasp of all the factors that influence the market. Also, let's suppose, the experts are divided on the issue. Moreover, John's family is pressuring him into making a decision so he does not have time to gather more evidence and study it more closely. Given this situation, John decides to come off the fence and

adopt the premising policy that the housing market is at a low and use this claim as a premise in his practical and theoretical reflections. He thus decides to take a certain epistemic risk but, in his view, the risk is not too high given his degree of confidence in the truth of the sentence.

On this picture deciding to believe is similar to deciding how to act. Both theoretical and practical deliberations are characterized as sensitive to pragmatic factors. However, we might wonder whether this is really the case. Can one make up one's mind about *p* just by reflecting on the fact that time is short and one has to make up one's mind about *p*? It seems plausible that belief is only motivated by the fact that the subject finds the evidence at his disposal convincing. Only reasons that have bearing on the truth of a proposition can make us give our assent to that proposition. Maybe reflection on the time we have to devote to deliberation would make us focus on evidence that we think will find convincing, but reflection on the fact that we are running out of time cannot directly influence belief. It does not seem to be phenomenologically accurate that such pragmatic factors are relevant in doxastic deliberation. This is not to deny the obvious fact that doxastic deliberation is constrained by pragmatic factors like time and the amount of cognitive resources one is able to use. The remark that deliberation is constrained in this way is perfectly consistent with the fact that we judge that *p* only based on the fact that we find the reasons for *p* convincing (Owens 2000).

These facts about belief point to an asymmetry between theoretical and practical deliberation. It is not unusual to be forced into making a decision regarding how to act by reflecting on the fact that time is running out. Once this asymmetry is brought to light we can see that Frankish's account of full-belief is flawed because it describes full-belief as if it is an action (i.e. adopting a premising policy) but this ends up misrepresenting the way we form beliefs. Frankish is able to run the argument that belief can be influenced by practical reasons mainly because he misconstrues belief in the first place, and identifies belief with a kind of action.

2. Reflective Control

Some philosophers argue that we have a reflective control over our beliefs, a kind of control which licenses ascriptions of responsibility (McDowell 1998; Scanlon 1998; Korsgaard 1996). The basic idea behind talk of reflective control is that we influence our beliefs by reflecting on the quality of the reasons we have for holding them. For instance, when we discover that one of our beliefs is not justified we will change it. We perform this change by first forming a second-order normative judgment to the effect that the first-order belief is unjustified and then, if we are rational, the normative judgment will result in the first-order belief being rejected. Similarly, when we form a belief we reflect on the reasons or evidence at our disposal and then form a second-order normative judgment to that effect that the belief in *p* is well-justified and that moves us to form the respective belief.

The main problem with the view sketched above is that it is not clear exactly how the higher-order normative judgment is supposed to influence the first-order one. David Owens (2000) emphasizes the fact that a theory of reflective control should include a motivational aspect and he defines it as follows:

Reflective Motivation: if R is a *prima facie* reason to believe that *p*, reflection on R provides the rational subject with a motive to believe *p*.

Owens points out that this principle is problematic both on an evidentialist and on a pragmatist construal of epistemic reasons. According to evidentialism, only evidence can constitute reasons for belief. However, an agent makes up his mind about *p* only based on what he takes to be *sufficient evidence* for *p*. But then, Owens asks, how can we determine in evidentialist terms when the evidence is

sufficient?[25] It seems that non-evidential factors like how important finding out the truth of *p* is for us, and how much time and cognitive resources we are willing to invest in finding out whether *p*, determine when the evidence for *p* is sufficient. Thus, Owens argues, we cannot be motivated to form a belief that *p* just by reflecting on evidence that *p* because this reflection will not establish whether that evidence is sufficient.

On the other hand, allowing the explicit consideration of pragmatic factors in our reflection about reasons for belief will still not give the agent the motivation necessary for reflective control. As indicated in the proceeding section, reflection on the fact that time is running out will not, in and of itself, motivate an agent to embrace or reject the truth of a proposition. In consequence, Owens concludes that reflection on the quality of our reasons for belief does not have any motivational force and therefore, we have no reflective control over our beliefs. On Owens's view, reasons or evidence for belief are the only things that move us to form beliefs. Belief is only responsive to reasons for belief. Reflection on the quality of the reasons has no role in producing beliefs. When involved in deliberation, the agent uses a standard of what he takes to be conclusive evidence for a certain belief, and that standard is influenced by pragmatic considerations, but the agent does not reflect on those considerations when forming the belief.

Although their view that our first-order beliefs are the result of second-order beliefs regarding the quality of our reasons faces severe difficulties, the advocates of reflective control may still point to the fact that once we discover that one of our beliefs is not well-grounded, we normally tend to revise it. However, the proponents of this view do not offer any explanation why this is the case; of how the second-order belief influences the first-order one. One suggestion is that the agent may identify himself with the ideal of having justified beliefs and, thus, the effect of the higher-order belief is, in a sense, a product of his will (Frankfurt 1971). But this proposal leaves open the logical possibility of

[25] As it was mentioned in the first chapter, Nishi Shah faces this problem when he tries to derive the norm of evidence from the norm of truth. Given that we do not know anything about the context in which deliberation takes place (how important the issue is for the subject and so on) we cannot settle on a standard of justification and we are in danger of setting the bar too high.

an agent who does not value warranted beliefs. This agent will decide to go on believing that *p* while knowing that his evidence does not establish that *p*. This is, however, incoherent. It is in tension with the claim that believing *p* is equivalent to having settled the question *whether p* because one can settle this question only if one finds the evidential reasons in support of *p* to be convincing or conclusive.

The fact that believing that one is unwarranted in judging *p* gives rise to a revision of the belief that *p* can be straightforwardly explained within the framework developed in Chapter 2. Discovering that one's reasons for *p* are not good enough to entitle one to believe *p* shows that one's previous acceptance of *p* was mistaken. When the agent had first deliberated *whether p*, he had taken himself to be entitled to settle the question positively based on the reasons at his disposal because he had found those reasons convincing. But afterwards, further reflection showed that those reasons are not conclusive. Thus, assuming that the agent is minimally consistent and of one mind, he will give up his initial belief the moment he finds the reasons for it unconvincing[26]. In this way, we can explain the normative pressure the agent feels to change his mind in terms of the norms embedded in the practice of assertion. On this picture, the reflective control rational subjects exercise over their first-order doxastic states is just the internalization of the norms which shape the external practice of communication and argumentation.

3. Doxastic Control as Responsiveness to Reasons

The preceding sections have shown that two important views to the effect that we have control over our beliefs face severe objections. The claim that beliefs are not within our control throws the idea that we are responsible for them into doubt. In addition, in order to demonstrate that we are accountable for our beliefs we need to clarify the sense in which we are active with respect to them; the sense in which forming a belief is something the epistemic agent does. In a series of

[26] A similar solution is proposed by Hieronymi (2009, p. 33)

recent articles Pamela Hieronymi (2005, 2006, 2009) puts forward an intuitive view of the agency and control we exercise with respect to our beliefs according to which responsibility for a wide range of cognitive attitudes does not require voluntary or reflective control. Her view is closely connected with the linguistic account of belief developed in Chapter 2. In this section I will present her account, highlight its connection with the present project, and then evaluate it.

Like Shah, Hieronymi maintains that belief is the result of settling the question *whether p*. In her view, reflecting on this definition reveals the way in which we are active and have control over our doxastic states. We can form a belief just by considering *whether p* is the case, considering evidence for and against *p*, gathering more evidence and so on. When we make up our mind and accept *p* as true then we literally change our psychology. In other words, by doing something (settling a question) we literally make up our *minds*. We exercise this form of agency with respect to other cognitive attitudes besides belief. For instance, we normally form an intention to ø by settling the question *whether to* ø; whether an action is worthwhile. Hieronymi labels this form of agency and control *evaluative control*; it basically means that we can change our psychology by considering and answering questions like *whether p*, *whether to* ø and so on.

On this view, it becomes clear why believing is not an action and why it cannot be voluntary. Like Shah, Hieronymi emphasizes the fact that the only reasons that can motivate us to believe *p* are reasons that convince us of *p*'s truth; they are answers to the question *whether p.* On the other hand, pragmatic reasons show that undertaking an action has beneficial consequences; they are answers to the question *whether to* ø. Therefore, we can never believe, in full consciousness, for pragmatic reasons because belief is not an action. Accordingly, while it makes sense to say that an action is voluntary, in the sense that it is the result of the agent's forming and executing an intention, it is senseless to say that a belief is voluntary, since, given that beliefs are not actions, we do not form beliefs by first forming intentions about how to act.

Hieronymi contrasts evaluative control with *managerial control*, the control we exercise over ordinary objects by forming and executing intentions with regards to them (e.g. I can move around the furniture in my room by forming an intention about it and executing it). The type of agency we exercise when we manage ordinary objects is characterized by *voluntariness* (we form and execute intentions to act) and *reflective distance* (we think about the objects of our actions before we act on them). On the other hand, evaluative control does not have these features. We do not make up our minds that *p* is true because we want it to be true. Also, explicit intention on forming a belief is not a necessary condition for actually making up your mind about what is true. What is necessary for forming beliefs is reflecting on what is the case, not reflecting on the activity of forming a belief.

Although distinct, the two types of control are closely related. For instance, sometimes we can manage our doxastic attitudes by taking practical steps which aim at settling the question of truth. If someone is worried whether they have cancer then they can just go to the doctor and find out. In this way, they manage their situation in a way which will affect their doxastic states. Or, more clearly, if one wants to believe that the light is on in one's room then one can just turn on the light-switch[27]. Moreover, evaluative control plays an essential role in every exercise of managerial control because the latter involves the formation and execution of intentions. But, as indicated, forming an intention means settling the question whether to pursue a course of action. Also, when we want to manage or manipulate the cognitive states of other people, we aim at giving them reasons to settle in a particular way certain deliberative questions. Therefore, a complete description of this kind of managerial control has to also make reference to evaluative control.

As indicated, managerial control is characterized by voluntariness and reflective distance. The fact that evaluative control lacks these two features is usually taken to show that it is not a genuine form of agency. However,

[27] Hieronymi uses similar examples following Tom Kelly and Richard Feldman (2004).

Hieronymi argues that this conclusion is in conflict with our strong intuition that we are active with respect to our beliefs; that our opinions are up to us. Our deliberative beliefs are not things that just happen to us, things which we can only passively contemplate, like our sensations or emotions. In consequence, she proposes that we distinguish between two basic types of agency and control and supports the view that voluntariness and reflective control are essential features of the kind of agency we exercise in ordinary intentional action but they are not central to our general idea of agency.

Hieronymi argues that once we understand that we exercise evaluative control over our beliefs and intentions we can have a better grasp of why we are responsible for these cognitive states. She suggests that we can construe responsibility as *answerability*: one is responsible for an activity, an attitude or a state of affairs if one can rightly be asked for the reasons why one engages in that activity, adopted that attitude or brought about that state of affairs. In this sense, we are responsible both for the intentional actions we perform and for the intentional states we are in. We can be asked for reasons both for our actions and for our beliefs, for example. Asking for reasons in these cases is appropriate because intentional action and belief are essentially connected with the agent's having settled the questions *whether to* ø and *whether p*. The reasons that the agent gives for belief are open to assessment by others: his belief may be justified or unjustified, reasonable or unreasonable and so on.

Hieronymi emphasizes the fact that the notion of responsibility as answerability is continuous with the more intuitive notion of moral responsibility. Someone's reasons for thinking or acting in certain ways reveal their moral personality and can be the subject of *reactive attitudes* like resentment and indignation. For instance, when someone explains blowing up a kindergarten by reference to the fact that the explosion gave him aesthetic pleasure he gives us access to the quality of his will or his moral self and it is this will or mind which we assess when we decide whether someone is praise- or blameworthy.

Similarly, in the case of belief, someone may prove gullible when he believes reports published in *National Enquirer* about an imminent alien invasion. By justifying one's belief by reference to such unreliable sources one shows that one is an unreasonable believer and becomes vulnerable to criticism. Our assignment of epistemic blame is triggered by the carelessness our interlocutor reveals when he offers reasons for his beliefs.

It is commonplace that our intentions may be subject to moral judgment. Someone is blameworthy for seriously considering and planning killing innocent people even if his intention is not realized. However, since we form intentions by exercising evaluative control, they are not subject to our voluntary control. Hieronymi illustrates this point by reference to Kavka's toxin puzzle (Kavka, 1983). Kavka's imagined scenario goes roughly like this. A subject is offered a deal: if he succeeds in forming the intention to drink a poisonous liquid he gets a large sum of money. The agent does not need to act on his intention, but just form it. Let's suppose that the person offering the deal has an 'intention detector' in her hands and can detect the physical correlate of the intention to drink the liquid. Now, you may want to form the respective intention but the only way to do that is by answering the question whether drinking the poison is a good thing to do. You cannot form the intention to drink poison just because of the reward. That is not the right kind of reason for forming intentions. Rather, the only way to form an intention to act in a certain way is to answer the question whether it is worthwhile to act in that way; settling this question positively commits you to a certain course of action. But, the point of the examples is that, given that it is hard for one (unless one is suicidal) to convince oneself that drinking poison is worthwhile, intention is not under the direct control of the will. Thus, even if in this case we want to form the intention in order to get the reward, it turns out that we cannot. This shows that we cannot intend at will because intention, like belief, is not a kind of action. But since we are blameworthy for our thoughts and intentions and not only for our actions, it follows that we can be held responsible for things that are not within our voluntary control.

The account developed by Hieronymi fits well with the view of belief proposed in this work. It is intuitively true that accepting an assertion or making a judgment are things that we do, but not actions we perform at will. If someone tells me that the world is going to end in 2012 and bases her assertion on an article in *National Enquirer*, I cannot force myself into accepting her claim. Normally, we accept an assertion when we find the reasons for it persuasive. Thus, given that we exercise a kind of agency and control over our judgments and that judgments result in beliefs, it follows that our beliefs are, in a sense, within our control.

However, I do not think that Hieronymi's notion of evaluative control captures the sense in which we have control over our doxastic attitudes. The fact that we exercise evaluative control over our beliefs seems to be true by definition, since belief is construed as a state in which the answer to the question *whether p* is embedded. Moreover, although Hieronymi is right in pointing out that our demonstrations of entitlement for our beliefs are subject to reactive attitudes, she does not explain why this is the case.

These shortcomings affect the plausibility of her account of doxastic responsibility, which is subject to immediate counterexamples. For instance, some beliefs are the result of neurotic disorders like paranoia or different kinds of phobias. Also, a belief may be the result of hypnosis or systematic indoctrination. Now, a case can be made that, if we equate responsibility to answerability, the neurotic subject is responsible for his beliefs. He is responsible because he can offer reasons for his beliefs, given that he settled the respective questions for himself. But the fact that his reasons are not good reasons does not show, on Hieronymi's account, that the person is not responsible. In other words, as long as the agent is answerable, he is responsible.

These counterexamples show that answerability is a necessary condition of responsibility, but not a sufficient one. The account should be supplemented with further conditions so that it fits our intuitions. In the cases mentioned above, it is also clear that although the agents exercise evaluative control over their

mental life, in the sense that when they settle certain questions they change their psychology, they are not really in control of what they believe. Furthermore, once we discover that a subject's beliefs are rooted in a neurotic disorder, we inhibit our reactive attitudes towards his epistemic behavior.

One plausible suggestion is that, in addition to answerability, the epistemic agent should be able to *rationality respond to criticism or epistemic reasons* articulated by his interlocutors. One problem with agents affected by neurotic disorders is that they are unable to change their beliefs even when confronted with conclusive evidence against them. Their condition makes them incapable to react to the force of reasons and evidence according to the standards which shape the epistemic practices of their community (Steup 2008). For instance, someone who has a pathological fear of flying will think that their plane will crash even if the probability of that event happening is extremely low. We can even imagine that, as long as they can vividly imagine the plane crashing, their phobia can be triggered even if there were no such accidents in the history of flying.

One difficulty with this suggestion is that although an agent may react rationally to criticism in one instance, he can move on and make the same mistake in other circumstances. For instance, even if a subject recognizes that he cannot make hasty generalizations about all asian people based on claims about his two asian neighbors he can continue to make such mistakes in reasoning about other subject matters. The question is: does he have doxastic control or not?

Before answering this problem it is worth recalling that our public criteria for doxastic control and responsibility, like those for moral responsibility, are not precise. There are always problematic borderline cases. However, in our ordinary attributions of belief responsibility, it is important that the subject *understands* that he makes a mistake in reasoning; that an inference he considers correct is not in fact entitlement-preserving. But the criteria of such understanding involve making sure that the agent does not make the same mistake in other contexts. In other words, if the agent repeats his error then we can say that he did not understand the initial criticism and he is not yet in control of his doxastic states.

On the other hand, if the agent meets the public criteria which signal understanding a mistake in reasoning then we ascribe him more theoretical responsibility.

This proposal is in line with the linguistic account of belief developed in Chapter 2. There, the notion of responsibility was correlated with that of authority. When a speaker puts forward a claim as true he licenses or authorizes others to rely on it. Because the speaker exercises this authority he is held responsible for his assertion and can be asked to prove his entitlement to it. However, if it is discovered that a speaker's reports are not based on a good sense of reasons and entitlement but rather are the outcome of a cognitive malfunction, then the agent temporarily loses his authority. His reports are considered unreliable and, in consequence, the question of responsibility loses its point.

A similar remark can be made in relation to the act of accepting a claim. If we learn that a person who is extremely afraid of flying accepted without question certain implausible statistics regarding the probability of plane crashes, we do not accept the information based on her authority but we rather tend to double-check it. This mistrust is explained by the fact that we know that the subject's belief formation is not fully within her control and, therefore, she does not exercise full responsibility with respect to it. In consequence, the dimensions of authority and responsibility are connected with both the act of making a claim and in that of accepting a claim (and forming a belief). In both these instances, if we lose our authority, if we fail to react appropriately to the force of reasons, then we also stop being held responsible.

4. Conclusion

This chapter aimed at proposing an account of belief responsibility which is not based on the controversial idea that we have voluntary control over our beliefs. Following Pamela Hieronymi, it was argued that we are responsible for our beliefs in the sense that we are answerable for them; that we have to offer reasons

in their support. The quality of our reasons for belief may be the subject of reactive attitudes like disapproval, outrage, praise and so on. Since responsibility requires control, this minimal notion of responsibility has to be supplemented with an account of doxastic control. Based on our ordinary practice of attributing doxastic responsibility, I argued that we take epistemic agents to have doxastic control when they are able to understand criticism by correctly responding to epistemic reasons. If subjects are incapable of readjusting their assessment of evidence in light of appropriate criticism then they lack doxastic freedom and control. It is noteworthy that this is an account of doxastic control and responsibility in terms of the subject's abilities to successfully participate in the public practice of assertion. Since exercising this ability does not require that the agent be able to form conscious beliefs or reflect on his actions, the view should not give rise to concerns about circularity.

Chapter V - The Teleological Account of Belief

The proponents of the teleological account of belief unpack the metaphor that beliefs 'aim at truth' in non-prescriptive terms. In the first section of this chapter I will present and evaluate Velleman's view (2000) concerning the definition of the 'aim of truth' and how this aim can be realized. Nishi Shah (2003) objects to Velleman's account by arguing that the connection it establishes between belief and truth is too weak to explain the phenomenon of doxastic transparency. Then I will present and assess two more refined versions of the teleological account, Steglich-Petersen's (2006) and Hamid Vahid's (2006). Although these more recent teleological accounts avoid some of the problems faced by Velleman's original view, they are still vulnerable to some important objections.

1. The Aim of Truth and Transparency

As indicated in the first chapter, David Velleman tries to distinguish between belief and other cognitive states, like assuming or imagining,[28] which involve accepting a proposition as true. Velleman's main point is that what differentiates belief from other cognitive states is the fact that it aims at truth: "What distinguishes a proposition's being believed from its being assumed or imagined is the spirit in which it is regarded as true, whether tentatively or hypothetically, as in the case of assumption; fancifully, as in the case of imagination; or seriously, as in the case of belief" (Velleman 2000, p. 252). He offers the following characterization of the specific aim of belief: "believing involves regarding a proposition as true with the aim of so regarding it only if it really is. Thus, to believe a proposition is to accept it with the aim of thereby accepting a truth" (Velleman 2000, p. 251).

Velleman argues that there is a broad spectrum of ways in which the aim of belief may be realized. At one end of the spectrum, the aim of belief is the aim

[28] Velleman argues that imagining is taking true in the context of pretending.

of the epistemic agent. The agent has an explicit intention to find out the truth about a subject-matter. In this case, the agent deliberates whether to believe that *p* or not and he either accepts or rejects the propositions depending on the evidence and arguments at his disposal. The intention to accept *p* only if *p* is true *regulates the formation and maintenance of deliberative belief*.

However, Velleman is aware that many of our beliefs are not the result of explicit deliberation but rather the result of the workings of some automatic cognitive mechanisms. In these cases, he argues, the aim of truth is an aim those mechanisms have as a result of the way they were designed by natural selection or by education or both. We say that a cognitive system aims at truth when it *regulates* its cognitions in ways designed to ensure that they are true, by *forming, revising and extinguishing them in response to evidence and argument.*

There are two attitudes an epistemic agent may have with regards to the cognitive mechanisms working within him. First, the epistemic agent may endorse the working of these cognitive systems and, in this case, their outcome may count as his doing and may be attributable to him. Second, there are cases in which the agent may dissociate himself from one of these mechanisms and in these cases there is no sense in which the mechanisms are regulated by the individual but they still count as producing beliefs.

It is noteworthy that, in Velleman's view, there are cases in which belief might be influenced by psychological mechanisms which do not aim at truth. He writes: "There are probably psychological mechanisms that cause, and are designed to cause, beliefs that happen to diverge from the truth. Evolution or education may have given us dispositions to err on the side of caution in perceiving predators, to overestimate our own popularity and so on. But my thesis is not that belief is completely shielded from mechanisms that tend to make it false; my thesis is that belief is necessarily subject to mechanisms which tend to make it true" (Velleman, 2000).

Nishi Shah argues that Velleman's account faces a dilemma. On the one hand, Velleman has to account for cases in which a belief is the result of a

psychological process like wishful thinking. In this case, one mechanism which is not truth-tracking *interferes* with the working of a truth-tracking mechanism. It follows that the fact that 'the aim of truth' is realized in the cognitive mechanism designed to track truth *is compatible with it being influenced by non-truth-tracking mechanisms*. In the case of wishful thinking, *practical, non-evidential considerations influence the workings of the cognitive mechanism which aims at truth.* All that aiming at truth requires is that the cognitive mechanism designed to represent reality accurately still have *some* influence over the resulting belief. However, Shah argues, this account is unable to explain why in the case of deliberation about belief *only* evidential considerations are taken to be relevant. Why is it that when we explicitly reason about belief we do not take into account practical, non-evidential considerations and we only focus on facts that bear on the truth of the proposition we are considering? Put differently, from the first-person point of view, the question *whether to believe that p* is equivalent to the question *whether p is true* and only considerations bearing on the truth of *p* are relevant in answering the latter question. But, if we construe 'the aim of truth' in the way suggested by Velleman, then we cannot account for this fact about doxastic deliberation. This is the first horn of the dilemma.

On the other hand, if Velleman chooses to construe 'the aim of belief' in such a way that it does not conflict with the phenomenon of doxastic transparency then he cannot leave room for the influence of non-evidential considerations on belief. But then, and this is the other horn of the dilemma, Velleman's account cannot make justice to the common-sense view that sometimes beliefs are influenced by non-evidential considerations. If the connection between belief and truth is too strong then Velleman's account does not get the extension of 'belief' right because there are certainly cases like wishful thinking, self-deception and so on, in which pragmatic, non-evidential considerations do influence beliefs (Shah, 2003).

2. Refining the Teleological Account

(a) Steglich-Petersen's view

Steglich-Petersen defends the teleological account against Shah's objection. He argues that the dilemma formulated by Shah has a mistaken presupposition: "that the concept of belief involved in, and thus framing, intentional doxastic deliberation must be exactly the same concept of belief as that used to pick out beliefs regulated by subintentional mechanisms" (Steglich-Petersen 2006, p. 510). Steglich-Petersen distinguishes between the deliberative concept of belief and the empirical concept. The deliberative concept of belief refers to the activity of doxastic deliberation as it is conceived by the epistemic subject himself. On the other hand, the empirical concept of belief refers to the activity of belief formation construed from a third-personal point of view. Steglich-Petersen argues that the primary, deliberative concept of belief is the one described by Velleman as follows: belief is the attitude of accepting a proposition with the aim of thereby accepting a truth. In his view, we can make sense of the empirical concept of belief on the basis of the deliberative concept of belief but not the other way around.

Steglich-Petersen illustrates his point by making an analogy between believing and concealing something. It would be a mistake to define what we try to do when we want to conceal an object from someone in terms of the empirical concept of concealing. It is a fact that sometimes we fail to conceal something and thus that there is a weak relation between our attempt at concealing something and success. However, this does not imply that when we want to conceal something from somebody we do not want to do it successfully. The point is that we should not use the empirical, third-personal concept of concealing in trying to explain what concealing is from the first-person point of view. The correct strategy, Steglich-Petersen argues, is to start by defining concealing in terms of someone's intention to put an object in a place where someone else may not

discover it. With this definition at hand, we can construe the empirical concept of concealing as what all action-types executed with the intention of concealing have in common. What we get, from the third-person perspective, is a weaker connection between concealing and success.

Similarly, while we cannot explain deliberative belief by reference to the empirical, descriptive concept of belief, we can follow the reverse order of explanation and account for beliefs which are not the result of explicit intentions on the part of the agent in terms of deliberative beliefs. Steglich-Petersen notes:

> So the primary concept of believing *p* is that of accepting *p* with the aim of doing so only if *p* is true. Even if they fail to achieve that goal and are merely weakly responsive to truth, cognitive processes can count as instances of beliefs in virtue of being brought about by someone who has this aim, with the important qualification of being at least to some degree conducive to this aim. Furthermore, cognitive states and processes that are not connected with any literal aim or intention of a believer can nonetheless count as 'beliefs' in virtue of sharing precisely the descriptive characteristic of being to some degree conducing to the hypothetical aim of someone intending to form a belief in the primary strong sense (Steglich-Petersen 2006).

Now, Steglich-Petersen claims that the phenomenon of doxastic transparency can be explained in terms of the primary, deliberative concept of belief. Moreover, this explanation does not appeal to the idea that belief involves any norms. The basic idea is that deliberation is defined as an intentional activity aimed at truth. When we deliberate whether *p* is the case or not we have an interest to find out the truth about *p* and this is why we immediately consider whether *p* is true. As Steglich-Petersen puts it: "The present account builds an interest in truth into the very concept of belief"

Furthermore, according to Steglich-Petersen, epistemic norms like the norm of evidence are instrumental norms in the sense that they are just *means to*

our end, which is the truth. "Beliefs ought not merely to be true – they also ought to be formed in ways that ensure or make it likely that they are true. A natural explanation of such epistemic norms is that following them promotes the aim of believing truly. Epistemic norms thus turn out to be instrumental norms, deriving their normative force from their ability to guide us to achieve our aims" (Steglich-Petersen 2008).

(b) Vahid's view

Like Steglich-Petersen, Vahid (2006) also departs from Velleman's original teleological account. Vahid draws the important distinction between (a) belief aiming at truth and (b) aiming at true beliefs. While (a) is a claim about the nature of belief, (b) is a description of the epistemic project of having true beliefs. With this distinction at hand, Vahid analyses Velleman's initial characterization of intentionally aiming a belief at truth: "When the subject intentionally sets out to accept whichever propositions is true, he will be guided by methodological beliefs about how to discriminate truth from falsehood" (Velleman 2000, fn. 17, p. 252). Based on these remarks, Vahid sums up Velleman's view as follows:

(B) Believing *p* = (i) regarding *p* as true

(ii) regarding the methodology for regulating (i) as truth-conducive.

However, Vahid argues, this is a characterization of justified belief rather than belief itself and it is directly connected with the epistemic project of having true beliefs, not with elucidating the nature of belief itself.

Vahid defends a deflationary view of belief. He suggests that we should only keep the component (i) of the above definition of belief. In his view, what differentiates belief from other cognitive attitudes is that belief involves regarding a proposition as true for its own sake rather than for the sake of something else.

While assuming, for instance, involves regarding a proposition as true with the aim of seeing what follows from it, believing involves having the aim of regarding a proposition as true seriously. In other words, regarding something as true is the primary aim of belief and it is just a means to some other aims in the case of other cognitive states. Vahid writes: "our primary aim in believing *p* is its very acceptance as a true proposition. In this case, our regarding *p* as true is not intended to serve a different aim. Rather, our primary aim in believing *p* is to regard it as true for its own sake (period). It is, in other words, to regard it as true seriously" (Vahid 2006, 326).

3. Evaluation of the Teleological Account

We can start the evaluation of Vahid's and Steglich-Petersen's versions of the teleological account by noticing that it is not entirely clear that their views manage to capture the relation between truth and belief in non-prescriptive terms. Although Vahid argues that we should distinguish the epistemic project of having true beliefs from the idea that belief 'aims at truth', I do not think this clearly implies that belief is not still subject to norms of evidence. Vahid's claim that to believe *p* is to aim at regarding *p* as true *seriously* may be interpreted as suggesting that the subject takes himself to be warranted in believing *p*. This would be in contrast with regarding *p* as true for no good reason. However, if being guided by evidence is not part of what Vahid means by seriously holding *p* true then his characterization of belief is in conflict with the ordinary use of the concept. His account is unable to explain a basic fact about belief: that it is something for which reasons can be asked[29]. Although aiming at truth is not to be

[29] Vahid's account faces another problem. As pointed out, both Velleman and Steglich-Petersen argue that belief can be directed at truth by the agent when he accepts *p* with the aim of accepting a truth. In this case, we have a distinct idea about how the agent may fail in realizing his aim: the agent fails when his belief turns out to be false. However, given Vahid's characterization of the aim of belief, it is not clear how an agent can fail in his attempt to believe the truth. According to Vahid, our aim in believing *p* is regarding *p* as true for its own sake. But how can this be an aim? Given that aims are by definition states of affairs which may not be realized it follows that it is

conflated with pursuing the epistemic project of having true beliefs, this does not imply that belief can be divorced from the notions of evidence and justification. If we do not take evidence as a guide when we aim our beliefs at the truth then it is not clear why our activity can be described as 'seriously aiming at truth'.

In consequence, I think that Vahid's account can gain some plausibility only if it adds Steglich-Peterson's remark that norms of evidence guide our belief formation. But then, it can be argued that, given that the 'aim of truth' gives rise to evidential norms, on this view, truth is prescriptive for belief. For even if we accept Steglich-Petersen's idea that norms of evidence are hypothetical norms, it does not follow that belief is not intrinsically subject to these norms. The force of these norms is contingent on aiming at truth and belief necessarily involves such an aim; an interest in truth is built into the concept of belief. Therefore, following these evidential norms is constitutive of belief. If someone does not follow evidential norms while reflecting on what to believe it is doubtful whether his activity can actually be described as deliberation about belief because, given that his reflection is not guided by evidence, it is not clear that he is aiming at truth.

One plausible consequence is that the subject, since he is interested in finding out the truth about something, should be sensitive to criticism from others. He ought to make public the reasoning behind his opinions and discuss it critically. In this way, the notion of aiming at truth is close to having the same normative import as the notion of commitment to truth developed in Chapter 2. However, Steglich-Petersen might insist that his teleological account does not, strictly speaking, give rise to such social epistemic obligations. Aiming at truth is something that the individual can do in isolation. I think that this position is deeply problematic.

---

possible that, when believing *p*, we fail to regard it as true. We can, in full awareness, believe *p* while thinking that *p* is false. However, this is in conflict with the phenomenon of doxastic transparency and with the intuitions regarding the nature of belief made vivid by Moore's paradox. Thus, I think that Vahid's deflationary account of belief cannot be a genuine teleological account because his description of the aim of belief is in tension with a constitutive feature of belief, belief transparency.

First, it is plausible to argue that the norms of evidence the agent follows in his private reflection are initially acquired from the social practice of argumentation. Recall that on the view developed in Chapter 2, deliberation about belief is the internalization of abilities used in assessing claims put forward by others (Frankish 2004; Sperber et. al. 2010). When we deliberate whether *p* we consider reasons which can be offered in support of *p* and against it. If we find the reasons for *p* to be powerful, then we accept it as information, thus settling the question whether *p*. This act, it was argued, is analogous to the act of accepting a claim in the public sphere as a result of the persuasive reasons offered in support of it. For instance, in trying to find out whether global warming is real, one may study meteorological data and the analysis of experts. Appealing to the reports of acknowledged experts is a standard way of proving entitlement to a claim in the public sphere and it therefore becomes a tool used in private research. Similarly, suppose someone wants to determine whether his favorite football team will win the next game. To this end, he might consult statistics, check their opponent's record, the scores in the previous matches between the two teams, and find out whether there are injured players. If it turns out that his favorite team has won the last ten games against that specific opponent, that they play at home, and that they have no injured players, then the agent will induce that they will win again. But the fact that the epistemic subject finds this evidence convincing is the result of his ability to recognize strong inductive arguments. This ability was first acquired in the course of his participation in the public practice of assertion. As a participant in this practice he was often faced with the problem of whether to accept someone else's claim; whether the reasons they offered in support of it were good. In this context, acquiring the ability to distinguish strong reasons from weak ones becomes useful for the agent. In short, if Steglich-Petersen is committed to the claim that 'aiming at truth' is something the individual can do in isolation, then the onus is on him to offer an individualist account of the subject's ability to evaluate the force of epistemic reasons.

A second important point has to do with the elements of authority and trust involved in deliberation about belief. Conducting inquiries only in isolation implies that the agent has a high degree of trust in his intellectual capacities and he takes himself to be an expert and have special authority. But the status of an expert is a social status. It implies a specific authority and responsibility and it is acquired *after* the subject has proved his reliability in the public sphere. Thus, a subject is rarely warranted in fully trusting his ability of getting to the truth and refusing to answer challenges to his views on the basis of his authority. As long as the inquirer does not benefit from the privileges of an expert, engaging in public debates of his opinions may be a useful tool for discovering truths. To sum up, as long as the agent aims at truth and social practices like argumentation and critical discussion are designed to track truths and eliminate falsehoods, the epistemic agent has an obligation to participate in these practices. Thus, Steglich-Petersen is forced to take into consideration the social aspect of aiming at truth. But then, his teleological view of belief will have the same normative implications as the claim that belief involves a commitment to truth.

Before developing a criticism of the main insight animating the teleological account of belief, I will present a brief assessment of the thesis that the force of norms of evidence is contingent on the agent's intention to find out the truth. As mentioned in the preceding section, Steglich-Petersen maintains that norms of evidence guide us in our deliberation about belief simply because this kind of reflection is a goal-oriented activity. However, it might be argued that our obligation to follow epistemic norms does not depend on whether or not we have the goal of finding the truth about a certain subject matter (Kelly 2003). In general, we are interested in obtaining the truth in connection with certain subjects but other matters are of complete indifference to us. Suppose that someone is completely indifferent as to whether Bertrand Russell was left-handed. Furthermore, suppose that that person stumbles upon strong evidence for the conclusion that Russell was actually left-handed. Then, intuitively, we think that the agent ought to believe that proposition independently of whether he has

an interest in the matter. As Kelly puts it: "Once I come into possession of evidence which strongly supports the claim that *p*, then I have epistemic reasons to believe that *p*, regardless of whether I presently have or previously had the goal of believing the truth about *p*" (Kelly 2003, p. 625).

I think that this intuition can be explained by reference to the fact that the purpose and proper function of our epistemic practices is independent of the epistemic aims of particular individuals. As indicated in Chapter 2, the social practice of assertion aimed at facilitating the exchange of information and coordination between the members of a group. On this picture, it is plausible that participants should be able to store bits of information that are not of immediate concern for them as individuals.

In the following I will develop my main critical point against the teleological account. Although the view is developed differently by Velleman, Steglich-Petersen and Vahid, I think it is still vulnerable to the basic objection that it is circular. Before stating the objection, recall that Steglich-Petersen's reply to Shah's objection involved taking the definition of the aim of truth in terms of the literal aim or intention of the agent to be our primitive concept of belief. Taking the empirical concept of belief as primitive makes the teleologist unable to explain doxastic transparency, as Shah argues. Thus, the proponents of the teleological account basically argue that belief aims at truth in the sense that epistemic agents have an aim, intention or purpose when they believe something. Believing is performing an intentional action; it is accepting a proposition as true. However, no matter how that aim is specified, this thesis presupposes the controversial claim *that the idea of doing something intentionally is more primitive and has explanatory priority over belief.*[30] But how can we make sense

[30] It might be argued that this presupposition does not make the teleological account strictly circular. After all, the proponents of the account define the cognitive state of belief in terms of a different cognitive state, that of intention. Recall that Dennett (1978) and Stalnaker (1984) claim that the concepts of belief and desire are interconnected. However, in a forthcoming paper called "Desires are Conceptually Prior to Beliefs" Steglich-Petersen explicitly rejects such an account. He writes: "Trivially, if believing P entails having the aim of believing P truly, believing P will be successful if and only if P is true – at least as far as the truth aim is concerned. Obviously, such an analysis can be successful only if aims or desires are conceptually prior to beliefs" (Paper

of the idea of intentional action and of having literal aims without the idea of explicit belief? According to the standard accounts of intentional action, both belief and desire are involved in explaining action (Anscombe 1963; Goldman 1970; Davidson 1980). Belief has an essential role in guiding action. When someone has a goal or desire he is guided by beliefs regarding how to realize that goal. In our case, aiming a belief at truth presupposes some prior beliefs about what is true.

Furthermore, the concept of intention is also closely related to the concept of belief. Intending that *p* involves wanting to make *p* true. But this implies the belief that *p* is not true. Also, it might be argued, having an intention involves the additional belief that one is able to successfully achieve his purpose and that one is able to determine when what was intended became true and the action was successful. Thus, it seems that the concepts of belief and intention are interrelated and we cannot make sense of the latter without the former. All these considerations point to the conclusion that, if the aim of truth is an intentional aim, the teleological account is circular because the notion of having an intentional aim presupposes the notion of belief.

One way in which the proponents of the teleological account can answer this objection is by appealing to the distinction between two levels of mentality. This would involve construing the act of accepting a proposition with the aim of accepting a truth as the outcome of first-level desires and beliefs.[31] As mentioned, first level beliefs and desires are unconscious and independent of language.

I think this suggestion is implausible. The acceptance which characterizes belief should be based on the desire to hit the truth and also on beliefs regarding the means to the truth. In other words, the subject desires the truth and has beliefs about how to realize that desire. Then, it seems that the content of these lower-

---

presented at The Normativity of Belief and Epistemic Agency Workshop, Mexico City, October 2009). However, as indicated in section 2, in his paper "No Norm Needed: the Aim of Belief" (2006) Steglich-Petersen makes the stronger claim that *intentional action* is conceptually prior to belief.

[31] This possibility is one of the main ideas of Keith Frankish's account of conscious belief (2004). I will present and evaluate his account in Chapter 7.

level states should include concepts like truth, falsity and evidence. However, it is not plausible that these concepts can be acquired by a cognitive system which does not use language. Recall that lower-levels beliefs or desires are the product of cognitive mechanisms and modules which we share with non-human animals. From an evolutionary perspective, it can be argued that the content of these states should be determined by simple concepts. The acquisition of sophisticated abstract concepts like truth, falsity or evidence is facilitated by the acquisition of natural language.

A more plausible route is combining the insight that believing *p* is accepting *p* with the aim of accepting a truth with the characterization of the practice of assertion provided in Chapter 2. Given that the proper function of the speech act of assertion is the dissemination of information, it follows that when a subject accepts an assertion he implicitly aims at accepting a truth. However, his aim is not to be explained in terms of his individual desires or intentions, but in terms of the biological purpose of the practice as a whole. Since the practice of assertion got selected in order to facilitate the transmission of information between the members of a group, participating in this practice involves aiming at truth. In this way, the teleological account of belief can answer the circularity objection and can fit the teleological account of assertion described in Chapter 2. Nonetheless, this implies that the teleological account is no longer an alternative to the prescriptivity view but it becomes part of the normative account.

4. Conclusion

This chapter showed that the teleological account of belief is not incompatible with the thesis that truth is prescriptive for belief, but actually presupposes it. The fact that belief involves intentionally aiming at accepting a truth means that belief is subject to evidential norms. Moreover, since doxastic deliberation is the internalization of social argumentative practices, it follows that the evidential norms which shape deliberation originate in public institutions. Then, the

individual is still subject to the epistemic standards of his community. I also argued that the teleological account is vulnerable to the objection that it is circular. One plausible way of solving this difficulty is by construing the act of aiming at truth, not as an act of the individual considered in isolation, but as a function of his ability to communicate effectively with others.

Chapter VI – Inferential Commitments

1. Introduction

The previous chapters articulated the idea that truth is both constitutive and prescriptive for belief. Many philosophers also argue that the logical relations between beliefs are *constitutive* of belief; the fact that an intentional state has a place in a logical pattern of beliefs determines that it is a belief (Davidson 2004; Zangwill 2005; Wedgwood 2007). This chapter will aim at specifying the prescriptions associated with these constitutive logical norms. First, however, let us clarify the sense in which logical norms are constitutive by reference to some examples. It might be argued that the fact that believing *p & q* implies believing *q* plays a constitutive role in the sense that if this relation of logical implication did not hold then we would not be able to say that we identified the belief that *p & q*. Similarly, it is essential to the belief that *p* that it is inconsistent with the belief that not-*p*. If this logical relation did not hold we would not be able to identify the belief that *p*. It is plausible to maintain that these logical relations play a constitutive role for belief because belief is defined by its relation to truth. To illustrate: believing *p & q* means taking it to be true. But if *p & q* is true then *p* must be true. Thus, if the belief that *p & q* did not entail the belief that *p* then we would not be warranted in taking it to be a belief, because believing *p* and *q* means taking it to be true and if *p & q* is true then *p* is true. Similarly, believing *p* implies taking *p* to be true and, thus, taking not-*p* to be false. If these beliefs were not inconsistent then we would not be entitled to call them beliefs, because believing *p* necessarily involves taking it to be true and if *p* is true then not-*p* is false.

Interpretivists about intentionality take the fact that belief is defined by logical requirements to mean that someone can count as a believer only if he follows these constitutive logical norms (Davidson 2004, Brandom 1994). In the next section this general claim will be revised in order to accommodate the obvious fact that many agents are taken to be believers although they do not have

perfect logical abilities. The claim will be that certain minimal deductive abilities are necessary for having the status of an agent and, thus, believer.

Section 3 will present an objection to the claim that logical norms can play both a constitutive and a prescriptive role for belief. As in the case of the relation between belief and truth, I will answer this objection by appealing to the idea that belief involves a commitment to truth. This commitment, it will be argued, involves a commitment to the logical consequences of a belief and also, it forbids one from believing incompatible claims. Although this inferential commitment is essential to belief, there are cases in which we fail to discharge it. However, failure to discharge it does not immediately lead to losing one's status as a cognitive agent.

Sections 4 to 6 are dedicated to clarifying and refining the account of the inferential commitment essential to belief, in light of some objections and counterexamples. Specifying the form of this inferential commitment has to meet a basic condition of adequacy: it should take into account that we are finite cognitive agents with limited logical abilities. In addition, any account which claims that belief is governed by deductive constraints should answer to two important counterexamples: the preface and lottery paradoxes.

2. Deductive Norms and the Finitary Predicament

Davidson is an interpretivist with regards to meaning and intentionality and he explains having intentional states in terms of attributing intentional states. He argues that when we interpret someone's behavior by reference to intentional states we use the principle of charity: we assume that their beliefs are mostly correct and coherent. In his view, the principle of charity is an essential part of a theory of interpretation. A theory of interpretation is assigning meaning to one's utterances and also intentional states to the agent. But then, in order for someone to count as having a belief, he should have many other beliefs which are logically connected with the first one. There are two reasons why someone cannot have an isolated belief. First, the content of that belief is semantically connected with the

content of other beliefs. Davidson defends *semantic holism*, the view that the meaning of a word is partially determined by the meanings of other words. Moreover, beliefs are identified by their place in a logical structure; by the way in which they logically interact with other beliefs. Davidson writes:

> A belief that it is about to rain would lose much of its claim to be just *that* belief if it did not have some tendency to cause someone who had it and wanted to stay dry to take appropriate action, such as carrying an umbrella. Nor would a belief that it is about to rain plausibly be identified as such if someone who was thought to have that belief also believed that if it rains it pours and did not believe that it was about to pour. And so on: these obvious logical relations amongst beliefs; amongst beliefs, desires and intentions; between beliefs and the world, make beliefs the beliefs they are; therefore, they cannot in general lose these relations and remain the same beliefs. Such relations are constitutive of propositional attitudes (Davidson 2004, p.196).

One kind of constitutive logical relations for belief are relations of logical consequence. Davidson offers the example of *modus ponens*: it is constitutive of the beliefs that *p* and *if p then q* that they logically imply *q*. Similarly, it is constitutive of the belief that *p and q* that it logically implies the belief that *p*. The relation of incompatibility that leads to inconsistency also plays a constitutive role for belief. Thus, it is an essential fact about the belief that *p* that it is incompatible with the belief that *not-p*. If this logical relation did not hold then we would not be able to identify that belief.

One standard objection against the view that deductive rules of reasoning are constitutive of belief is that while the logical abilities of normal agents are far from ideal we have no problem attributing beliefs to them. For instance, it seems that Davidson's view implies that an agent should be perfectly rational in order to be interpretable; that he should believe all the logical consequences of his beliefs and be absolutely consistent. However, real agents are in a finitary predicament; we have a limited amount of time, memory and computational resources to devote

to reasoning (Cherniak 1982). For example, it is impossible for an epistemic agent to think of all the logical consequences of her beliefs because there are infinitely many such consequences. Similarly, when someone acquires a new belief she is not checking *all* her beliefs to see whether there is an inconsistency. Usually, we do not have enough time to do such a detailed investigation and also it is always possible due to our cognitive imperfections that some of the inconsistencies go unnoticed.

Both Cherniak and Davidson respond to this objection by emphasizing the fact that following some basic rules of reasoning is constitutive of personhood or agenthood. Cherniak develops the idea of minimal rationality and argues that a system has to have a minimal deductive ability in order to count as a cognitive system. In the context of discussing the assent theory of belief, the view that to count as believing a proposition an agent only has to give his assent to it, Cherniak writes:

> A cognitive theory with no rationality restrictions is without predictive content; using it, we can have virtually no expectations regarding a believer's behavior. There is also a further metaphysical, as opposed to epistemological, point concerning rationality as part of what it is to be a *person*: the elements of a mind – and, in particular, a cognitive system – must fit together or cohere. A collection of mynah bird utterances or snippets of the *New York Times* are chaos, and so at most just a sentence set, not a belief set. Again, no rationality, no agent (Cherniak, 1986, p. 6).

Thus, according to Cherniak, the rationality constraints for belief are constitutive of our idea of an epistemic agent. In consequence, Cherniak tries to find a middle way between the view that there are no logical constraints on belief and the view that belief requires ideal deductive abilities. His answer is that for a mental state to count as a belief, it must interact in some *significant subset* of the

ways that such a belief would interact if it were held by a fully rational agent. He formulates the following minimal inference condition on deductive ability:

> If A has a particular belief-desire set, A would make *some*, but not necessarily all, of the sound inferences from the belief set that are apparently appropriate.

For Cherniak, an appropriate inference is one that is useful in a certain circumstance and, in conjunction with a certain desire, leads to an appropriate action.

> Suppose the person's putative belief set included the beliefs 'If it rains, then the dam will break' and 'it is raining'. The person would not conclude that the dam would break, even if this conclusion would be obviously useful – for instance, when the person also believed he was below the dam and would be drowned if it broke, was not suicidal and so on. Therefore, the person would not be able to undertake any appropriate action on the basis of his beliefs that depended on this information, such as fleeing. The putative agent's deficit of logical insight, and so of rational action, would exclude him from having beliefs according to the minimal general rationality constraint." (Cherniak, 1986, 10-11).

Cherniak construes a similar argument in connection with the notion of logical consistency. He rejects the extreme case that there are no consistency requirements on a set of beliefs. This, again, would lead to epistemological and metaphysical problems. First, we would not be able to predict an agent's behavior if there are no consistency requirements on his beliefs. Secondly, maximal inconsistency would lead to the dissolution of the idea of an agent. No rationality, no agent. Cherniak argues also against the idea that an epistemic agent should have a belief set that is perfectly consistent. He argues that the occurrence of an inconsistency in someone's belief set is not enough to make us doubt his ability to

have beliefs. In his attempt to find a middle way between these two extremes Cherniak formulates the minimal consistency condition as follows:

> If A has a particular belief-desire set, then if any inconsistencies arose in the belief set, A would sometimes eliminate *some* of them.

According to Cherniak, these minimal rationality requirements are context-sensitive and should be used in conjunction with an empirical theory of feasible inferences and of memory structure. Simply put, fixing the level of minimal rationality for an agent requires knowing both what inferences are easier for him and the way his memory works. For instance, normally, inferring $\sim q \rightarrow \sim p$ from $p \rightarrow q$ is easier than inferring $(\forall x)\, Fx \rightarrow (\forall x)\, Gx$ from $(\exists x)\, (\forall y)\, (Fx \rightarrow Gy)$. Also, a subject is not expected to infer q from p and $p \rightarrow q$ in a certain situation if his memory structure is such that he cannot activate these two pieces of information at the same time and hold them before his mind.

Given these observations, Cherniak argues that the requirements of minimal rationality would be very different for a person who has an alternative feasibility ordering from that of 'normal' human beings. Someone who computes easily inferences which we find very hard and has a hard time grasping rules of reasoning like *modus ponens* can still count as a minimally rational agent. Thus, Cherniak infers, minimal rationality is context-sensitive and has a cluster structure. The agent must be able to perform a set of correct inferences from his belief-set, but no inference in particular. The agent can fail to perform simple inferences and compute complicated ones and vice-versa. Thus, some inferences are always constitutive of having a belief but identifying which ones play a constitutive role depends on the context of attribution.

Like Cherniak, Davidson argues that some inferential norms play a constitutive role with regards to belief even if we accept that finite agents can make logical mistakes. Davidson makes the point that attributions of irrationality make sense only against a background of attributions of rationality. He writes:

"We make sense of aberrations when they are seen against the background of rationality." (Davidson 2004, p. 196) For instance, we can understand that someone does not draw the inference from *p*, and *if p then q* to *q* only if we assume that he understands the propositions involved. In other words, among other things, we assume that he does not also believe that not-*p*, and that *it is not the case that if p then q*. If the agent holds many inconsistent beliefs then we would get to a point where we could not attribute any beliefs to him. "The essential point is that the more flamboyant the irrationality we ascribe to an agent, the less clear it is how to describe any of his attitudes" (Davidson, 2004, p. 196).

To sum up, both Cherniak and Davidson argue that some rules of reasoning are constitutive of belief; that an agent does not count as having a belief unless he is aware of some of its logical relations with other beliefs. The main difference between the two philosophers is that Cherniak allows for the possibility of agents with different feasibility orderings from ours whereas Davidson argues that being able to master simple inferences is essential to grasping the meaning of logical constants and, thus, to the status of a logical agent.

It might still be objected that there may be cases in which subjects do not even follow these basic deductive norms of reasoning and, on the view proposed here, it follows that they do not have beliefs. These agents, the objection might go, do not have a different feasibility ordering, as described by Cherniak, but they have very poor reasoning skills. One example of this type was discussed in Chapter 3. Recall that according to Evans-Pritchard (1937, p. 24), the Azande hold contradictory beliefs in full consciousness and they do not seem to follow simple norms like *modus ponens*. They hold that if a man is proven to be a witch then the whole of his clan are witches but they resist drawing this inference in particular cases. However, as I argued in Chapter 3, if we apply the principle of charity we can avoid the radical conclusion that the Azande do not follow basic deductive norms. One way of preserving coherence is by drawing a difference between different meanings of the term 'witch'; between having an active witchcraft substance and having a cold one. Moreover, the conclusion that the

Azande do not follow the principle of non-contradiction or the basic principle *modus ponens* should be resisted on the grounds that in other instances they prove to be capable of following these norms and drawing valid inferences. For example, they are capable of applying to particular cases the general principle that if someone suffers an unexpected accident then they are the victims of witchcraft. Also, they are able to infer that if the oracle says that someone is a witch then certain rituals should be performed and if the ritual treatments do not work then either the oracle was wrong or the witchcraft was too powerful. Similarly, once they determine that someone is a witch they also reject the claim that he is not a witch. Identifying someone as a witch and trying to neutralize his witchcraft is an important social practice and, thus, being consistent about who is a witch is essential. In consequence, it seems that our understanding of the Azande's beliefs about witchcraft depends on implicitly attributing to them minimum deductive abilities.

3. The objection from normative force

If we accept the way Cherniak and Davidson develop the idea that some inferential norms are constitutive of belief then we regard these norms as playing a descriptive role: someone must meet some of these logical requirements in order to be described as a believer. Now, the question arises: is this descriptive function compatible with a normative one? Are there prescriptions associated with these constitutive logical relations? Katrin Glüer and Åsa Wikforss argue that the constitutive role of these logical relations is incompatible with their playing a prescriptive role:

> This further point can be put in terms of the notion of *internal relations*. The idea is that beliefs stand in basic internal relations to one another, such that being a believer requires that certain general patterns of very basic rationality are instantiated by one's beliefs. However, to say that beliefs stand in various internal relations to one another is not to say that these connections are *normative*. On the contrary, there is a clear sense in

> which states or performances that are internally related *cannot* stand in normative relations. To illustrate this, let us for a moment simplify considerably and examine an isolated example: assume that there in fact is an internal relation between the beliefs that *p* and that *if p then q* on the one hand, and the belief that *q* on the other. Then, the claim is, it is *not* the case that a subject who believes that *p*, and believes that *if p then q,* ought to believe that *q* (Katrin Glüer and Åsa Wikforss 2009, p. 49).

This objection applies to both Cherniak's and Davidson's views, although the former is not committed to the idea that some simple inference like *modus ponens* always plays a constitutive role. As indicated, in his view, a cluster of inferences are essential for belief and, depending on the context, they may be either simple or complicated. The cluster view is affected by the objection from normative force because, once it is settled that a cluster of inferences are essential for counting an agent as minimally rational, those inferences are precluded from playing a normative role.

Glüer and Wikforss are aware of the fact that people can make basic logical mistakes, that actual people's reasoning capacities are far from ideal, but they formulate two arguments why this fact does not support the prescriptivity view. First, they argue that there might be further conditions which can be specified such that the agent may become fully aware of the internal connection between, for instance, believing *p* and *if p then q* and believing *q*. They write: "After all, the question whether to believe *q* is simply settled once I am completely awake, fully aware of all the relevant beliefs at the same time, and give it all my attention—provided, at least, that the relevant beliefs are not too numerous or too complicated. No prescription is required to establish the connection, and no prescription can be violated." (Katrin Glüer and Åsa Wikforss 2009, p. 50).

Second, Glüer and Wikforss argue that if logical norms are violated only when the subject does not pay enough attention and is not conscious of his errors then they are not genuine prescriptions. For instance, the prescription that one

should not smoke in public institutions is a genuine prescription because one can knowingly break it or follow it. However, this does not seem to be the case with a norm like *modus ponens*: "the same holds for any putative prescription to the effect that subjects believing that *p* and that *if p then q* ought to believe that *q*. Assume that there is such a prescription. The principle that ought implies the possibility of violation requires that it could be violated even by a calm, fully awake subject with rather simple states of mind. But it cannot." (Katrin Glüer and Åsa Wikforss 2009, p. 51).

Thus, as in the case of the norm of truth, there is an apparent conflict between norms of reasoning being both constitutive and prescriptive for belief. The fact that logical rules like *modus ponens* have a constitutive character implies that the beliefs that *p* and *if p then q* are constitutively, internally connected with the belief that *q*. Then, it follows that someone cannot have the first two beliefs without having the third one. We cannot prescribe what is necessary.

4. A solution: inferential commitments

As in the case of the relation between belief and truth, I think we can look at the notion of doxastic commitment in order to show how rules of reasoning can be both constitutive and prescriptive for belief. It might be argued that undertaking an inferential commitment is constitutive of having a belief in the sense that one *cannot* have beliefs without undertaking an inferential commitment. In other words, the necessary or internal relation is between having a belief and making an inferential commitment rather than between beliefs themselves. This suggestion is a natural consequence of the account of belief defended in chapter two. The main thesis was that belief essentially involves a commitment to truth and that this commitment parallels the one undertaken in making an assertion. An important aspect of making an assertion, it was suggested, following Brandom, is putting the claim forward as a premise for the others to use in their inferences. Given that the claim is considered true, its truth will be preserved by commitment-preserving

inferences. Thus, the asserter *authorizes* others to base their theoretical and practical inferences on his claim. Undertaking a commitment to truth of *p* thus involves being committed to the truth of what follows from *p* as a result of a commitment-preserving inference. Furthermore, being committed to the truth of *p* involves not being entitled to accept the claim that not-*p*.

Given the parallel between belief and assertion, it follows that believing *p* involves a commitment to what follows deductively from *p*. The necessary or constitutive relation is between having beliefs and undertaking this inferential commitment. However, this commitment is prescriptive and thus we have a way of reconciling the constitutive and prescriptive aspect of inferential norms for belief. However, the idea that we ought to believe the logical consequences of our beliefs has been subject to intense criticism. In the following I will discuss those critical points and refine the account of inferential commitment in light of them.

Let us simplify and focus on the paradigmatic rule of reasoning *modus ponens*. The suggestion is that there is a constitutive relation between believing *p* and *if p then q* and undertaking a commitment to believe *q*. In other words, if someone asserts that *p* and that *if p then q* he is committed to the claim that *q* (Brandom 1994; 2008). However, there is an obvious objection against this proposal. First, it is possible that an agent believes *p* and *if p then q* and he realizes that *q* is false. For instance, someone may believe that if she opens the kitchen closet she will find a box of Cheerios and that she opens the closet. If she finds that there is no box of Cheerios it would not be rational for her to form the belief that there is one based on *modus ponens*. In this case we would not say that the agent ought to believe *q*, but rather that he ought to revise[32] at least one of her

[32] One difficulty here is that the obligation to revise at least one of the initial beliefs might be too demanding. After all, it is not always clear which beliefs are false and should be revised and the process of identifying the mistaken beliefs may require a lot of time and cognitive resources. Thus, one apparent solution is to formulate the obligation by using the word 'reevaluate' rather that 'revise'. The logical commitment would thus be to reassess one's own initial beliefs with the goal of finding alternatives to them. However, I do not think this suggestion is satisfactory because it implies that if someone reevaluates his own initial beliefs but does not find anything wrong with them then he is as he should be, although he rejects their logical consequence. I think it is more plausible to hold that his belief-set is still not in order. In consequence, I will defend the stronger

two initial beliefs (Harman 1986). John Broome (1999) takes these counterexamples to show that we should reconsider the form these logical requirements should have. The suggestion that we are considering now is that the requirement associated with *modus ponens* has the following form:

(F1) $(Bp \ \& \ B(p \rightarrow q)) \rightarrow OBq$

This requirement has the form of a conditional. It follows that if someone satisfies the antecedent of the conditional, has beliefs of the form *p* and *if p then q*, then he ought to believe *q*. The operator signifying 'obligation' has narrow-scope in the sense that it applies only to the consequent of the conditional. In this case, Broome observes, the obligation to believe *q* is detachable from the conditional if we know that the antecedent of the conditional is true. However, as Harman points out, it is possible that the agent finds out the *q* is false so she should not believe it.

In order to avoid problems of the kind described above, Broome suggests that we should interpret the rational requirements associated with rules of reasoning like *modus ponens* as having wide-scope, that is, as applying to the whole conditional:

(F2) $O((Bp \ \& \ B(p \rightarrow q)) \rightarrow Bq)$

The advantage of (F2) over (F1) is that the obligation is non-detaching. If someone accepts the norm captured by (F2) and he happens to believe *p* and *if p then q,* then it does not follow that he ought to believe *q*. In other words, the obligation operator in (F2) is non-detaching because it applies to the whole conditional, not only to its consequent. (F2) only means that one ought to see to it that *if* one believes *p* and *if p then q* then one believes *q*. If an agent happens to believe *p* and *if p then q* and he wants to follow the norm expressed by (F2), he is

formulation of the rule and describe cases when one does not know which beliefs to change as cases in which one fails to discharge his inferential commitment.

not obligated to believe *q* but he has the option of revising one or both of his initial beliefs. What the norm forbids is this pattern: {P, P→Q, ~Q}. But, obviously, there is more than one way in which someone can avoid holding these three beliefs at the same time.

In consequence, we should construe the logical commitments or obligations involved in having beliefs as having wide-scope rather than narrow-scope. If an agent does not hold true a deductive consequence of his beliefs then he ought to revise his initial beliefs because he cannot hold them true while rejecting their logical consequence. For instance, suppose someone believes that if it is raining then streets are wet and that it is raining while he does not accept the logical conclusion that the streets are wet. In this case, we think that he is under the obligation to revise at least one of his initial beliefs because the fact that he does not accept their logical consequence implies that he is not fully committed to their truth. Thus, one plausible way of describing the logical commitments involved in having beliefs is as commitments with a disjunctive form: either accept the logical consequence of your beliefs or revise your initial beliefs (Broome, 1999; Millar 2004).

In reply to the objection from normative force, 'normativists' can argue that there is a constitutive relation between having beliefs and undertaking the inferential commitments involved in having those beliefs while, at the same time, emphasizing that the notion of commitment is a prescriptive notion. For the notion of inferential commitment to be genuinely normative there ought to be cases in which an agent does not discharge his commitments and these cases should not be in conflict with the constitutive role played by these inferential commitments. In the case of relations of logical consequence, the suggestion is that an agent does not discharge his inferential commitment when he does not believe what follows deductively from his beliefs but he is not ready to revise any of his initial beliefs. This situation occurs when the agent still perceives his initial beliefs as true and well justified. For instance, by starting from intuitive premises about the deterministic nature of the physical universe someone may reach the

conclusion that we have no free will, given that our actions are just the effects of the laws of physics coupled with physical states of the universe we have no control over. However, while he may regard the disturbing conclusion as deeply implausible he may be at a loss as to which of the initial premises he should reject. In this case the agent is under normative pressure for failing to discharge his logical commitments.

Similarly, there is a constitutive relation between a belief that *p* and a belief that not-*p*: someone cannot believe *p* if he does not *understand* that his commitment to *p* being true forbids him from endorsing not-*p*. As Brandom puts it, if an agent undertakes a commitment to *p* then he is not entitled to undertake a commitment to not-*p* (Brandom 1994). However, let us suppose that someone discovers that he believes both *p* and not-*p* and that he has good reasons to hold both beliefs. This is a case in which the agent breaks an inferential norm in the sense that he has incompatible beliefs and he is under normative pressure to revise them. But it may not be straightforward to him how to solve this tension in his belief system and the revision may take a lot of time and cognitive resources. Nonetheless, the fact that he *understands* that there is something wrong with his beliefs is all that the constitutive relation requires. Furthermore, the fact that he did not discharge the commitments implicit in his beliefs shows that the notion of logical commitment has genuine normative force. This example also demonstrates that the prescriptive force involved in the notion of logical commitment is not in conflict with the constitutive relation between beliefs and logical commitments.

5. The objection that 'ought' implies 'can'

As indicated in the first section, one standard objection against normative accounts of rationality is that they construe the norms of reasoning as they apply to ideal agents with perfect logical abilities. However, there is a gap between such ideal agents and finite human beings with limited time and cognitive resources. Thus, if we accept the intuitive premise that it makes sense to talk about a norm

being in force for an agent only if that agent is capable of following it, it follows that the ideal norms of reasoning are not in force for finite agents because they cannot follow those norms.

There is a conflict between the principle that 'ought' implies 'can' and the idea that belief necessarily involves an inferential commitment. The present proposal is that an agent is subject to two main rational requirements. First, he ought either to believe the logical consequences of his beliefs or to revise his initial beliefs. Second, if an agent believes that *p* then he is not entitled to believe not-*p*. One problem with this suggestion is that a finite agent does not have the cognitive power to compute all the logical consequences of his beliefs and to decide whether he should either embrace these consequences or reassess his belief-set. Similarly, a real agent cannot hold all his beliefs before his mind and decide whether he is entitled to add a new belief to his belief-set.

In response to this objection, 'normativists' may insist that there is a difference between undertaking a commitment and acknowledging a commitment (Brandom 2004). Undertaking a commitment is explained in terms of the attribution of commitment: an agent undertakes a commitment when he does something that entitles his interlocutor to attribute him a certain commitment. Acknowledging a commitment is attributing a commitment to oneself. In many cases, speakers undertake commitments that they fail to acknowledge, given that they are not perfectly rational. However, this reply is unsatisfactory because it presupposes that we are entitled in holding a finite agent responsible for something he cannot do. But the principle that 'ought' implies 'can' advises against attributing to an agent a commitment the fulfillment of which is beyond his cognitive capabilities.

Another way of responding to the objection described above is by construing an agent's inferential commitments in such a way that he is able to follow them. For instance, it might be argued that an agent ought to eliminate an inconsistency once he becomes *aware* of it. Similarly, if someone realizes that his beliefs imply a belief *p* then he has to either believe *p* or change his original

beliefs. However, this does not imply that agents can be totally ignorant with respect to rules of reasoning and that they incur inferential commitments only when they acquire explicit knowledge of these rules. As Cherniak and Davidson emphasize, an agent should display a minimal deductive ability in order to make it possible for us to understand and predict his actions; to interpret him as a cognitive agent (Cherniak, 1986; Davidson 2004).

Similarly, an agent cannot choose to revise his beliefs all the time. This becomes clear if we consider Brandom's idea that rationality consists in the ability to participate in the game of giving and asking for reasons (Brandom 1994). According to this view, in order to count as a participant in the game of giving and asking for reasons an agent cannot be massively inconsistent and he, usually, ought to be aware of and accept *some* of the consequences of his beliefs. Otherwise, if someone is never aware of the most obvious consequences of his beliefs then he loses his authority as an asserter; his interlocutors will not adopt the asserter's claims as premises for their inferences. Similarly, someone who makes only incompatible claims ends up being unable to make claims because his interlocutors are unable to attribute commitments to him.

Nevertheless, it might be objected that the proposed inferential obligations are too weak. Since they are in force only when the agent *becomes aware* of the consequences of his beliefs or of an inconsistency in his belief-set, it follows that the agent is not subject to any norms as long as he actively avoids such awareness. Examples of such epistemically irresponsible actions are not hard to find. For instance, someone might believe that the Bible is the revealed word of God and also she might believe in the theories advanced in biology and geology regarding the age of the Earth and the evolution of species. Although she might know deep down that there are many unresolved tensions in her belief-system, she might decide to avoid formulating them explicitly and confronting them. Thus, for instance, she might keep away from debates about the relation between religion and science, and from books regarding the dispute between creationism and evolution. Similarly, someone may wholeheartedly believe that killing innocent

people is wrong but also find himself supporting the war on terrorism. Although the agent might be aware that trying to exterminate terrorists involves a lot of collateral damage, he may successfully avoid directing his conscious attention to the resulting inconsistency between his beliefs.

I think these cases are similar to the examples of self-deception discussed in Chapter 1, section 5, and I will propose a similar solution. Like in the case of self-deception, in the examples described above the agent wants to avoid directly confronting a difficult problem and pursues a shielding strategy. There are two possibilities: either the desire to pursue a shielding strategy is conscious or it is unconscious. If the desire is unconscious then the agent cannot be asked to conform to the norms of rationality, since he has no control over his unconscious mind. Such an inferential obligation would be in tension with the tenet that 'ought' implies 'can'. On the other hand, if the agent consciously decides to pursue the shielding strategy then he is, in some sense, conscious of the tension in his belief system and deliberately tries to avoid it. However, given that the agent has awareness of the difficulty, it follows that the inferential norms described in this section apply to him. The norms dictate that the agent should try to solve the inconsistency rather than pursue a shielding strategy. Note that this norm is not in conflict with the 'ought' implies 'can' principle because it does not imply that the agent should always scan his belief-set in search for inconsistencies. Rather, it claims that once the agent becomes aware of an inconsistency or suspects that there is a tension between his beliefs, he ought to solve it, rather than avoid considering it.

There are also cases in which subjects outright deny that there is any conflict between their beliefs. For instance, someone might deny that holding the views expressed in the Genesis to be true is in conflict with scientific findings in biology and geology. I think that in these cases the agent ought to articulate an explanation why his beliefs are not inconsistent. One such explanation, for instance, is advanced by day-age creationists who claim that each of the six-days of creation lasted thousands or millions of years. Put differently, if someone's

interlocutor is able to detect an explicit contradiction in his claims, the speaker cannot simply deny it, but he should either explain why it is only an apparent contradiction or withdraw and reconsider his assertions.

To sum up, given the parallel between the commitment to truth specific of belief and that involved in making an assertion described in chapter two, it follows that a believer ought to accept some of the obvious logical consequences of his claims and reject statements which are explicitly incompatible with them. The description of this minimal deductive ability is vague because it is shaped by the implicit norms governing the practice of assertion; by the way participants establish when a speaker has the authority and responsibility necessary for becoming part of the community of asserters. In addition to displaying this minimal deductive ability an asserter or believer should follow two additional logical norms. First, once the agent becomes aware of a logical consequence of his beliefs, he should either accept it or revise some of his beliefs. Second, when the agent discovers an inconsistency in his belief-set he should revise his beliefs in order to reestablish their consistency. These doxastic norms are not in tension with the fact that we are in a finitary predicament and cannot compute *all* the logical consequences of our beliefs and eliminate *all* the inconsistencies in our belief-set.

6. Counterexamples: the lottery and preface paradoxes

The lottery[33] and preface paradoxes[34] challenge the idea that belief is subject to deductive constraints. I will first present the paradoxes and then discuss some solutions to them based on the assertion view of belief defended in the second chapter. The lottery paradox goes as follows: suppose that there is a fair lottery with 1000 tickets and you have one ticket. Now, given the high probability that you will lose the lottery, you are justified in believing that your ticket is a losing ticket. But the same is true of all lottery tickets: you would be justified in

[33] The lottery paradox was first formulated by Kyburg (1961).
[34] A version of this argument was first advanced by Makinson (1965).

believing that you will lose no matter what ticket you got. However, if you form this belief in connection with all the tickets in the lottery and if belief is closed under conjunction then you end up believing that *all* the lottery tickets will lose. This last belief, however, is not justified given that the lottery is fair and at least one ticket will win. Thus, although you are justified in believing of each ticket that it is going to lose, you cannot believe that of all tickets. In other words, you are not justified in believing the conjunction of your individual beliefs. But the conjunction of your beliefs is one of their deductive consequences. Therefore, the lottery paradox shows that belief is not closed under deduction. Put differently, not only does the lottery player not believe the deductive consequence of his beliefs but he is under no obligation to revise or reassess his initial beliefs. This is why the lottery paradox posses a challenge to the idea that inferential commitments have wide-scope. Moreover, it seems that the lottery player has a set of inconsistent beliefs because he believes of each ticket that it will not win but, at the same time, he thinks that at least one ticket will win.

The preface paradox also challenges both the idea that belief is closed under deduction and that it is irrational to have inconsistent beliefs. Suppose that a historian writes a book which is the product of years of intense research. Now, he knows that these kinds of books are bound to contain errors, no matter how careful the researcher is. Thus, in the preface of the book he naturally asserts that there are some false propositions in his book. In other words, although the author believes all the claims in his book, he does not believe the conjunction of those claims. Moreover, when he makes the preface assertion the set of the author's beliefs becomes inconsistent because he believes all the sentences in his book but also that some of these sentences are false. However, it is argued, the author's beliefs are entirely rational and we feel that there is no normative pressure on him to revise them. On the contrary, we have a strong intuition that the author would be irrational if he chose to be consistent and believe that his book is error free, given all the evidence against this possibility.

In order to answer the challenge posed by these objections we should recall the reasons why belief is construed as subject to logical requirements. Deductive norms govern belief because belief involves a commitment to the truth of the proposition believed and deductive constraints are truth-preserving. Thus, the following critical discussion will just show that if we weaken the force of deductive requirements on belief then we alter the relation between belief and truth. However, as it was emphasized in the first two chapters, belief is defined by its relation to truth. In consequence, if a cognitive state is governed by loose and weak logical requirements, then that state is not a full belief. In the following I will present some answers to the two paradoxes which start from the established premise that belief is a commitment to truth.

Jonathan Adler (2002; 2007), a defender of the idea that there is a significant parallel between assertion and belief, formulates a *reductio ad absurdum* argument in reply to the paradoxes. Then, he argues that we can avoid deriving the absurd conclusion from the paradoxes only if we accept that lottery player's beliefs about his tickets and the historian's preface belief are partial beliefs rather than full beliefs.

The distinction between full and partial belief is central to Adler's response. Full belief or all-or-nothing belief involves full commitment on the part of the subject with respect to the truth of the proposition believed. As it was emphasized in the first two chapters, full belief has the property of being transparent in the sense that the believer, from his first-person point of view, 'sees through the belief' the objective states of affairs represented by it. As Adler puts it, our attitude of accepting the proposition as true is *detachable* from the proposition. Plus, full belief is parallel to unqualified assertion, putting forward a sentence as true. In contrast, partial belief does not involve full commitment to the truth of a proposition, but rather a weaker commitment to the proposition being probable. Partial belief is not transparent, and the proposition thus believed is not detachable from our attitude towards it. Partial belief is connected with qualified

assertions of the form 'It's probable that *p*' or '*p*, but I'm not sure that *p*' and so on.

Now, Adler notes that only full belief is subject to deductive constraints and, if the paradoxes challenge the idea that belief is subject to deductive requirements then the beliefs involved should be full beliefs rather than partial ones. Both in the lottery and the preface paradox, the subject believes a set of propositions but does not believe their conjunction. We can represent this as follows:

(Bp & Bq & Br …..) & B ~(p & q & r ….)

The sentences 'p, q, r' and so on represent the claims about the lottery tickets and the claims in the history book. The subject believes each of those claims but do not believe their conjunction. Given that full belief is transparent and detachable from its content, it follows that, from the point of view of the agent, he believes the following content:

(p & q & r…..) & ~ (p & q & r…)

However, Adler observes, this statement has the form of a contradiction, p & ~ p. But then, given that full belief is connected with unqualified assertion, it follows that an agent can put forward a contradiction, thus committing himself to it being true. However, this is incoherent both from the first-person perspective and from the third-person perspective. First of all, an agent cannot believe, in full consciousness, a statement which is necessarily false. This is in conflict with the transparency of belief, the idea that from the first person perspective a full belief is true. Second, his interlocutor cannot attribute a commitment to the truth of the sentence asserted because it is *impossible* for that sentence to be true.

In consequence, Adler attempts to solve the two paradoxes by challenging the crucial premise that they involve full beliefs. Rather, he maintains, the lottery

player's beliefs about the tickets are partial beliefs or probability claims. Similarly, the historian assigns the preface-assertion a high subjective probability but he is not fully committed to its truth (Adler 2002, p. 201). Adler's suggestion gains some plausibility once we reflect on the fact that the lottery player does not have knowledge about the tickets and the historian does not know that some of the claims in his book are false. Thus, given that assertions are normally used to express knowledge claims (Williamson 1996), it follows that qualified assertions are the proper vehicles for lottery and preface claims. But qualified assertions are expressions of partial belief, not full belief.

In the following I will concentrate on the preface paradox and formulate another argument against its conclusion. I will argue that if we assume the assertion model of belief we can make prominent the fact that the author of the history book is under some normative pressure to revise or qualify his beliefs. If the author asserts all the sentences in his book, then he is committed to all of them being true. This implies, for instance, that if he is challenged with regards to a particular claim he is ready to defend it and show why he was entitled to making that claim. On the other hand, if the author puts forward as true the claim in the preface then he undertakes a commitment to the effect that some of the sentences in the book are false. But then, he does not count as asserting those sentences because asserting implies a commitment to truth and, by making the preface claim, he disavowed such a commitment. Now, one consequence of making the preface assertion is that the author does not have to defend some of the claims he makes in the book, because he is not committed to their being true. However, I think that this consequence is unintuitive because we naturally assume that the author intended his book to be a contribution to the advancement of knowledge with regards to some historical period and it is part of the normal scientific practice that his ideas can be challenged and that he should justify his claims. The point is that if the author wants to participate in the normal scientific practice then he should not make the preface assertion; he is therefore under normative pressure to change or qualify the belief expressed in the preface. If the author wants to

stick to the truth of the preface assertion then he does not authorize the members of his scientific community to rely on the truth of some of the claims in his book, and implicitly, he admits that he was not entitled to making those claims. However, this is in conflict with his intention of making a contribution to our knowledge of history[35].

In reply, it might be argued that the author's claims can be assessed by the other members of the scientific community even if he is not committed to all of them being true (Sorensen 2004). One problem that arises in this case is whether the claims put forward in the book are qualified or unqualified. It might be suggested that the claims in the book are put forward as probable rather than true. But, in this case, the assertions in the book are expressions of partial belief rather than full belief and then the author's set of beliefs is not inconsistent. Partial belief and implicitly qualified assertion are not subject to the same deductive requirements are full belief and unqualified assertion. In this case, the author is rational, for instance, in not believing the conjunction of the claims in his book because probability is not preserved by deduction from multiple premises.

However, I agree with the intuition that the preface assertion does not imply that the author is going to avoid critical discussion of his ideas. But this fact shows that, as Adler points out, in practice, these kinds of claims do not carry a full commitment to truth. Let us reflect on the following case. A historian claims that *some* of the claims in his long book are wrong. 'Some' usually means more than one. Then, supposing that he makes many bold claims and he puts forward some courageous hypotheses, we can infer that he expects about 20% of the claims in the book to be proven wrong. In addition, let us suppose that some of his colleagues manage to corroborate 80% of the claims and theses of the book. At this point, given the author's commitment to the preface assertion, everyone should just conclude that the rest of the claims in the book are false. Any further evaluation or critical discussion of these theses would be superfluous, given the

[35] A similar argument is put forward by Mike Kaplan (1996). Kaplan suggests that asserting whole theories is required for doing science and, therefore, accepting the risk that one may defend something false is necessary for participating in scientific practice (Kaplan 1996, 118).

preface assertion. But this is implausible. Neither the author nor his colleagues will just accept that the rest of the book is wrong and they will continue their assessment of the author's proposed ideas. From this we can infer that the preface assertion did not really carry a commitment to truth and did not function like a genuine assertion – did not aim at transmitting information or knowledge. Thus, either the author actually made a probability claim or he just expressed his intellectual modesty. In either case, the challenge of the paradox to the idea that belief is governed by deductive constraints vanishes.

Finally, let us assume for the sake of argument that the preface assertion represent a claim to truth and knowledge. Does it follow from this that the historian is under no obligation to revise some of his beliefs? How can that conclusion follow if he knows for a fact that some of his views are wrong? It might be argued that, in his present situation, he cannot discern which ones are false, given that they all seem well-grounded to him. However, it does not follow from this that he cannot take some steps in the direction of revising his beliefs. As John MacFarlane notes, if you wanted to obey the logical consequence norm, "you would step up critical examination of these claims. You would do more studies, try harder to embed them in established theory, publish them so that others can scrutinize them, and do all the things a good scientist does. […] This is the only course of action open to you that could conceivably count as seeing to it that your beliefs and disbelieves are revised for coherence." (MacFarlane, 2004, p. 15).

To sum up, the preface paradox represents a challenge to the view that belief is subject to deductive rational requirements only if we construe belief as involving a commitment to truth. Starting from the premise that belief involves a commitment to truth I formulated two critical points against the conclusion of the paradox. First, if the author were really committed to the preface assertion then he would disavow some of the claims made in the book and not engage in their critical discussion. But this is in conflict with our intuition that the author sees his book as a contribution to knowledge and wants other scholars to consider it

seriously. Furthermore, his research colleagues do not interpret the preface claim as a knowledge claim and they diligently assess each of the claims in the book. Thus, given that the preface claim is not taken by any of the participants in the language-game to have the force of a knowledge claim, it is more natural to interpret it as a probability statement, an expression of graded belief. Second, if we assume that the preface statement *is* a knowledge claim, then it follows that the author *ought* to take steps towards identifying the false theses in his book. This is also what normal scientific practice dictates.

7. Belief, truth and logical commitments

The purpose of this chapter has been to specify the horizontal logical norms governing belief starting from the idea that belief, like assertion, involves a commitment to truth. The following two main ideas were suggested. First, following Cherniak, Davidson and Brandom, I argued that in order to count as an asserter or believer, an agent has to have a minimal deductive ability. Given that making an assertion involves authorizing others to draw logical consequences from the claim put forward as true, it follows that the asserter should be aware of *some* of the deductive consequences of his beliefs and be prepared to endorse them. Thus, a commitment to some of the basic logical consequences of one's beliefs is essential for gaining the authority necessary for making assertions. Second, I argued that if an agent becomes aware of the fact that a proposition follows deductively from his beliefs then he *ought* to either believe that proposition or revise some of his initial beliefs. Furthermore, if an agent becomes aware of an inconsistency in his belief system then he *ought* to reassess his beliefs to establish coherence because he is not permitted to believe incompatible statements. The lottery and preface paradoxes call into doubt the idea that belief is subject to these deductive constrains. Following Adler, I argued that the thesis that belief, like assertion, involves a commitment to truth can be used in resisting the conclusion of these paradoxes.

Chapter VII – Full and Partial Belief

In this chapter I will focus on the relation between full belief and partial belief. While full belief is the analog of unqualified assertion, partial belief is expressed in qualified assertions of the form 'I'm pretty sure that *p*', 'It is probable that *p*', 'I suspect that *p*' and so on. Many epistemologists have tried to either reduce full belief to partial belief or, if full reduction fails, show how it can be derived from it. In the first section I will assess two ways of reducing flat-out belief to degrees of confidence. In the second section I will focus on two attempts of demonstrating how flat-out belief can be derived from a more basic notion of partial belief (Foley 1993; Frankish 2004). The central idea of the latter accounts is that full belief is a result of a decision to accept a proposition as true on the basis of the high degree of confidence we have in it. Based on the view proposed in Chapter 2, I will argue that talk of full belief is irreducible to talk of degrees of confidence. First, flat-out belief is intrinsically connected with the speech-act of assertion, which normally implies a claim to knowledge (Adler 2002; Brandom 1994; Williamson 1996). Qualified assertion, the linguistic expression of partial belief, does not carry such an implication. Second, as indicated in Chapter 1, flat-out belief is defined by the property of transparency. An agent believes *p* when he has settled the question *whether p*. However, partial belief shows that the agent has not yet settled the question *whether p*. Hence, graded belief is essentially distinct from full belief. Furthermore, in reply to Foley and Frankish, I will submit that full belief cannot be the result of a conscious decision to accept a proposition as true based on one's degree of confidence in it because pragmatic reasons cannot have a direct influence on full belief.

1. Is full belief a species of partial belief?

In this section I will present and evaluate two ways of reducing full belief to partial belief. As indicated, in contrast with full belief, partial belief admits of

degrees; one can be more or less confident in the truth of a proposition. Degrees of confidence are standarldly measured on a scale of 0-1. The first proposal is that we can define full belief as a belief of which we are certain, a belief of degree 1. Second, full belief may be characterized as a belief that meets a certain threshold of confidence that is lower than certainty. The threshold does not have to be indicated precisely but it should be more than 0.5 because believing a proposition implies assigning it a higher probability than its negation[36].

There is one strong objection against the view that full belief involves a maximum degree of confidence. If fully believing *p* were to involve being certain that *p* is the case then we would not even consider the possibility that we may be wrong. In consequence, given that we are absolutely certain that *p*, we must be willing to bet everything we have on the truth of *p* while getting nothing in return. We should rationally accept the bet because we dismiss the possibility that we are wrong (Kaplan 1996; Maher 1993). However, although there are things we fully believe we are not ready to bet everything we have on them. Fully believing is compatible with having a degree of doubt in situations where the stakes are very high. For instance, although a boy's father may be quite certain that his gun is not loaded, he would still be unwilling to aim it at his son and press the trigger.

The lottery and preface paradoxes represent a challenge to the view that full belief is reducible to having a high degree of confidence. One way of formulating this objection is as a *reductio ad absurdum* of the view. The argument relies on two additional assumptions: the conjunction principle for rational belief and the idea that a contradiction has probability 0. According to the conjunction principle an agent ought to believe the conjunction of his beliefs. Now, imagine a fair lottery with 100 tickets. Regarding each ticket, one can form the belief that it will not win given that the probability of it not winning is 0.99. However, given the conjunction rule, one then has to believe that no ticket of the lottery will win. But since the lottery is fair, the agent must believe that at least one of the tickets

[36] This view is sometimes attributed to Richard Foley (1993). However, as it will be shown in the next section, Foley's account of the relation between full and graded belief is more complex. The view is also traced back to Locke and Hume and it is also attributed to modern philosophers like Chisholm (1957) and Sellars (1964).

will win or, in other words, that it is not the case that no ticket will win. Thus, the agent ends up believing a contradiction, which is irrational given that it has probability 0. In consequence, the threshold 0.99 turns out to be too low for full belief. This conclusion is obviously counterintuitive.

The preface paradox represents a different challenge to the threshold view. The author of the history book has different degrees of confidence in the claims he makes in the book. Let's suppose that all these degrees are above 0.6. Now, the conjunction principle implies that the author ought to believe that conjunction of the claims in his book. However, the probability of this large conjunction may be well below 0.5. But then it follows that the threshold for full belief may be well below 0.5.

A number of philosophers argue that in order to answer these objections we should give up the conjunction rule (Kyburg 1961; Foley 1993; Hawthorne and Bovens 1999). This suggestion is supported by the fact that the conjunction rule is a rule of deductive logic but if we construe full belief as a species of partial belief then full belief is not governed by deduction but rather by rules of probabilistic coherence. Probabilistic logic does not require one to believe the consequences of one's beliefs because the probability of a conjunction is less then the probability of its conjuncts. Thus, it may well happen that the conjunction is below the threshold we choose for belief while all conjuncts considered separately are above that threshold.

One objection against this suggestion is that it is in conflict with our practice of rational argumentation. In particular, it is in conflict with the widespread use of *reductio* arguments (Kaplan 1996). When we use a *reductio ad absurdum* argument against someone's view we show that a false or absurd claim follows from the set of beliefs which constitute his view. We expect someone to revise their beliefs once they understand that their view leads to absurdity. However, once we reject the conjunction principle it is hard to see why the *reductio* argument should have any force. If an agent does not follow the conjunction rule then he is not committed to the conjunction of his beliefs and to

what follows from that conjunction. Thus, he should be undisturbed by the fact that an absurdity follows from his beliefs taken together.

In response to this objection Foley argues that a *reductio* is not always a decisive argument; it has various degrees of strength in different situations. For instance, Foley observes, a *reductio* is more effective when the inconsistency follows from a small number of premises. In this case, it is more likely that at least one of the premises is inaccurate. However, if the absurdity is based on a large number of premises then the argument is not perceived as damaging because, given that we attach degrees of probability to the premises, it is no surprise that their conjunction may be false or entail an inconsistency. Thus, Foley argues, the view that full belief is high confidence is compatible with the *reductio ad absurdum* arguments having a limited applicability. For example, if my degree of confidence in each of the propositions p, q and r is very high then I should be pressured to reconsider these propositions if someone showed me that they entertain a contradiction. This is perfectly consistent with the confidence threshold view of full belief. If I assigned probability 0.8 to each of the three propositions and probability 0.5 to their conjunction, I would be surprised to find out that the probability of their conjunction is actually 0. So, assigning high degrees of belief to a small number of sentences entails a weak commitment to their conjunction. The commitment may only by that their conjunction is not definitely false. However, our commitment to the conjunction of our beliefs becomes weaker and ends up disappearing once the set of our beliefs gets larger (Foley 1993, 167-169).

As mentioned, the proponents of the confidence threshold view of full belief want to reject the conjunction principle as a way of responding to the objection based on the lottery and preface paradoxes. However, some philosophers have argued that the lottery case in particular represents a problem for the threshold view even without the assumption of the conjunction principle. Recall that we usually express our full beliefs by making unqualified assertions and our partial beliefs by qualified assertions. Now, in connection with the lottery

case, many philosophers have argued that the lottery player cannot assert that his ticket will not win without qualification (Maher 1993; Kaplan 1996; Williamson 1996). Some epistemologists explain this fact by reference to the thesis that unqualified assertion is a knowledge claim (Williamson 1996; DeRose 1996). They observe that the standard form of challenge to the assertion that *p* is "How do you know that *p*?" If the lottery player flatly asserts that his ticket will not win then he is putting this claim forward as information or knowledge, implying that he has additional information to the effect that he bought a losing ticket. However, given that he has no such information, it follows that the unqualified assertion is not appropriate in this context.

It follows that, given the account of full belief in terms of assertion described in Chapter 2, what is expressed by the qualified lottery assertion is not full belief but partial belief. This point can be made clear by reference to doxastic transparency – the essential feature of full-belief. Chapters 1 and 2 highlighted the fact that, from the first-person perspective, to believe that *p* is to take *p* to be true or be a fact. Flat-out belief is transparent in the sense that the agent can look through it to the state of affairs it represents. From the point of view of the agent, the question *whether p* is settled. That is, the agent takes himself to know that *p*. But belief transparency can be explained, as it was shown in Chapter 2, only by explaining belief in terms of unqualified assertion because only this type of assertion aims at conveying knowledge or information. Thus, given that our belief that our ticket will not win is a partial belief and cannot only be expressed by a qualified assertion, it follows that this type of belief cannot be transparent and thus is essentially different from full belief.

These remarks bring to light the fact that there is an essential phenomenological difference between full belief and partial belief. In introspection, full belief has the essential characteristic of knowledge: facticity. If an agent fully believes that *p* then it is a fact for him that *p*. However, this quality of full-belief can only be explained by reference to the thesis that believing *p* means giving your assent to the unqualified assertion of *p*; that is, accepting *p* as

knowledge or information. In consequence, given the essential connection between partial belief and qualified assertion, we can infer that partial belief will never have the phenomenal characteristics of full belief. Although we can be very confident in the truth of these beliefs, we will not see them as settled facts. The relation with confidence and probability is essential to the notion of degree of belief. In conclusion, we can infer that full belief is irreducible to partial belief because partial belief cannot have the essential phenomenological feature of full belief: transparency or facticity.

One consequence of the intrinsic connection between full belief and knowledge is that the relations between full belief and justification and action are different from those which characterize partial belief. Consider again the lottery paradox. Since we have a strong intuition that the agent does not posses knowledge that his ticket will not win, it follows that his belief is not fully justified (BonJour 1985). Some epistemologists infer from this and similar cases that purely statistical support is insufficient for full belief (Nelkin 2000). Independently of the value of this suggestion, we can safely argue that there is something wrong with the statistical support for the claim that my ticket will not win, something which precludes the player from having knowledge of the fact. But then, this type of justification is different in kind from the one which figures in a successful knowledge claim. However, if full belief were reducible to partial belief then there would be no difference between these types of justification. The type of warrant that the player has for his belief that his ticket will not win will be very strong, given that it establishes that the probability of the proposition is 0.99. Nonetheless, in ordinary practice, we treat this kind of justification as insufficient for knowledge and the flat-out belief and assertion of the claim. But these features of our epistemic practice and the distinction between types of justification which occurs in our ordinary attribution of knowledge would remain unexplained if we were to equate full belief with belief which comes in degrees.

Furthermore, we can easily imagine someone formulating the prescription that we should not give our assent to propositions which are not fully justified.

Although this is a high epistemic standard to follow, there is nothing incoherent in someone trying to follow this rule. For instance, in his attempt to give a solid foundation to our knowledge, Descartes decided against giving his assent to statements which are not certain. But if full belief is just a species of partial belief, these epistemic prescriptions would be senseless. It would be impossible for us not to believe something that is probable enough. The fact that we fully believe it *is* the fact that we think it is probable (Kaplan 1996).

The close connection between full belief and action is another factor which stops us from claiming that the lottery player fully believes that his ticket will not win. For instance, we ordinarily say that someone did not play the lottery because they did not believe they could win. But, if one believes that the lottery is fair and one plays the lottery then it is problematic to say that one does not believe that one's ticket will win. If he does not have that belief, we are inclined to say: why did he play the lottery in the first place?[37] It might be argued that someone can play the lottery because they believe that there is a chance they may win. But this means that they fully believe that they can win. This belief, however, is in conflict with their flat-out belief that their ticket will not win. Thus, it seems that in lottery cases our ordinary belief ascriptions and our degrees of confidence attributions come apart.

Although reflection on the lottery example reveals many differences between all-or-nothing belief and partial belief, the strongest argument against replacing talk of belief with talk of degrees of belief is the indispensability argument; the argument for the conclusion that full beliefs are indispensable. The main premise of the argument is that if all our beliefs were partial beliefs then we would soon become cognitively overwhelmed (Adler 2002; Foley 1993; Harman 1986). First, reasoning and argumentation would involve constantly adjusting the

[37] It is noteworthy that there are cases in which someone acts out of desperation, without really believing in the success of their action. However, this is consistent with the fact that there is a criterial relation between forming the intention to act in a certain way and believing that that action has the desired consequences. Indicating one's desire and corresponding belief is the standard way of giving reasons for action.

probabilities we assign to different propositions. As it was already pointed out, basic argumentative strategies like *reductio ad absurdum* would have limited applicability and their relevance would depend on contextual factors. In addition, given that partial belief is not closed under deduction, the fact that your interlocutor accepts the premises of your argument and thinks that the conclusion logically follows from them will not mean that he has to accept the conclusion. The probability of the conclusion might be just below the threshold for partial belief. Furthermore, if our minds were filled with partial beliefs then we would make excessive use of qualified assertion. Thus, all the information would be qualified and this would greatly impoverish communication and reflection. Our reasoning would be based on probability assignments and we would have to keep track of all these assignments in order to be able to revise them in light of new information.

Consider the case of a police detective trying to solve a murder case. There are a few witnesses who are pretty sure that they saw the suspect enter the victim's house on the night of the murder. The detective has to keep track of their probability assignments and also to press the witnesses to offer precise numerical values. Then, let's suppose that the doctor informs him that, based on the autopsy, it's probable that the victim has died because of the knife stabs but she was suffering from other conditions which might have contributed to her death. Also, there is a chance that the fingerprints on the murder weapon are not those of the suspect and it's also possible that someone has framed him by spilling his blood on the victim's clothes. The doctor is also pretty sure that there is another set of fingerprints on the knife. Now, let's suppose that the defective finds out that some of the witnesses may be enemies of the suspect. His informer claims that he's pretty sure that they were involved in a bar fight a few years ago. In light of all this, the detective has to readjust all the probability assignments for the statements made by the witnesses. However, there is a real possibility that his informer misremembers, given the probability that his memory is not reliable. Thus, we can see that in a world in which qualified assertion is the only way of transmitting

information, every new development gives rise to new possibilities forcing us to continuously alter our degrees of belief. Any new belief, instead of eliminating possibilities, adds new complications and difficulties. In consequence, in such a universe our working memory will be cluttered and we won't be able to keep track of all the alternatives, let alone process and draw inferences from our belief-set. In addition, recall that in order to draw an inference, the agent should have the probability assignments in front of his mind in order to determine whether the inferred proposition passes the threshold for belief. To do this properly, the agent should keep track of all the evidence because it is based on this evidence that he assigned different probabilities to his beliefs. In consequence, on this scenario, both our short-term memory and our working-memory will be cluttered with qualified information, making it impossible for us to function properly as cognitive agents.

Second, as mentioned in the last section, deciding to act on the basis of degrees of confidence usually involves additional deliberation and calculation of risk (Adler 2002; Frankish 2004). Intentional action is normally the result of the interaction between our beliefs and our desires. This functional role of full belief is partially constitutive of our ordinary concept of belief. On the other hand, the relation between partial belief and action is not as direct and simple. For instance, if two people have a sudden craving for ice-cream and one of them says 'I suspect that there is an ice-cream place two blocks away but I'm not sure', then we would not know for sure what they will do. Maybe they want to be sure of finding an ice-cream place. However, if he would simply say 'There is an ice-cream place down the road' then we would be warranted in expecting them to go there.

To sum up, given that full belief is the analog of simple assertion and that partial belief corresponds to qualified assertion, the cases described above show that there is no straightforward way of reducing full belief to graded belief and that our use of unqualified assertions serves certain important functions, like making knowledge claims and facilitating action, which cannot be realized by the use of qualified assertions.

2. The indispensability of full belief

Although most theorists agree that full belief is indispensable and cannot be eliminated in favor of partial belief, they interpret this fact in different ways. On the one hand, the indispensability of full belief is taken to show that full belief refers to a self-contained phenomenon or practice which cannot be analyzed in terms of degrees of beliefs. This kind of account is usually called a bifurcation account (Adler 2002; Kaplan 1996; Maher 1993). On the other hand, some epistemologists argue that the fact that full belief plays a central role in our doxastic practices is compatible with the possibility that full belief is derivable from graded belief.

Both Richard Foley (1993) and Keith Frankish (2004) argue that our full belief is the result of our decision to accept a proposition as true, usually based on the degree of confidence we have in it. In the following I will present their central arguments in support of this thesis. According to Foley, in our intellectual lives we are often forced to choose between three attitudes we can take toward a proposition: belief, disbelief and suspension of judgment. Usually, our epistemic situation is such that we should take an intellectual stand and make a commitment. Thinking in terms of degrees of confidence is not useful in such circumstances. However, the choice we make is directly connected with the degree of confidence we have in the proposition we are considering. Foley captures this relation by what he calls 'the Lockean thesis': it is rational for S to believe *p* just in case it is rational for S to have a degree of confidence in *p* that is sufficient for belief. It is noteworthy that the degree of confidence which licenses full belief is left unspecified and it depends on the decisions of the epistemic subject. However, the subject is bound by the norm that the probability assigned to a proposition he believes should be higher then the probability of its negation. Thus, the threshold of full belief should not be lower than 0.5 (Foley 1993, p. 170-174).

In Foley's view, all rational agents have the epistemic goal of forming a comprehensive and accurate system of beliefs. This epistemic project has a

practical justification: only accurate beliefs can lead to successful action and a comprehensive system of beliefs offers guidance in a variety of circumstances. However, it is up to the epistemic agent to decide whether he should emphasize accuracy over comprehensiveness or vice-versa. This decision determines what degree of confidence he considers to be sufficient for full belief. If, for instance, the agent assigns a great value to avoiding error then he will need to have a high degree of confidence in a proposition before assenting to it. In consequence, he will need more evidence in support of a proposition before deciding to take it true but this implies that he will have to spend more time and cognitive resources in gathering evidence. By contrast, an agent may value comprehensiveness over accuracy and aim at having an opinion about a wide variety of subjects. In this case, his threshold for full belief will have to be lower (Foley 1993, p. 197-201).

Like Foley, Frankish emphasizes the fact that all-or-nothing belief is the result of accepting a proposition as true for a certain purpose. In order to grasp Frankish's view of full belief, we should first look at his general project. Frankish puts forward the thesis that there are two kinds of belief that should be clearly distinguished and they correspond to two levels of cognition. At the lower level, the type of belief that humans share with animals, has the following features: it is non-conscious, passive, partial and non-verbal. At the higher level, belief is conscious, active, flat-out and verbal (Frankish 2004, p. 23-24).

Frankish's central idea is that flat-out beliefs are acceptances. When we accept a proposition as true we decide to adopt it as a premise in future deliberations. This is why, for Frankish, having a flat-out belief involves adopting a premising policy; the policy is that we should take that proposition as a premise in our reasoning about what to think and what we should do. Now, not every acceptance issues in a full belief. The distinguishing mark of acceptances which result in beliefs is that the agent is prepared to use those propositions as premises in deliberations which are *truth critical with respect to premises* (TCP, for short). These are the kinds of deliberations in which truth is the only criterion for the selection of premises. Thus, for example, a lawyer may accept the proposition that

his client is not guilty for professional purposes but that acceptance constitutes a belief only if the lawyer is prepared to use it as a premise in his TCP deliberations.

The policy adoption which results in full belief may take place either at a conscious level or it may be unconscious. In the first case, as a result of theoretical reasoning we decide to accept something as true and use it in our TCP deliberations. However, more often, this decision takes place at an unconscious level, it is unthinking and automatic. For example, most of us routinely accept what we are told unless we have reason to suspect that the speaker is deceitful and uninformed. Another difference between the two levels concerns the *motivation* we have for the policy. Normally, the motivation for accepting a proposition as a premise in TCP deliberations is based on the high degree of confidence we have in that proposition being true. However, at the subpersonal unconscious level, we may accept propositions for non-epistemic reasons. This level of cognition is not within our control and can be influenced by mechanisms which are not truth-conducive.

Now, Frankish argues, both the conscious and unconscious acceptances are realized or grounded in lower level beliefs and desires. Even if one's acceptance is the result of explicit thinking and decision, in time the control over that premising policy is transmitted down to subpersonal, automatic mechanisms. Adopting the full belief that *p*, for instance, is realized in a partial belief that one adopts *p* as a premise and in the desire to stick to that premising policy. In other words, the acceptance which results in flat-out belief is grounded in lower level, graded beliefs and desires. Our conviction that adopting the policy will have long-term benefits for us makes it desirable.

Flat-out beliefs can also be analyzed, Frankish submits, as intentional dispositions. They are dispositions to act in certain ways. For instance, if I accept the proposition 'Smoking is damaging for your health' as true then this acceptance will influence the way I act. I would be disposed to smoke less often or to warn my friends who are smokers and so on. Surely, the relation between

full belief and action is mediated by desires, but Frankish's point is that this relation can be captured in terms of higher-level intentional dispositions realized in lower-level graded beliefs and desires.

Frankish's view concerning the conscious adoption of premising policies is very similar to Foley's. They both argue that some flat-out beliefs are the result of our conscious decisions. These decisions are based both on our high degree of confidence in the respective proposition and on the desire we have to adopt a policy in connection to it. One initial problem that can be raised in connection with this view is whether the decision to fully believe a proposition involves other full beliefs. This problem takes the form of a dilemma. If the decision to accept a proposition as true involves some background beliefs which are full and are the results of other decisions then the account faces the problem of regress. On the other hand, if the background beliefs involved in making a decision about belief are partial beliefs then it is not clear that the acceptance view of belief avoids the problem of cognitive overload. To illustrate the latter point, suppose that you consider who is guilty of a murder and you want to make a decision about it based on the evidence. This decision is based on many background beliefs. For instance, you should believe that the situation is one which requires making up your mind. Plus, you should believe that what is considered to be evidence is actual evidence rather than something designed to mislead the investigators. The fingerprints of the murder weapon are those of the criminal and it is not the case that the real criminal just used someone else's fingerprints. In addition, you have to accept that the suspect you want to put in jail today is the person that you interviewed yesterday and not someone who looks just like him. Thus, these doxastic decisions are based on countless other beliefs. If these background beliefs are conscious partial beliefs then we have the problem of cognitive overload. As Foley mentions, we have to ignore many possibilities in order to be able to reach a decision. That is, we have to decide to ignore those possibilities and take their truth for granted. But then, in order to avoid the problem of infinite regress,

defenders of this view should accept Frankish's suggestion that most of these doxastic decisions take place at a non-conscious level.

The point of my remarks is that, at the conscious level of cognition, full belief must be more fundamental than partial belief. Otherwise, we cannot avoid the problem of cognitive overload. In the following I will formulate a different argument for this conclusion. As it was pointed out, the linguistic analog of full belief is unqualified assertion and the linguistic expression of partial belief is qualified assertion. But it is clear that the ability to make qualified assertions depends on the ability to make unqualified ones. One cannot claim '*p*, but I don't know it' or '*p*, but I'm not sure about it' if one cannot assert *p*. One should understand the commitment involved in simple assertion and be a participant in the practice of giving and asking for reasons in order to grasp what it means to qualify a claim and make it more specific. Qualifying a claim implies knowing that it normally involves a commitment to truth and to proving entitlement to truth and explicitly indicating to the hearer that no such full commitment is made. Thus, on this picture, the practice of making commitments more specific makes sense and is derivative from the core practice of simple assertion.

There is another objection which can be raised against both Foley's and Frankish's views. They both presuppose that beliefs can be influenced by pragmatic factors. In Foley's view we realize that we want a black-and white picture of the world, a system of beliefs that is both accurate and comprehensive, and this epistemic goal makes us take intellectual stands. As a result of doxastic deliberation, we can undertake a commitment to the truth of a sentence although we know that the evidence in its favor only makes it probable. Similarly, for Frankish, premising policies are influenced by pragmatic factors. However, these characterizations of doxastic deliberation are in conflict with the idea that such deliberation exhibits transparency. Transparency implies that only evidential factors influence belief in full consciousness; that we cannot, in full awareness base our opinions on pragmatic factors (Shah 2003; Adler 2002). As indicated in Chapter 4, in the context of discussing Frankish's argument in support of doxastic

voluntarism, David Owens (2000) formulates an argument against both pragmatism and evidentialism about belief. In his view, we make up our minds that *p* because we find the reasons for *p* to be convincing, to establish its truth. However, the evidence appears convincing in a specific context of deliberation and there are limits to the amount of time and cognitive resources we want to invest in deliberation about belief in that context. Although deliberation has this practical aspect, Owens points out that we cannot reach a certain conclusion as a result of reflection on such matters. We cannot end up believing *p* based on inconclusive evidence just because we realize that we do not have enough time to collect more evidence. In that case we can just accept *p* as a hypothesis but not as a belief.

However, Foley and Frankish might reply that we never have conclusive evidence for our beliefs. If doxastic deliberation results in fully believing the truth of *p* then the agent should believe that he has conclusive evidence for *p*, evidence which guarantees or necessitates *p*'s truth beyond any doubt. But we rarely have such evidence for our beliefs. For instance, the requirement for conclusive evidence might be construed as implying that statements about our evidence should deductively imply the conclusion. But this is in conflict with the fact that we often base our beliefs on inductive or abductive pieces of reasoning.

In the following I will present a few interrelated ways of replying to this objection. Adler, for instance, denies that only deductive inferences guarantee truth. Even when we base our conclusions on inductive reasoning, we take our evidence to be conclusive. For example, seeing his dad's car pull into the driveway may constitute a conclusive reason for a kid to form the full belief that his dad is home. Numerous times in the past, the sight of the car was followed by the appearance of his dad. In this context, the son does not reason that it's very probable that his dad arrived, but he knows that his dad is home. How can we explain that? Owens, for example, argues that the standards of evidence vary from one context of deliberation to another. Thus, the pragmatic factors which shape deliberation dictate how much evidence is conclusive in that specific context. In

his view, these pragmatic factors operate in the background of our reasoning and they give us the impression that we have conclusive grounds for full belief, but when we reflect on the evidence outside of a context, we realize that it is insufficient evidence. This is why, Owens concludes, we do not form beliefs by reflecting on the quality of our reasons or evidence but our beliefs are the results of the force of reasons in a particular context, not of rational appreciation of those reasons.

One of Frankish's suggestions is convergent with this observation. Frankish claims that many of our basic full beliefs are the results of decisions we make at a non-conscious, subpersonal level; that they are realized in non-conscious graded beliefs and desires. Whether this is the case is partly a matter of empirical research but I will make an argument against the idea that these non-conscious, non-verbal, partial states can have explanatory priority over linguistic full belief. It is noteworthy that Frankish describes the work of these systems by using words like 'premising policy', 'decision making', 'policy commitment' and so on. For instance, he argues that when we adopt a policy we make a commitment to stick to it because it is desirable. However, it appears that the home of the act of adopting a policy is the public language and our social practices. Moreover, a decision is based on reasons and the home of the concept of reason is our public practice of giving and asking for reasons. Reasons can be offered only in the form of assertions. Similarly, adopting a policy is an intentional action and an action is something for which reasons can be asked. My point is that we can understand all these subpersonal automatic processes only on the model of the social linguistic practices in which we participate, specifically, the basic game of giving and asking for reasons. The intentionality of subpersonal mechanisms is derivative intentionality and it can only be understood by linguistic creatures who posses original intentionality. But, as it was demonstrated in the second chapter, simple assertion plays a fundamental role in these discursive practices. Thus, if simple assertion is the expression of full belief, it follows that full belief plays such a fundamental role.

The puzzle we are trying to solve is that when we fully believe something we take it to be true but when we reflect on the evidence we have we usually realize that our evidence does not guarantee truth. It appears that, contrary to Foley and Frankish, we cannot explain this gap between what the evidence warrants and our beliefs in terms of the conscious or unconscious decisions of the subjects to take epistemic risks. Are there other solutions? My suggestion is that the epistemic risks are embedded in the practice of assertion. Taking a strong inductive argument to establish the truth of a conclusion makes sense from an evolutionary perspective. It saves us time and cognitive resources by not forcing us to keep track of probabilities. This is not a decision we each make for ourselves when we think inductively, but it is the way the practice of assertion got shaped by evolution, because this form of cooperative linguistic behavior proved to be more efficient. Inductive reasoning is based on our observation of regularities in the environment and it results in expectations we form regarding how things will unfold. These expectations form the basis of our planning. To illustrate, suppose that the members of a tribe notice that a nearby river floods every year as the wet-season rain fills it. As a result, when they decide to build a settlement they want to construct it at a certain distance from the river. However, suppose that one of the members of the tribe is skeptical about whether the river will flood again, although he offers no reason why the future will be different from the past. It seems clear that this kind of skepticism does not lead to fruitful or successful cooperation and thus, in the long run, it becomes a useless conventional move and is selected out. By contrast, reasoning based on strong induction normally results in successful cooperation and planning because it leads to the formation of reasonable expectations. As a result, this form of argumentation becomes stabilized and turns into a standard way of proving entitlement.

3. Conclusion

In this chapter I argued for the thesis that full belief and partial belief are distinct, irreducible phenomena. Full belief is transparent and essentially connected with the more complex state of knowledge. Partial belief, on the other hand, shows that the agent did not make up his mind on a specific subject matter. His evidence was not convincing or conclusive. This is why, from his perspective, the issue is not settled and the cognitive state he is in is not transparent.

It is noteworthy that, like full belief, partial belief is prescriptive. When an agent puts forward the claim that "It is probable that *p*", he ought to demonstrate why his evidence entitles him to have a high degree of confidence in the truth of *p*. Since logical consequences of *p* do not have lower probability it follows that the agent should accept these consequences with at least the same degree of confidence. Moreover, if *p* is probable then not-*p* is improbable. Hence, the agent should not be confident that not-*p*, on pain of inconsistency. As the norms governing full belief, the rules of graded belief are originated in our public practices of communication and argumentation.

Conclusion

The main thesis defended in this dissertation is that conscious beliefs are essentially prescriptive. They are defined by the fact that they 'aim at truth' and by their logical relations with other beliefs. These constitutive relations have normative aspects. First, conscious belief involves a commitment to truth, and the agent ought to demonstrate his entitlement to that commitment, if appropriately challenged. Second, if the agent is aware of the fact that his belief that *p* logically implies the belief that *q*, then he should either accept *q* or revise his initial belief. Third, if an agent believes that *p* then he is not entitled to accept as true a proposition incompatible or inconsistent with *p*. As it was argued, this observation implies that once the subject becomes aware of an inconsistency within his belief-set, he ought to solve it, since he cannot be entitled to incompatible commitments.

These obligations play a constitutive role in the sense that if someone does not grasp them then he has an incomplete understanding of the concept of belief. For instance, if a speaker does not comprehend that he has to offer reasons in support of his beliefs and that they are subjects to requirements of consistency and coherence then we say that he does not get the ordinary notion of belief. Put differently, *an adequate explanation of the meaning* of the ordinary notion of belief should make reference to the norms of entitlement and inference.

In this dissertation I also argued for the claim that we can understand the normativity of conscious belief only by examining its relation with the linguistic practice of assertion. I defended the thesis that conscious, deliberative belief is just the acceptance of an assertion *in foro interno*. In other words, deliberation about belief is only inner-speaking, the internalization of the capacity to communicate with others. Private reflection is governed by the same norms as public communication and argumentation. In the social setting, the subject ought to offer reasons for the claims he puts forward as true or accepts as true. Also, he has to be able to challenge his interlocutor's claims and evaluate their responses. The subject's private cogitation counts as doxastic deliberation only as long as it

resembles public debate or argumentation. Put differently, if his inner thoughts were made public, the members of his linguistic community would recognize them as the expression of an argument; they would meet the criteria of coherence and relevance which shape social communication. If the agent's inner process does not meet these criteria then we claim that he was not deliberating what to believe.

This work examined the nature of conscious belief. This type of belief has different properties from unconscious belief, the kind of informational state we normally attribute to non-human animals or small children. My argument was that the only way in which we can explain the features of conscious belief is by studying its relation to language-use. But this argument implies that we can offer an account of the emergence of the linguistic practice of assertion in terms of lower-level or animal belief. Developing such a theory was beyond the scope of this thesis but it is a line of research worth pursuing. In Chapter 2 I suggested, following Millikan, that the practice of assertion is a type of coordination convention which got selected for its benefit: sharing information. Next, it needs to be demonstrated that lower-level cognition is sufficient to allow non-human animals to engage in this type of cooperative practice.

Bibliography:

Adler, J. (2002) *Belief's Own Ethics*, MIT Press.

Adler, J. and Armour-Garb, B. (2007) "Moore's Paradox and Transparency" in *Moore's Paradox*, eds. Mitchell Green and John Williams, Oxford: Clarendon Press, p. 146-162.

Anscombe, G.E.M. (1957) *Intention*, Oxford: Basil Blackwell.

Baldwin, T. (2007) "The Normative Character of Belief", in *Moore's Paradox*, eds. Mitchell Green and John Williams, Oxford: Clarendon Press, p. 76-89.

Bennett, J. (1964) *Rationality*, London: Routledge and Kegan Paul.

Bloor, D. (1991) *Knowledge and Social Imagery*, second ed., The University of Chicago Press.

Boghossian, P. (2003) "The normativity of content", *Philosophical Issues*, 13, 31-45

——— (2006) *Fear of Knowledge*, Oxford University Press.

Brandom, R. (1994) *Making it Explicit*, Harvard, Harvard University Press.

——— (1997) "Replies" *Philosophy and Phenomenological Research*, Vol. 57, No 1, p. 189-204

——— (2008) *Between Saying and Doing. Towards an Analytic Pragmatism*, Oxford, Oxford University Press.

Bratman, M. (1987) *Intention, Plans, and Practical Reasoning*, Cambridge, MA: Harvard University Press.

Broome, J. (1999) "Normative Requirements", *Ratio*, 12, p. 398–419. Reprinted in Dancy (2000) p. 78–99.

Bykvist, K. and Hattiangadi, A. (2007) "Does *thought* imply *ought*?" *Analysis,* 67: 277- 285.

Buleandra, A. (2008) "Normativity and Correctness: a reply to Hattiangadi", *Acta Analytica*, Vol. 23, No. 2. p. 177-186.

Carruthers, P. (1996) *Language, Thought and Consciousness*, Cambridge: Cambridge University Press.

Cherniak, C. (1986) *Minimal Rationality*, Cambridge, MA: MIT Press.

Christensen, D. (2004) *Putting Logic in Its Place*, Clarendon Press, Oxford.

Dancy, J. (2000) ed., *Normativity*, Oxford: Blackwell

Davidson, D. (2001a), *Essays on Actions and Events*, Oxford: Clarendon Press, 2nd edn,

——— (2001b) *Inquiries into Truth and Interpretation*, Oxford: Clarendon Press, 2nd edn.

——— (2004) *Problems of Rationality*. Oxford University Press.

Dennett, D. (1987) *The Intentional Stance*, Cambridge, Mass.: MIT Press.

——— (1991) *Consciousness Explained*, Boston: Little, Brown and Company.

de Santillana, G. (1955) *The Crime of Galileo*, University of Chicago Press.

Dretske, F. (2000) "Norms, History and the Constitution of the Mental" in *Perception, Knowledge and Belief*, Cambridge University Press.

Engel, P. (2001) "Is Truth a Norm?" in *Interpreting Davidson*, ed. P. Kotatko et al., Stanford, CA; CSLI Press.

—— (2005) "Truth and the Aim of Belief", in D. Gillies, ed., *Laws and Models in Science*, King's College, London, 2005, pp. 77-97.

——— (2006) “Epistemic norms and rationality” in W. Strawinski, ed. *Festsschrift for Jacek Jadacki*, Springer.

——— (2007) “Belief and Normativity”, *Disputatio*, Vol. II, No. 23, p. 179-203.

Evans-Pritchard, E. E. (1937) *Witchcraft, Oracles and Magic among the Azande*, Oxford: Clarendon Press.

Foley, R. (1993) *Working Without a Net*, Oxford University Press.

Frankish, K. (2004) *Mind and Supermind*, Cambridge University Press.

Frankish, K. (2007) “Deciding to Believe Again”, *Mind*, vol. 116: 523-547.

Gendler, T. S. (2008) “Alief and Belief” *Journal of Philosophy*, 105 (10), p. 634-663.

Glüer, K. and Pagin, P. (1999) “Sense and Prescriptivity”, *Acta Analytica*, 14, p. 111–128.

Glüer, K. and Wikforss, Å. (2009) “Against Content Normativity”, *Mind,* Vol 118, 469, p. 31-70.

Grice, P. (1989) *Studies in the Way of Words.* Cambridge: Harvard University Press.

Harman, G. (1986) *Change in View*. MIT Press.

Hattiangadi, A. (2007) *Oughts and Thoughts*, Oxford, Clarendon Press.

Hawthorne, J. and Bovens, L. (1999) “The Preface, the Lottery, and the Logic of Belief”, *Mind* 108 (430): 241-264.

Heal, J. (1994) “Moore’s Paradox: a Wittgensteinian Approach”, *Mind*, Vol. 103, 409.

Hieronymi, P. (2006) "Controlling Attitudes", *Pacific Philosophical Quarterly*, 87, 45-74.

——— (2008) "Responsibility for Believing", *Synthese*, 161: 357 – 373.

——— (2009) "Two Kinds of Agency", in *Mental Agency*, Lucy O'Brien, ed. Oxford University Press.

Horton, R. (1967) "African Traditional Thought and Western Science", *Africa*, 38, p. 50-71 and 155-187.

Jeffrey, R. (1992) *Probability and the Art of Judgment*, Cambridge: Cambridge University Press.

Kaplan, M. (1996) *Decision Theory as Philosophy*, New York: Cambridge University Press.

Kavka, G. (1983) "The Toxin Puzzle", *Analysis* 43 (1): 33-36.

Kelly, T. (2003) "Epistemic Rationality as Instrumental Rationality: a Critique", *Philosophy and Phenomenological Research,* Vol. 66, Issue 3, p. 612-640.

Kolodny, N. (2005) "Why be Rational?", *Mind* vol. 114, 509-562.

Korsgaard, C. (1996) *The Sources of Normativity*, Cambridge University Press.

Kripke, S. (1982) *Wittgenstein on Rules and Private Language: An Elementary Exposition*, Oxford, Blackwell.

Kyburg, H. (1961) *Probability and the Logic of Rational Belief*, Middletown: Wesleyan University Press.

Lorenz, K. Z. (1966) *Evolution and Modification of Behavior*. London: Methuen.

MacFarlane, J. (2004) "In What Sense (if any) is Logic Normative for Thought?" Draft presented at the Central Division APA symposium on the normativity of logic. http://johnmacfarlane.net/normativity_of_logic.pdf

Maher, P. (1993) *Betting on Theories*, Cambridge University Press.

Makinson, D. C. (1965) "The Paradox of the Preface", *Analysis* 25, 205-207.

Mercier, H. and Sperber, D. (2009) "Intuitive and Reflective Inferences", in *In two Minds: Dual Processes and Beyond*, eds. Jonathan Evans and Keith Frankish, Oxford University Press.

Millar, A. (2004) *Understanding People*: *Normativity and Rationalizing Explanation.* Clarendon Press, Oxford.

Millikan, R. (1993) *White Queen Psychology and Other Essays for Alice*. Cambridge, MIT Press.

——— (2005) *Language: a Biological Model*, Oxford University Press.

Moran, R. (1988) "Making up your Mind: Self-Interpretation and Self-Constitution", *Ratio*, NS ½, p.135-151.

——— (2001) *Authority and Estrangement: An Essay on Self-Knowledge,* Princeton University Press.

Morton, A. (2003) *The Importance of Being Understood: Folk-Psychology as Ethics*, Routledge.

Owens, D. (2000) *Reason without Freedom: the Problem of Epistemic Normativity,* Routledge University Press.

Papineau, D. (1999) "Normativity and Judgement", *Aristotelian Society Supplementary Volume*, 73, 16-43.

Reisner, A. (forthcoming) "Anchoring Diachronic Rationality".

Rosen, G. (2001) "Brandom on Modality, Normativity and Intentionality", *Philosophy and Phenomenological Research,* 63, (3): 611-623.

Scanlon, T.M. (1998) *What we Owe to Each Other*, Cambridge, MA: Harvard University Press.

Searle, J. (1969). *Speech Acts: An Essay in the Philosophy of Language.* Cambridge: Cambridge University Press.

——— (1983) *Intentionality: an Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.

Shah, N. (2003) "How Truth Governs Belief" *Philosophical Review* 112, 447-482.

Shah N. and Velleman J. (2005) "Doxastic Deliberation" *The Philosophical Review*, 114, 497-534.

Sperber, D. (1982) "Apparently Irrational Beliefs" in *Rationality and Relativism*, Ed. Martin Hollis and Steven Lukes, MIT Press.

Sperber, D. Clement, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G. and Wilson, D. (2010) "Epistemic Vigilance", *Mind and Language*, Vol. 25, No. 4, p. 359-393.

Stalnaker, B. (1984) *Inquiry*, MIT Press.

Steglich-Petersen, A. (2006) "No Norm Needed: the Aim of Belief", *The Philosophical Quarterly*, October, 56, 225, 500-516.

Stein, E. (1996) *Without Good Reason.* Oxford: Oxford University Press.

Steup, M. (2008) "Doxastic Freedom", *Synthese*, 161, 375-392.

Taylor, C. (1982) "Rationality", in *Rationality and Relativism*, Ed. Martin Hollis and Steven Lukes, MIT Press.

Velleman, J. D. (2000) "The Aim of Belief", in *The Possibility of Practical Reason*. Oxford University Press.

Warhaft, S. (ed.) (1965) *Francis Bacon: A Selection of His Works*, Toronto.

Watson, G. (2004) "Asserting and Promising", *Philosophical Studies*, 117, 57–77.

Wedgwood, R. (1999) "The A Priori Rules of Rationality", *Philosophy and Phenomenological Research* LIX, 113-131.

——— (2002) “The Aim of Belief”, *Philosophical Perspectives* 16, 276-297.

——— (2007) *The Nature of Normativity*, Oxford University Press.

Williams, B. (1970) “Deciding to Believe” in *Language, Belief, and Metaphysics*, ed. Howard E. Kiefer and Milton K. Munitz, 95-111. Albany: SUNY Press.

——— (2002) *Truth and Truthfulness*, Princeton: Princeton University Press.

Williamson, T. (1996) “Knowing and Asserting”, *The Philosophical Review*, 105, p. 489-523.

Wittgenstein, L. (1953) *Philosophical Investigations*, G.E.M. Anscombe and R. Rhees (eds.), G.E.M. Anscombe (trans.), Oxford: Blackwell.

Zangwill, N (2005) “The Normativity of the Mental”, *Philosophical Explorations*, 8, 1-19