

representations in certain ways causes other representations to be distributed in other ways. This, I submit, is characteristic of institutions.

Some sets of representations include representations of the way in which the set should be distributed.

An institution is the distribution of a set of representations which is governed by representations belonging to the set itself.

This is what makes institutions self-perpetuating. Hence, to study institutions is to study a particular type of distribution of representations. This study falls squarely within the scope of an epidemiology of representations.

Let me end by illustrating this characterization of an institution with an example. Consider the Malinowski Memorial Lecture. It is, as everyone will agree, an institution. A representation was put on paper when the Lecture was first instituted; unwritten additions were made in the course of time. This representation calls for the yearly distribution of invitations, to a speaker on the one hand, to an audience on the other; it represents the speaker distributing to the audience the complex representation called a lecture; it represents the lecturer including in his lecture some deferential references to Malinowski; it represents the lecturer ending the oral representation after an hour or so, so that the, by then, thirsty audience can go for a drink. It represents the lecturer, a few weeks later, submitting a written version of his oral representation, to the journal *Man*, thus ensuring a wider, more lasting distribution of it. When all these representations have been distributed according to one of them, then you have – or, rather, you have had – a Malinowski Memorial Lecture.

The Epidemiology of Beliefs

I would like to bring together two sets of speculations: anthropological speculations on cultural representations and psychological speculations on the cognitive organization of beliefs, and to put forward, on the basis of these speculations, fragments of a possible answer to the question: how do beliefs become cultural? I will not apologize for the speculative character of the attempt. At this stage, either the question is answered in a vague, fragmentary and tentative way, or it must be left alone: there is not enough sound theorizing and well-regimented evidence in the domain to do otherwise.

Anthropological Speculations

I use 'cultural representation' in a wide sense: anything that is both cultural and a representation will do. Thus, cultural representations can be descriptive ('Witches ride on broomsticks') or normative ('With fish, drink white wine'); simple, as in the above examples, or complex, like the common law or Marxist ideology taken as a whole; verbal, as in the case of a myth, or non-verbal, as in the case of a mask, or multi-media, as in the case of a Mass.

To begin with, two remarks about the notion of a representation.

This chapter is a revised version of 'The Epidemiology of Beliefs', published in Colin Fraser and George Gaskell (eds), *The Social Psychological Study of Widespread Beliefs* (Oxford: Clarendon Press, 1990), 25–44. Reprinted by permission of Oxford University Press.

First, as we saw in chapter 2, 'to represent' is not a two-place predicate: something represents something; it is a three-place predicate: something represents something for someone. Second, we should distinguish two kinds of representations: internal, or *mental representations* – for example, memories, which are patterns in the brain and which represent something for the owner of that brain – and external, or *public representations* – for example, utterances, which are material phenomena in the environment of people and which represent something for people who perceive and interpret them.²⁰

Which are more basic: public or mental representations? Most cognitive psychologists (see Fodor 1975) see mental representations as more basic: for public representations to be representations at all, they must be mentally represented by their users; for instance, an utterance represents something only for someone who perceives, decodes and comprehends it – that is, associates with it a (multi-level) mental representation. On the other hand, mental representations can exist without public counterparts; for instance, many of our memories (and all or nearly all the memories of an elephant) are never communicated. Therefore, it is argued, mental representations are more basic than public ones.

Most social scientists (and also philosophers such as Ludwig Wittgenstein (1953) and Tyler Burge (1979)) do not agree: they see public representations as more basic than mental ones. Public representations are observable, both by their users and by scientists, whereas mental representations, if they exist at all, can only be surmised. More importantly, it is claimed (e.g. by Vygotsky (1965)) that mental representations result from the internalization of public representations and of underlying systems (e.g. languages and ideologies) without which no representation is possible. But if so, public representations must be more basic than mental ones. (This denies mental representations to non-social animals, but holders of this view don't mind.)

There is an obvious sense in which public representations do come before mental ones: a child is born into a world full of public representations, and is bombarded with them from the first moments of her life. She does not discover the world unaided, and then make public her privately developed representations of it; rather, a great many of her representations of the world are acquired vicariously,

not through experience, but through communication, or through a combination of experience and communication. Moreover, her very ability to communicate effectively is contingent upon her acquiring the language and the other communication tools of her community. However, those who see mental representations as basic are not (or should not be) denying this point. What they are (or should be) denying is that public representations could be of use to the child if she did not have, to begin with, some system of mental representations with which to approach the public ones.

Conversely, those who see public representations as basic are not (or should not be) merely making the trivial point that each individual is born into a world full of public representations and crucially relies on them. They are (or should be) claiming not only that the physical shape of public representations is public, outside people's heads, there for people to perceive, but also that the *meaning* of public representations is public, out there for people to grasp. On this view, meaning – the regular relationship between that which represents and that which is being represented – is social before being individually grasped; hence, in the relevant sense, public representations are more basic than mental ones. This leads anthropologists, in particular, to consider that 'culture is public because meaning is' (Geertz 1973:12). Most anthropologists study culture as a system of public representations endowed with public meanings, without any reference to the corresponding mental representations.

I have a bias: I am a materialist – not in the sense this word too often has in the social sciences, where a materialist is one who believes that the economic 'infrastructure' determines the ideological 'superstructure', but in the sense of philosophy and the natural sciences: that all causes and all effects are material. I then wonder: what kind of material objects or properties could public meanings possibly be? I am not persuaded by Geertz when he dismisses the issue thus:

The thing to ask about a burlesqued wink or a mock sheep raid [two of his examples of public representations], is not what their ontological status is. It is the same as that of rocks on the one hand and dreams on the other – they are things of this world. The thing to ask is what their import is: what it is . . . that, in their occurrence and

through their agency, is getting said. (Geertz 1973:10)

I am not persuaded, because the task of ontology is not so much to say which things are 'of this world' and which are not, as to say in what manner, or manners, things can be of this world; and regarding cultural things, the problem stands.

We understand reasonably well how material things fit into the world; but we don't know that there are immaterial things, and if there are, we don't know how they fit. Hence, for any class of entities — rocks, memories, or cultural representations — a materialist account, if it is available at all, is preferable, on grounds of intelligibility and parsimony.

In the case of mental things such as memories or reasonings, cognitive psychologists accept at least a form of minimal materialism, which has precise methodological implications. From a cognitive point a view, an appropriate description of a mental phenomenon must, *inter alia*, show that this phenomenon can be materially realized. For example, cognitive psychologists may try to show how a reasoning process could be materially implemented on a computer. Thus, with the development of cognitive psychology, we begin to grasp what kind of material objects mental representations might be.

Now, when it comes to cultural representations allegedly endowed with public meanings, whether we pay lip-service to materialism and declare them to be material too or resign ourselves to ontological pluralism, the truth of the matter is that we have no idea in what manner they might be 'things of this world'.

The materialist alternative is to assume that both mental and public representations are strictly material objects, and to take the implications of this assumption seriously. Cognitive systems such as brains construct internal representations of their environment partly on the basis of physical interactions with that environment. Because of these interactions, mental representations are, to some extent, regularly connected to what they represent; as a result, they have semantic properties, or 'meaning', of their own (see Dretske 1981; Fodor 1987b; Millikan 1984). Public representations, on the other hand, are connected to what they represent only through the meaning attributed to them by their producers or their users; they have no

semantic properties of their own. In other words, public representations have meaning only through being associated with mental representations.

Public representations are generally attributed similar meanings by their producers and by their users, or else they could never serve the purpose of communication. This similarity of attributed meaning is itself made possible by the fact that people have similar enough linguistic and encyclopaedic knowledge. Similarity across people makes it possible to abstract from individual differences and to describe 'the language' or 'the culture' of a community, 'the meaning' of a public representation, or to talk of, say, 'the belief' that witches ride on broomsticks as a single representation, independently of its public expressions or mental instantiations. What is then described is an abstraction. Such an abstraction may be useful in many ways: it may bring out the common properties of a family of related mental and public representations; it may serve to identify a topic of research. Mistake this abstraction for an object 'of this world', however, and you had better heed Geertz's advice and ignore its ontological status.

From a materialist point of view, then, there are only mental representations, which are born, live and die within individual skulls, and public representations, which are plain material phenomena — sound waves, light patterns and so forth — in the environment of individuals. Take a particular representation — witches on broomsticks — at an abstract level: what it corresponds to at a concrete level is millions of mental representations and millions of public representations the meanings of which (intrinsic meanings in the case of mental representations, attributed meanings in the case of public ones) are similar to that of the statement: 'Witches ride on broomsticks'. These millions of mental and public representations, being material objects, can and do enter into cause-effect relationships. They may therefore play a role both as *explanans* and as *explanandum* in causal explanations. The materialist wager is that no other causal explanation of cultural phenomena is needed.

Consider a human group: it hosts a much wider population of representations. Some of these representations are constructed on the basis of idiosyncratic experiences, as, for instance, my memories of the day on which I stopped smoking; others are based on common experiences, as, for instance, our belief that coal is black; others

still derive from communication rather than from direct experience, as, for instance, our belief that Shakespeare wrote *Macbeth*. Common experience and communication bring about a similarity of representations across individuals; or, loosely speaking, they cause some representations to be shared by several individuals, sometimes by a whole human group. This loose talk is acceptable only if it is clear that when we say that a representation is 'shared' by several individuals, what we mean is that these individuals have mental representations similar enough to be considered versions of one another. When this is so, we can produce a further version – a public one this time – to identify synthetically the contents of these individual representations.

When we talk of cultural representations – beliefs in witches, rules for the service of wines, the common law, or Marxist ideology – we refer to representations which are widely shared in a human group. To explain cultural representations, then, is to explain why some representations are widely shared. Since representations are more or less widely shared, there is no neat boundary between cultural and individual representations. An explanation of cultural representations, therefore, should come as part of a general explanation of the distribution of representations among humans – as part, that is, of an *epidemiology of representations*.

The idea of an epidemiological approach to culture is by no means new. It was suggested by Gabriel Tarde (1895, 1898). Several contemporary biologists have developed it in various ways. The value of an epidemiological approach lies in making our understanding of micro-processes of transmission and macro-processes of evolution mutually relevant. However, if the micro-processes are fundamentally misunderstood, as I believe they have been in previous epidemiological approaches, the overall picture is of limited value. Whatever their differences and their merits, past approaches share a crucial defect: they take the basic process of cultural transmission to be one of replication, and consider alterations in transmission as accidents.

The view of cultural transmission as a process of replication is grounded not only in a biological analogy – a genic mutation is an accident, replication is the norm – but also in two dominant biases in the social sciences. First, as we have already seen, individual differences are idealized away, and cultural representations are too

often treated as identical across individuals throughout a human group or subgroup. Second, the prevailing view of communication, as a coding process followed by a symmetrical decoding process, implies that replication of the communicator's thoughts in the minds of the audience is the normal outcome of communication.

In *Relevance: Communication and Cognition*, Deirdre Wilson and I have criticized this code model of human communication, and developed an alternative model which gives pride of place to inferential processes. One of the points we make – a commonsensical point, really, which would hardly be worth making if it were not so often forgotten – is that what human communication achieves in general is merely some degree of resemblance between the communicator's and the audience's thoughts. Strict replication, if it exists at all, should be viewed as just a limiting case of maximal resemblance, rather than as the norm of communication. (The same is also true of imitation, another well-known but little understood means of cultural transmission, which I won't go into here.) A process of communication is basically one of transformation. The degree of transformation may vary between two extremes: duplication and total loss of information. Only those representations which are repeatedly communicated *and* minimally transformed in the process will end up belonging to the culture.

The objects of an epidemiology of representations are neither abstract representations nor individual concrete representations, but, we might say, strains, or families, of concrete representations related both by causal relationships and by similarity of content. Some of the questions we want to answer are: what causes such strains to appear, to expand, to split, to merge with one another, to change over time, to die? Just as standard epidemiology does not give a single general explanation for the distribution of all diseases, so there is no reason to expect that these questions will be answered in the same way for every kind of representation. The diffusion of a folk-tale and that of a military skill, for instance, involve different cognitive abilities, different motivations and different environmental factors. An epidemiological approach, therefore, should not hope for one grand unitary theory. It should, rather, try to provide interesting questions and useful conceptual tools, and to develop the different models needed to explain the existence and fate of the various families of cultural representations.

Though which factors will contribute to the explanation of a particular strain of representations cannot be decided in advance, in every case, some of the factors to be considered will be psychological, and some will be environmental or ecological (taking the environment to begin at the individual organism's nerve endings and to include, for each organism, all the organisms it interacts with). Potentially pertinent psychological factors include the ease with which a particular representation can be memorized, the existence of background knowledge in relationship to which the representation is relevant, and a motivation to communicate the content of the representation. Ecological factors include the recurrence of situations in which the representation gives rise to, or contributes to, appropriate action, the availability of external memory stores (writing in particular), and the existence of institutions engaged in the transmission of the representation.

Unsurprisingly, psychological and ecological factors are themselves affected by the distribution of representations. Previously internalized cultural representations are a key factor in one's susceptibility to new representations. The human environment is, for a great part, man-made, and made on the basis of cultural representations. As a result, feedback loops are to be expected both within models explaining particular families and between such models. The resulting complexity is of the ecological rather than of the organic kind. Though 'organicism' has disappeared from the anthropological scene, the organicist view of a culture as a well-integrated whole still lingers. The epidemiological approach departs from such cultural holism; it depicts individual cultures as wide open, rather than almost closed, systems and as approximating an ecological equilibrium among strains of representations, rather than as exhibiting an organic kind of integration. It is then of interest to find out which strains of representations benefit one another, and which, on the contrary, compete.

The identification of epidemiological phenomena in classical epidemiology often arises out of the study of individual pathology, but the converse is also true: the identification of particular diseases is often aided by epidemiological considerations. Similarly, when types of mental representations have been identified at a psychological level, the question of their epidemiology arises; and, conversely,

when particular strains of representations, or mutually supportive strains, have been epidemiologically identified, the question of their psychological character arises. More generally, as with the pathology and epidemiology of diseases, the psychology and epidemiology of representations should prove mutually relevant.

Psychological Speculations

Anthropologists and psychologists alike tend to assume that humans are rational — not perfectly rational, not rational all the time, but rational enough. What is meant by rationality may vary, or be left vague, but it always implies at least the following idea: humans beliefs are produced by cognitive processes which are on the whole epistemologically sound; that is, humans approximately perceive what there is for them to perceive and approximately infer what their perceptions warrant. Of course, there are perceptual illusions and inferential failures, and the resulting overall representation of the world is not totally consistent; but, as they are, the beliefs of humans allow them to form and pursue goals in a manner which often enough leads to the achievement of those goals.

Anthropologists and psychologists tend to assume that humans are rational, without explaining why. I assume some degree of rationality because it makes good biological sense. Why did vertebrates evolve so as to have more and more complex cognitive systems, culminating, it seems, in the human one, if not because this makes their interaction with the environment (e.g. feeding themselves, protecting themselves) more effective? Now, only an epistemologically sound cognitive system (i.e. one that delivers approximations of knowledge rather than pretty patterns or astounding enigmas) can serve that purpose, and, for that, it must be rational enough. This way of explaining why humans are rational implies that there is an objective reality, and that at least one function of human cognition is to represent in human brains aspects of that reality.

Fitting together reality and reason in this manner may seem commonsensical to psychologists, but many anthropologists — not so long ago, most of them — know better. People of different cultures have beliefs which are not only very different, but even mutually

incompatible. Their beliefs from our point of view, ours from theirs, seem irrational. If we want to maintain, nevertheless, that both they and we are rational, then an obvious way out is to deny that there is an objective reality to begin with. Reality on that view is a social construct, and there are at least as many 'realities' or 'worlds' as there are societies. Different beliefs are rational in different socially constructed worlds. I have argued at length against this view (see Sperber 1974, 1985b). Here I will merely state my bias: I find a plurality of worlds even less attractive than a plurality of substances; if there is a way, I would rather do without it.

There is a way, but first we must do a bit of conceptual house cleaning.²¹ What are we referring to when we talk of 'beliefs'? Take an example: we tend to assume that Peter believes that it will rain if he says so, or assents to somebody else saying so, or, in some cases, if he takes his umbrella on his way out. We do not mistake these behaviours for the belief itself; we take them, rather, as caused in part by Peter having the belief in question and, therefore, as evidence of the belief. We might be tempted to say, then (as many philosophers have – e.g. Ryle 1949), that a belief is a disposition to express, assent to, or otherwise act in accordance with, some proposition. As psychologists, however, we will want to go deeper and find out what kind of mental states might bring about such a disposition. An answer often heard nowadays is that humans have a kind of 'data base' or 'belief box' (Steven Schiffer's phrase) in which some conceptual representations are stored.²² All representations stored in that particular box are treated as descriptions of the actual world. When the occasion is right, this yields the usual behavioural evidence for belief: assertion and assent in particular.

The belief box story, however attractive, cannot be the *whole* story. Many of the propositions to which we are disposed to assent are not represented at all in our mind – a well-known point – and many of the propositions we are disposed not only to assent to but also to express and, in some cases, to act in accordance with are not, or not simply, stored in a data base or belief box – a more controversial point.

You have long believed that there are more pink flamingos on Earth than on the Moon, but no mental representation of yours had, until now, described that state of affairs. We may well have an infinity of such unrepresented beliefs, and a large proportion of these are

widely shared, though of course they have never been communicated. It is reasonable, however, to assume that what makes them unrepresented beliefs (more specifically, propositions to which we are disposed to assent) is that they are inferable from other beliefs which *are* mentally represented. What we need to add for this to the belief box is some inferential device allowing subjects to accept as theirs these unrepresented beliefs on the basis of the actually represented ones. The inferences in question are not made consciously, so the inferential device hooked up to the belief box must be distinct from, and need not resemble, human conscious reasoning abilities (see Sperber and Wilson 1986: ch. 2).

Besides accounting for unrepresented beliefs, hooking the belief box up to an inferential device introduces a factor of rationality in the construction of beliefs. Suppose that some of the representations in our belief box come from perception (broadly understood to include the 'perception' of one's own mental states), and that all other beliefs are directly or indirectly inferred from the perceptually based ones. This will already ensure areas of consistency among our beliefs. Suppose, furthermore, that the inferential device recognizes an inconsistency when it meets one, and corrects it. Then you get a tendency to enlarge areas of consistency (even though contradictory beliefs may still be held, provided they are never used as joint premises in an inference).

While perception plus unconscious inference might be the whole story for the beliefs of elephants, it could not be for the beliefs of humans. There are two interconnected reasons for this: first, many – possibly most – human beliefs are grounded not in the perception of the things the beliefs are about, but in communication about these things. Second, humans have a meta-representational, or *interpretive*, ability. That is, they can construct not only *descriptions* – that is, representations of states of affairs – but also *interpretations* – that is, representations of representations.²³ Now, humans use this interpretive ability to understand what is communicated to them and, more generally, to represent meanings, intentions, beliefs, opinions, theories and so on, whether or not they share them. In particular, they can represent a belief and take a favourable attitude to it, and therefore express it, assent to it, and generally show behaviours symptomatic of belief, on a basis quite different from belief box inclusion.

incompatible. Their beliefs from our point of view, ours from theirs, seem irrational. If we want to maintain, nevertheless, that both they and we are rational, then an obvious way out is to deny that there is an objective reality to begin with. Reality on that view is a social construct, and there are at least as many 'realities' or 'worlds' as there are societies. Different beliefs are rational in different socially constructed worlds. I have argued at length against this view (see Sperber 1974, 1985b). Here I will merely state my bias: I find a plurality of worlds even less attractive than a plurality of substances; if there is a way, I would rather do without it.

There is a way, but first we must do a bit of conceptual house cleaning.²¹ What are we referring to when we talk of 'beliefs'? Take an example: we tend to assume that Peter believes that it will rain if he says so, or assents to somebody else saying so, or, in some cases, if he takes his umbrella on his way out. We do not mistake these behaviours for the belief itself; we take them, rather, as caused in part by Peter having the belief in question and, therefore, as evidence of the belief. We might be tempted to say, then (as many philosophers have – e.g. Ryle 1949), that a belief is a disposition to express, assent to, or otherwise act in accordance with, some proposition. As psychologists, however, we will want to go deeper and find out what kind of mental states might bring about such a disposition. An answer often heard nowadays is that humans have a kind of 'data base' or 'belief box' (Steven Schiffer's phrase) in which some conceptual representations are stored.²² All representations stored in that particular box are treated as descriptions of the actual world. When the occasion is right, this yields the usual behavioural evidence for belief: assertion and assent in particular.

The belief box story, however attractive, cannot be the *whole* story. Many of the propositions to which we are disposed to assent are not represented at all in our mind – a well-known point – and many of the propositions we are disposed not only to assent to but also to express and, in some cases, to act in accordance with are not, or not simply, stored in a data base or belief box – a more controversial point.

You have long believed that there are more pink flamingos on Earth than on the Moon, but no mental representation of yours had, until now, described that state of affairs. We may well have an infinity of such unrepresented beliefs, and a large proportion of these are

widely shared, though of course they have never been communicated. It is reasonable, however, to assume that what makes them unrepresented beliefs (more specifically, propositions to which we are disposed to assent) is that they are inferable from other beliefs which *are* mentally represented. What we need to add for this to the belief box is some inferential device allowing subjects to accept as theirs these unrepresented beliefs on the basis of the actually represented ones. The inferences in question are not made consciously, so the inferential device hooked up to the belief box must be distinct from, and need not resemble, human conscious reasoning abilities (see Sperber and Wilson 1986: ch. 2).

Besides accounting for unrepresented beliefs, hooking the belief box up to an inferential device introduces a factor of rationality in the construction of beliefs. Suppose that some of the representations in our belief box come from perception (broadly understood to include the 'perception' of one's own mental states), and that all other beliefs are directly or indirectly inferred from the perceptually based ones. This will already ensure areas of consistency among our beliefs. Suppose, furthermore, that the inferential device recognizes an inconsistency when it meets one, and corrects it. Then you get a tendency to enlarge areas of consistency (even though contradictory beliefs may still be held, provided they are never used as joint premises in an inference).

While perception plus unconscious inference might be the whole story for the beliefs of elephants, it could not be for the beliefs of humans. There are two interconnected reasons for this: first, many – possibly most – human beliefs are grounded not in the perception of the things the beliefs are about, but in communication about these things. Second, humans have a meta-representational, or *interpretive*, ability. That is, they can construct not only *descriptions* – that is, representations of states of affairs – but also *interpretations* – that is, representations of representations.²³ Now, humans use this interpretive ability to understand what is communicated to them and, more generally, to represent meanings, intentions, beliefs, opinions, theories and so on, whether or not they share them. In particular, they can represent a belief and take a favourable attitude to it, and therefore express it, assent to it, and generally show behaviours symptomatic of belief, on a basis quite different from belief box inclusion.

Young Lisa is told by her teacher: 'There are male and female plants.' She understands 'male' and 'female' with respect to animals as more or less an extension of the distinction between men and women: females have children, males fight more easily, and so on. She does not see in plants anything resembling this distinction, and so she does not quite understand what her teacher is telling the class. On the other hand, she understands it in part; she understands that in some species there are two types of plants, and guesses that this difference has to do with reproduction, and so on. She trusts her teacher, and if he says that there are male and female plants, then she is willing to say so herself, to say that she believes it, and to exhibit various behaviours symptomatic of that belief.

Behind Lisa's belief behaviour, do we have a genuine belief? Not of the belief box kind, certainly, since such a half-understood idea (what I called a 'semi-propositional representation' in Sperber 1985b: ch. 2) could not have emerged from perception or from inferential communication. Remember, too, that the inferential device must be able to operate freely on beliefs in the belief box so as to yield more mutually consistent beliefs; but in that case half-understood ideas should not be allowed directly in the box, since their consistency with other representations and their implications are largely indeterminate. But how, then, might Lisa's half-understood idea of there being male and female plants be represented in her mind? Well, she might have in her belief box the following representations:

What the teacher says is true.

The teacher says that there are male and female plants.

Lisa's partial understanding of 'there are male and female plants' is now embedded in a belief box belief about what her teacher said. This belief, together with the other belief that 'what the teacher says is true', provides a validating context for the embedded representation of the teacher's words. This gives Lisa rational grounds for exhibiting many of the behaviours symptomatic of belief – but grounds quite different from plain belief box inclusion.

What this example suggests is that the beliefs we attribute to people on the evidence provided by their behaviour do not belong to a

single psychological kind; in other words, quite different types of mental states can bring about identical belief behaviour.

I maintain that there are two fundamental kinds of beliefs represented in the mind. There are descriptions of states of affairs directly stored in the belief box; let us call this first kind *intuitive beliefs*. Such beliefs are intuitive in the sense that they are typically the product of spontaneous and unconscious perceptual and inferential processes; in order to hold these intuitive beliefs, one need not be aware of the fact that one holds them, and even less of reasons for holding them. Then there are interpretations of representations embedded in the validating context of an intuitive belief, as in the above example; let us call this second kind *reflective beliefs*. These beliefs are reflective in the sense that they are believed in virtue of second-order beliefs about them.²⁴

Intuitive beliefs are derived, or derivable, from perception by means of the inferential device. The mental vocabulary of intuitive beliefs is probably limited to *basic concepts*: that is, concepts referring to perceptually identifiable phenomena and innately pre-formed, unanalysed abstract concepts (of, say, norm, cause, substance, species, function, number, or truth). Intuitive beliefs are on the whole concrete and reliable in ordinary circumstances. Together they paint a kind of common-sense picture of the world. Their limits are those of common sense: they are fairly superficial, more descriptive than explanatory, and rather rigidly held.

Unlike intuitive beliefs, reflective beliefs do not form a well-defined category. What they have in common is their mode of occurrence: they come embedded in intuitive beliefs (or, since there can be multiple embeddings, in other reflective beliefs). They cause belief behaviours because, one way or another, the belief in which they are embedded validates them. But they may differ in many ways: a reflective belief may be half-understood but fully understandable, as in the above example of the sex of plants; or, as I will shortly illustrate, it may remain half-understood for ever; or, on the contrary, it may be fully understood. The validating context may be an identification of the source of the reflective belief as a reliable authority (e.g. the teacher) or an explicit reasoning. Given the variety of possible contextual validations for reflective beliefs, commitment to these beliefs can widely vary, from loosely held opinions to

fundamental creeds, from mere hunches to carefully thought-out convictions. Reflective beliefs play different roles in human cognition, as I will very briefly illustrate.

For Lisa, forming and storing the half-understood reflective belief that there are male and female plants may be a step towards a more adequate understanding of the male–female distinction. It provides her with an incomplete piece of information which further encounters with relevant evidence may help complete. After she achieves an adequate understanding of the matter, her reflective belief that there are male and female plants may well be transferred to, or duplicated in, her belief box as an intuitive belief. So, one role of reflective belief is to serve as a ‘hold’ format for information that needs to be completed before it can constitute an intuitive belief.

Now, consider the following case. Young Bobby has in his belief box the two representations:

What Mom says is true.

Mom says that God is everywhere.

Bobby does not fully understand how somebody, be it God or anyone else, can be everywhere. However his mother saying so gives him sufficient ground to exhibit all the behaviours symptomatic of belief: he will readily state that God is everywhere, will assent when the same statement is made by others, and may even refrain from sinning in places where (apparently) nobody can see him. That God is everywhere is for Bobby a reflective belief. As he grows older, he may keep this belief and enrich it in many ways, but, if anything, its exact meaning will become even more mysterious than it was at first. Here is a belief which, like most religious beliefs, does not lend itself to a final, clear interpretation, and which therefore will never become an intuitive belief. Part of the interest of religious beliefs for those who hold them comes precisely from this element of mystery, from the fact that you are never through interpreting them. While the cognitive usefulness of religious and other mysterious beliefs may be limited (but see Sperber 1975b), it is not too difficult to see how their very mysteriousness makes them ‘addictive’.

In the two examples considered so far – Lisa and the sex of plants,

Bobby and divine omnipresence – what made the reflective representation a belief was the authority granted to the source of the representation: the teacher and the mother respectively. Laymen accept scientific beliefs on authority too. For instance, most of us believe that $e = mc^2$ with only a very limited understanding of what this formula means, and no understanding of the arguments that led to its adoption. Our belief, then, is a reflective belief of mysterious content, justified by our trust in the community of physicists. It is not very different, in this respect, from Bobby’s belief that God is everywhere.

There is a difference, though. Even for theologians, that God is everywhere is a mystery, and they too accept it on authority. For physicists, on the other hand, the theory of relativity is not a mystery, and they have reasons to accept it which have nothing to do with trust. Well-understood reflective beliefs, such as the scientific beliefs of scientists, include an explicit account of rational grounds to hold them. Their mutual consistency and their consistency with intuitive beliefs can be ascertained, and plays an important, though quite complex, role in their acceptance or rejection. Still, even for physicists, the theory of relativity is a reflective belief; it is a theory, a representation kept under scrutiny and open to revision and challenge, rather than a fact that could be perceived or unconsciously inferred from perception.

Half-understood or mysterious reflective beliefs are much more frequent and culturally important than scientific ones. Because they are only half-understood and therefore open to reinterpretation, their consistency or inconsistency with other beliefs, intuitive or reflective, is never self-evident, and does not provide a robust criterion for acceptance or rejection. Their content, because of its indeterminacy, cannot be sufficiently evidenced or argued for to warrant their rational acceptance. But that does not make these beliefs irrational: they are rationally held if there are rational grounds to trust the source of the belief (e.g. the parent, the teacher, or the scientist).

This, then, is my answer to those who see in the great diversity and frequent apparent inconsistency of human beliefs, an argument in favour of cultural relativism: there are two classes of beliefs and they achieve rationality in different ways. Intuitive beliefs owe their rationality to essentially innate, hence universal, perceptual and

inferential mechanisms; as a result, they do not vary dramatically, and are essentially mutually consistent or reconcilable across cultures. Those beliefs which vary across cultures to the extent of seeming irrational from another culture's point of view are typically reflective beliefs with a content that is partly mysterious to the believers themselves. Such beliefs are rationally held, not in virtue of their content, but in virtue of their source. That different people should trust different sources of beliefs – I, my educators, you, yours – is exactly what you would expect if they are all rational in the same way and in the same world, and merely located in different parts in this world.

Different Types of Beliefs, Different Mechanisms of Distributions

Let us now bring together the anthropological and psychological speculations developed so far: If there are different kinds of beliefs, then we might expect them to be distributed by different mechanisms. More precisely, we might expect the distribution of intuitive beliefs, which are a relatively homogeneous kind, to proceed along roughly common lines,²⁵ and the distribution of reflective beliefs, which are much more diverse, to take place in many different ways. In this concluding section, I would like to suggest that such is indeed the case.

In all human societies, traditional or modern, with or without writing, with or without pedagogic institutions, all normal individuals acquire a rich body of intuitive beliefs about themselves and their natural and social environment. These include beliefs about the movement of physical bodies, the behaviour of one's own body, the effects of various body-environment interactions, the behaviour of many living kinds, the behaviour of fellow humans. These beliefs are acquired in the course of ordinary interaction with the environment and with others. They need no conscious learning effort on the part of the learner and no conscious teaching effort on the part of others (see Atran and Sperber 1991). Even without teaching, these beliefs are easily acquired by everybody. The more fundamental ones are acquired quite early, suggesting a very strong innate predisposition (see Keil 1979; Carey 1982, 1985; Gelman and Spelke 1981; Hirschfeld 1984, 1994).

Some intuitive beliefs are about particulars (particular locations, personal events, individual animals or people), and are idiosyncratic or are only shared very locally; others are general (or about widely known particulars such as historical events and characters), and are widespread throughout a society. General intuitive beliefs vary across cultures, but they do not seem to vary greatly. To mention just one piece of anecdotal evidence, one has yet to find a culture in which where intuitive beliefs about space and movement are so different from modern Western ones that the natives have inordinate problems in learning to drive a car. Much recent work in ethno-science shows, too, that cross-cultural differences in zoological, botanical, or colour classification are rather superficial, and that for each of these domains (and presumably for other domains, e.g. artefacts or mental states), there are underlying universal structures (see Berlin and Kay 1969; Berlin et al. 1973; Berlin 1978; Atran 1985, 1986, 1987).

What role does communication play in the construction of intuitive beliefs? The answer is not simple. Intuitive beliefs are (or are treated as) the output of perception and unconscious inference, either the subject's own perceptions and inferences or those of others in the case of intuitive beliefs acquired through communication. Even when an intuitive belief is derived from the subject's own perceptions, the conceptual resources and the background assumptions which combine with the sensory input to yield the actual belief have, in part, been acquired through communication. So, it seems, both perception and communication are always involved in the construction of intuitive beliefs. Perception is involved either as the direct source of the belief or as its assumed indirect source (which puts a strong constraint on the possible contents of intuitive beliefs). Communication is involved either as a direct source or, at the very least, a source of concepts and background.²⁶

What, now, is the relationship between the relative shares of perception and communication in the construction of an intuitive belief, on the one hand, and its social distribution, on the other? Is it the case that the greater the share of communication, the wider the distribution? Again, the answer is not that simple. A great number of very widespread beliefs owe their distribution to the fact that all

members of a society, or in some cases all humans, have similar perceptual experiences. However, as already suggested, the resources for perception are themselves partly derived from communication.

Take the widespread intuitive belief that coal is black: were you told it, or did you infer it from your own perception? Hard to know. But even if you inferred it from perception, in doing so, you used the concepts of black and of coal, and how did you acquire those? Regarding 'black', it seems that the category is innately pre-wired, so that, when you learned the word 'black', you merely acquired a way to express verbally a concept you already possessed (see Berlin and Kay 1969; Carey 1982). Regarding 'coal', no one would claim that the concept is innate; but what might well be innate is the structure of substance-concepts with the expectation of regular phenomenal features – in particular, colour. So, while you probably acquired the *concept* of coal in the process of learning the word 'coal', acquiring the concept meant no more than picking the right innate conceptual schema and fleshing it out. In the process of fleshing it out, either you were told, or you inferred from what you saw, that coal is black.

It does not make much difference, then, whether an individual's belief that coal is black is derived from perception or from communication: once the concept of coal is communicated, the belief that coal is black will follow one way or the other. This is generally true of widespread intuitive beliefs. These beliefs conform to cognitive expectations based on culturally enriched innate dispositions, and are richly evidenced by the environment. As a result, different direct perceptual experiences and different vicarious experiences acquired through communication converge on the same general intuitive beliefs.

Widespread intuitive beliefs, even exotic ones, are rarely surprising. They are not the kind of beliefs that generally excite the curiosity of social scientists, with the exception of cognitive anthropologists. Among psychologists, only developmental psychologists have started studying them in some detail. Yet intuitive beliefs not only determine much of human behaviour; they also provide a common background for communication and for the development of reflective beliefs.

Whereas widespread intuitive beliefs owe their distribution both

to common perceptual experiences and to communication, widespread reflective beliefs owe theirs almost exclusively to communication. The distribution of reflective beliefs takes place, so to speak, in the open: reflective beliefs are not only consciously held; they are also often deliberately spread. For instance, religious believers, political ideologists, and scientists, however they may differ otherwise, see it as incumbent upon them to cause others to share their beliefs. Precisely because the distribution of reflective beliefs is a highly visible social process, it should be obvious that different types of reflective beliefs reach a cultural level of distribution in very different ways. To illustrate this, let us consider very briefly three examples: a myth in a non-literate society, the belief that all men are born equal, and Gödel's proof.

A myth is an orally transmitted story which is taken to represent actual events, including 'supernatural' events incompatible with intuitive beliefs. Therefore, for a myth to be accepted without inconsistency, it has to be insulated from intuitive beliefs: that is, held as a reflective belief. A myth is a cultural representation; this means that the story is told (given public versions) often enough to cause a large enough proportion of a human group to know it (have mental versions of it). For this, two conditions must be met. First the story must be easily enough and accurately enough remembered on the basis of oral inputs alone. Some themes and some narrative structures seem in this respect to do much better cross-culturally than others. The changing cultural background affects memorability, too, so that the content of a myth tends to drift over time so as to maintain maximal memorability.

Second, there must be enough incentives to actually recall and tell the story on enough occasions to cause it to be transmitted. These incentives may be institutional (e.g. ritual occasions where telling the story is mandatory); but the surest incentive comes from the attractiveness of the story for the audience and the success the story-teller can therefore expect. Interestingly, though not too surprisingly, the very same themes and structures which help one remember a story seem to make it particularly attractive.

If the psychological conditions of memorability and attractiveness are met, the story is likely to be well distributed; but in order for it to be a myth, rather than, say, a mere tale recognized and enjoyed as

such, it must be given credence. What rational grounds do people have to accept such a story as true? Their confidence in those who tell it to them: typically, their confidence in elders whom they have many good reasons to trust and who themselves claim no other authority than that derived from *their* elders. The originator of the chain might be a religious innovator who claimed divine authority for a distinctly different version of older myths. Reference to elders provides a self-perpetuating authority structure for a story which already has a self-perpetuating transmission structure. Still, the authority structure is more fragile than the transmission structure, and many myths lose their credibility, though neither their memorability nor their attractiveness, and end up as tales.

The belief that all men are born equal is a typically reflective belief: it is not produced by perception or by unconscious inference from perception. Rather, except for a few philosophers who originated the belief, all those who have held it came to it through communication. Such a belief does not put any significant weight on memory, but it does present a challenge for understanding, and indeed it is understood differently by different people. As already suggested, the fact that it lends itself to several interpretations probably contributed to its cultural success.

Still, the most important factor in the success of the belief that all men are born equal is its extreme relevance — that is, the wealth of its contextual implications (see Sperber and Wilson 1986) — in a society organized around differences in birthrights. People who accepted, and indeed desired, the implications of this belief found there grounds to accept the belief itself and to try to spread it. However, there was a risk, not to holding the belief, but to spreading it, and so the belief spread only where and when there were enough people willing to take this risk. In other words, unlike a myth, which seems to have a life of its own and to survive and spread, as myth or as tale, in a great variety of historical and cultural conditions, the cultural destiny of a political belief is tied to that of institutions. Ecological factors (more particularly, the institutional environment) play a more important role in explaining the distribution of a political belief than cognitive factors.

Consider now a mathematical belief, such as Gödel's proof. Again, all those who hold it, except Gödel himself, arrived at it

through communication. However, the communication, and hence the diffusion, of such a belief meets extraordinary cognitive difficulties. Only people with a high enough level of education in mathematical logic can begin to work at understanding it. Outside scholarly institutions, both the means and the motivation to do that work are generally lacking. On the other hand, once the difficulties of communication are overcome, acceptance is no problem at all: to understand Gödel's proof is to believe it.

The human cognitive organization is such that we cannot understand such a belief and not hold it. To some significant extent, and with obvious qualifications, this is the case with all successful theories in the modern natural sciences. Their cognitive robustness compensates, so to speak, for their abstruseness in explaining their cultural success. The fact that successful scientific theories impose themselves on most of those who understand them is manifest to people who don't understand them. This leads, quite rationally, to lay persons believing that these theories are true and expressing as beliefs whatever they can quote or paraphrase from them. Thus Gödel's proof, and scientific theories generally, become cultural beliefs of a different tenor, accepted on different grounds by the scientists themselves and by the community at large.

We might contrast our three examples in the following way. The distribution of a myth is determined strongly by cognitive factors, and weakly by ecological factors; the distribution of political beliefs is determined weakly by cognitive factors, and strongly by ecological factors; and the distribution of scientific beliefs is determined strongly by both cognitive and ecological factors. However, even this exaggerates the similarities between the three cases: the cognitive factors involved in myth and in science and the ecological factors involved in politics and in science, are very different. The very structure of reflective beliefs, the fact that they are attitudes to a representation, rather than directly to a real or assumed state of affairs, allows endless diversity.

Notwithstanding their diversity, explaining cultural beliefs, whether intuitive or reflective, and if reflective, whether half-understood or fully understood, involves looking at two things: how they are cognized by individuals and how they are communicated within a group; or to put it in the form of a slogan: *Culture is the precipitate of cognition and communication in a human population.*