

(1925, 34, 118–22, 1965, 39–40; cf. Arberman 1939, 27). I have elsewhere tried to show in detail that here Durkheim errs: all people all over the world have intuitive expectations about the ordinary course of events, together with ideas that violate these intuitive expectations (Geary 2005; Pyysiäinen 2001b, 55–74, 2004d, 39–52, 81–89, 2005d). Moreover, the “regeneration” of the group is a cognitive act, consisting of persons’ mental representations of other persons’ mental representations; what he called “collective consciousness” thus is shared knowledge and can be studied from the cognitive psychological perspective (see Pyysiäinen 2005f).

Belief in supernatural agents forms an obvious recurrent pattern in various cultural traditions. As some of these beliefs help people express, process, and justify their central values and norms, we can try to combine the Tylo-rian and Durkheimian legacies in explaining how beliefs about supernatural agents are used in organizing a society (see Boyer 2000b, 2002; Pyysiäinen 2005f). Typical examples of supernatural agents are gods, spirits, ghosts, angels, demons, and so forth. The category of supernatural agents may seem so heterogeneous that one may well ask whether it is a genuine category at all. Some might want to claim that it is a pseudocategory just like “mysticism,” “ritual,” or even “religion” itself (see, e.g., Fitzgerald 1996, 1997; Penner 1983). I think this view is overly pessimistic. If we can find some theoretical depth to the ideas of agency and counterintuitiveness, the idea of supernatural agents might be operationalized for research. We need a theoretical frame of reference in which the various *emic* distinctions between different kinds of supernatural agents are seen as subdivisions within the general, *etic* category of supernatural agency. Using concepts at a mediating level between cultural particulars and the general category of supernatural agents, we may then conceptualize the various recurrent patterns within that category.

1.3 There Must Be *Somebody* Out There

I use quite consciously the expression “supernatural *agents*” instead of “supernatural *beings*.” By agents, psychologists and philosophers mean organisms whose behavior can be successfully predicted by postulating conscious beliefs and desires (Dennett 1993, 15–17) or entities whose behavior is caused by their mental states (Bechtel 2008, xii). I use the concept in the sense of an organism to which animacy (liveliness, self-propelledness) and mentality (beliefs and desires) are (correctly or incorrectly) attributed.

I follow Barrett (2008) in distinguishing between two components of agency: animacy and mentality. The domain of animacy is characterized by such things as self-propelledness and goal-orientation: organisms tacitly postulate animacy for other organisms when they feel that those organisms move by themselves and seem to be moving toward some goal (see Boyer and

Barrett 2005). When one also attributes to the organism a conscious intention and begins to simulate its mental states or to try to theorize about its beliefs and desires, one moves on to “mentalizing” or “mind reading” (Nichols and Stich 2003). This ability is often called a “theory of mind” (ToM), in the sense of folk-psychological theories about other minds (Carruthers and Smith 1996).⁷

I distinguish three overlapping cognitive mechanisms that contribute to agentic reasoning. The first is hyperactive agent detection (HAD, Barrett 2000): the tendency to postulate animacy—this mechanism is triggered by cues that are so minimal that it often produces false positives, for example, we see faces in the clouds, mistake shadows for persons, and so forth. Second is hyperactive understanding of intentionality (HUI): the tendency to postulate mentality and to see events as intentionally caused even in the absence of a visible agent. Third is hyperactive telefunctional reasoning (HTR): the tendency to see objects as existing for a purpose.

1.3.1 *Agency*

Barrett (2000) coined the acronym HADD (hyperactive agent detection device) to refer to the cognitive processes that help us recognize agents (actually, animacy) and distinguish them from nonagents. This mental “device” is hyperactive or hypersensitive, in that under certain conditions, it is triggered by very minimal cues.⁸ We see faces in the clouds and detect predators in rustling bushes because such ambiguous perceptions easily trigger the postulation of agency. According to Barrett, a normally functioning HADD is hyperactive by its very nature—hyperactivity is not something exceptional. From an evolutionary point of view, this is plausible, insofar as the costs of false positives that an overreacting detector produces are lower than the benefits it brings (Atran 2006).

The following most common direct cues for agency have been suggested (Blakemoore et al. 2003; Boyer and Barrett 2005; see Heider and Simmel 1944):

1. Animate motion that has as its input such things as nonlinear changes in direction, sudden acceleration without collision, and change of physical shape that accompanies motion (e.g., caterpillar-like crawling)
2. An object reacting at a distance
3. Trajectories that only make sense if the moving entity is trying to reach or avoid something, which leads to goal-ascription
4. An entity appearing to be moving by conscious intention to an apparent end result (intention-ascription)
5. The experience of joint attention (for which we develop a capacity between nine and twelve months of age)

These cues trigger the feeling or intuition of agency spontaneously and automatically, in the sense that this intuition can neither be rationally controlled nor initiated or terminated at will (see Bargh 1994; Pyysiäinen 2004c; see

the appendix). However, I suggest that the first three cues only trigger animacy assumptions, while the latter two trigger HUI. It may also be misleading to call HADD a (single) “device,” as several systems contribute to the perception that one is facing an agent (see Boyer and Barrett 2005). Once people start reasoning about the beliefs and desires of a postulated agent, they use mind reading: the capacity to make inferences about the beliefs and desires of others and to explain their behavior on that basis (see Carruthers and Smith 1996; Frith and Frith 2005; Nichols and Stich 2003; Premack and Woodruff 1978; Saxe and Baron-Cohen 2006; Tomasello et al. 2005; Tremlin 2006, 75). Mind reading is often understood to be innate—not a learned ability but rather something triggered in the course of normal development (Leslie 1996; Wellman and Miller 2006, 28). It has been seen as an innate modular algorithm, (e.g. Leslie), as an innate body of knowledge (e.g. Pinker and Sperber), and as modular in all three senses of the concept of a “module” (Baron-Cohen) discussed in the appendix (see Gerrans 2002, 308). It can also be understood to be based on a nonmodular conceptual competence (Wellman et al. 2001; see Yazdi et al. 2006).

The standard psychological test for ToM is the so-called false belief task (Wimmer and Perner 1983). In this test, children are shown a sketch in which a boy called Maxi first puts a chocolate bar in container A. Then, without Maxi’s knowing, his mother moves it to container B. The subjects are then asked where Maxi would look for the chocolate bar when he returns. Children around three years of age tend to say (incorrectly) that Maxi would look in container B (where the bar was hidden). Only in the fourth year does the tendency appear to say that Maxi would look in container A. It thus seems that younger children cannot understand that Maxi remains ignorant of his mother hiding the bar in container B. These results have been obtained in many replications of the experiment and across many variations of the original design. Scholarly opinions differ on what these results really tell about the ways children think (see Baron-Cohen 2000; Baron-Cohen et al. 1985; Fodor 1992; Perner 1995; Wimmer and Perner 1983).⁹

Simon Baron-Cohen, Alan Leslie, and Uta Frith (1985) used the same task to test autistic children’s ability to impute mental states to others. They tested twenty autistic children, fourteen children with Down syndrome, and twenty-seven clinically normal preschool children. Two dolls were used, Sally and Anne. First, Sally placed a marble in her basket, and then left the scene. Anne then transferred the marble into her own box. When Sally returned, the experimenter asked: “Where will Sally look for her marble?” Children with Down syndrome as well as normal preschool children answered correctly by pointing to the basket where Sally had put the marble. The autistic group consistently answered by pointing to the box where the marble really was; the autistic children did not seem to get the difference between their own and the doll’s knowledge. In the authors’ words, they “failed to employ a ToM”—a failure

consisting of an “inability to represent mental states” (Baron-Cohen et al. 1985, 43). Autism thus might seem to involve a lack of a ToM. The inability to make inferences about what other people believe to be the case in a given situation prevents one from predicting what they will do (Baron-Cohen 2000; Baron-Cohen et al. 1985, 39; see Carruthers and Smith 1996, 223–73; Frith 2001).

However, failure in the test does not necessarily indicate a lack of a ToM, because having a ToM does not necessarily require the ability to reason about false beliefs (Bloom and German 2000; see Wellman et al. 2001). Reasoning about false beliefs is just one way of using ToM (see Stone 2005). Philip Gerrans and Valerie Stone also argue that there is no hardwired ToM as a dedicated module at all and that belief attribution is not supported by a domain-specific, modular mechanism (see appendix). Instead there is interaction between such low-level, domain-specific mechanisms as tracking gaze and bodily movement, joint attention, and so forth and higher-level domain-general mechanisms for metarepresentation, recursion, and executive function. The output of the low-level mechanisms serve as input for the higher-level, domain-general processes (Gerrans 2002, 306, 311; Stone and Gerrans 2006a,b).

Perner et al. (2006), as well as Saxe et al. (2006) present neuroscientific evidence for the view that in the right and left temporo-parietal junction of the brain there is a specialized, domain-specific neural mechanism for reasoning about beliefs. Inhibitory control of mental contents, together with response selection, is mediated by domain-general mechanisms, while the domain-specific mechanism of the temporo-parietal junction is only recruited for processing *beliefs*. And reasoning about beliefs is faster than following domain-general rules in reasoning about other things than beliefs (Saxe et al. 2006; see Saxe and Baron-Cohen 2006).

Leslie et al. (2004, 2005) argue that ToM is a partly modular learning “mechanism” that only “kick-starts” the attribution of beliefs and desires. Just as color vision provides us with color concepts, ToM introduces belief and desire concepts. Reasoning about the contents of beliefs takes place through a selection process (SP) with inhibition. Developing slowly from the preschool period onward, SP acts by selecting a content for an agent’s belief and an action for the agent’s desire. In everyday thinking, taking a belief to be true is the default, so success in the false belief task requires an ability to inhibit the default attribution and to select the erroneous belief for the character looking for the hidden object.

The mentality of organisms that ToM processes is very early on understood as separate from the physical body. Kuhlmeier et al. (2004), for example, show that five-month-old infants apply the constraint of continuous motion to inanimate blocks but not to persons. They thus do not seem to view human agents as material objects. In the experiment, twenty infants watched a film in which a woman was standing on a stage that contained two large red screens separated by 1.21 meters. The woman then went behind the first screen; her

identical twin sister in identical clothing had been hiding behind the second screen and now emerged from behind it. What the infants saw was a woman moving behind a screen and then suddenly emerging from behind the other screen without being visible in the space between the two screens. The infants revealed that they were not surprised by this apparently miraculous event, although they were surprised when they watched mere physical objects (as distinct from humans) behaving in the same way. (They exhibited surprise by looking longer at the action they were witnessing.) This suggests that they do not consider agency to be constrained by the physical body.

Even as adults, people do not feel themselves to *be* bodies but instead feel that they *occupy* bodies (Bloom 2005, 191, 2007; Merleau-Ponty 1992; see Bronkhorst 2001, 402). Routinely, people spontaneously attribute their own agentive properties to the homunculus called the “self,” which is not merely the sum of a set of lower-level, “dumb[er]” homunculi (Nichols and Stich 2006, 9–10). To the extent that the agentive properties are detached from the physical body, it is perfectly natural for people to have intuitive beliefs about disembodied agents such as spirits. However, disembodiment does not mean a complete lack of bodily form; mentality may instead be attributed to various forms of “subtle” or otherwise nonstandard bodily forms.

1.3.2 *Intentionality*

There are also indirect cues that lead people to interpret an event as intentionally caused even when they do not perceive an agent. Hyperactive understanding of intentionality produces an automatic conclusion that an event or structure must have been caused or designed by an intelligent agent, even when no trace of an agent is evident (Barrett 2004b, 34). For example, normally, when one sees an artifact, one knows that it must have been designed and made by an agent, even though this agent is not present and one does not know her or his identity (see Czachesz 2007, 86).

This kind of reasoning contributes to beliefs about supernatural agency when it is extended to such cases as, for instance, seeing the image of Jesus falling off the wall right at the moment when someone says something blasphemous. The coincidence of the words uttered and the picture falling creates the feeling that the falling was intentionally caused, despite the fact that no one is present and one has no idea of the mechanism by which an agent could have caused the picture to fall (see Atran 2002a, 59–63; Bering 2006; Bering and Parker 2006). Bering (2003) suggests that there may even be a specific cognitive mechanism devoted to processing apparently intentional events (“existential meaning”) in the absence of a physical agent.

Such inferences are fast and automatic intuitions (see the appendix). The low-level intuitions about agency are dedicated to combining perceived movement with the perceiver’s reflective ideas of agency and free will. When

action consistently follows prior thought and when apparent causes of action other than somebody's thought are excluded, we have the experience of volitional action. Volition, in turn, presupposes an agent. Torsten Nielsen (1963) showed in the 1960s that when subjects were asked to put one hand in a box and draw a straight line with it while looking into the box through a tube, they could be fooled into believing that the experimenter's hand seen in the box through a mirror was their own (which they did not see). When both hands drew a straight line, the subjects perceived the alien hand as their own. When the experimenter's hand drew a curve to the right, the subjects still experienced the hand as their own, and now perceived it as making involuntary movements.

Similarly, Ramachandran and colleagues arranged an experiment in which a subject with an amputated arm was asked to put his real arm and what he experienced as his phantom limb in a box with a vertical mirror in the middle; the reflection of the real arm then appeared in the mirror where the phantom would have been if it were real. When the subject moved his real arm, he felt he was moving his phantom limb. Then the experimenter put his arm inside the box so that it appeared in the place of the phantom. Observing his real arm and the experimenter's arm in the box made the subject feel as though his phantom arm was moving (Ramachandran and Blakeslee 1999, 46–48; see Wegner 2002, 40–44).

Wegner (2002, 44) thus argues that one can think movement is intentional by watching *any* body move where one thinks one's own body is. But there are also cases in which one perceives a movement as guided by an intentional will not one's own. When the picture of Jesus Christ suddenly falls off the wall, the default explanation is not necessarily that one did it oneself (although this is possible). One can come to think that the movement is intentional by combining one's perception of it with a representation of an invisible agent.

Wegner (2002, 44) suggests in passing that the experience of another's movement as willful is mediated by the so-called mirror neuron system in the brain. This system mediates such intuitive reasoning as, for example, feeling disgust when witnessing someone drink a glass of milk with a face contracting in an expression of disgust. It cannot mediate such reflective ToM reasoning as in, for instance, trying to figure out what gift would please a foreign colleague (Bargh 1994, 3; Keysers and Gazzola 2007). Mirror neurons were first found in monkey brains in premotor area F5 and parietal area 7b. They are called mirror neurons because they are activated during both the execution of purposeful, goal-related hand actions (grasping etc.) and the observation of similar actions performed by conspecifics (see Gallese et al. 1996; Rizzolatti et al. 1988; Rizzolatti et al. 2000). A mirror neuron system has also been found in the human brain (although no special type of neuron seems to be involved). Implicit, automatic, and unconscious simulation processes establish a link between an observed agent and the observing agent, thus making imitation possible (Bremmer et al. 2001; Gallese and Metzinger 2003; Hari et al. 1998).

It has been suggested that social cognition is based on the mirror neuron system, which mediates observations about other agents' apparently intentional behavior (see Hurley 2008a). When two agents interact socially, the mirror neuron system is activated and creates a shared neural representation (Becchio et al. 2006, 66–67). It thus helps people understand the intentions of others and to relate them to their own (see Gallese and Goldman 1998; Gallese and Metzinger 2003; Gallese et al. 2004). In being restricted to motor processes and intuitions, though, the mirror neuron theory alone may not be sufficient to account for social cognition (see Bargh 1994; Hurley 2008b; Keysers and Gazzola 2007). Social cognition also involves complicated reflective processes related to understanding of intentionality.

Recognizing consciously goal-directed action involves recognizing intentionality in three senses: intention recognition, attribution of intention to its author (HUI), and understanding motivation behind the intention (HUI, HTR) (see Becchio et al. 2006). Intentionality as the mark of the mental is usually understood as “directedness” or “aboutness” (Brentano 1924; Dennett 1993, 67, 1997, 66; Husserl 1950a,b). Mental states are beliefs about the world, desires for things, and so forth (Wellman and Miller 2006, 34–35). Following Jaakko Hintikka (1975), I argue that it might be better understood as intensionality, however (see the appendix, section 2.1).

Intensionality relates to such things as meanings, properties, property-based relations, and propositions. When we describe a set intensionally, we list the properties the members of the set have. We may, for instance, try to list the various gods of the religions of world using the following formula: { x : x is a god}. An extensional listing, in contrast, is made by just listing the members: {Allah, Quetzalcoatl . . . Zeus}. These two lists are extensionally equivalent but intensionally distinct. This distinction is based on the fact that terms and expressions have an intension or meaning (*Sinn*) and an extension or reference (*Bedeutung*) (Frege 1966). For example, someone's beliefs about the Evening Star are not necessarily beliefs about the Morning Star, because the person in question may not know that these two are the same planet.¹⁰ The terms have different intentions but the same reference. Mere extensional concepts are not enough for conceptualizing agency because agency is defined by the mental act of directedness (instead of mere reactivity, Leslie 1994).

The idea of intentionality as intensionality brings to the fore what seems to be crucial in intentionality in a psychological sense. According to Hintikka, “a concept is intentional if and only if it involves the [logically] simultaneous consideration of several possible states of affairs or courses of events.” Intentionality as intensionality thus means that for an act to be intentional, it must involve a conscious choice made between several alternatives (1975, 195; see also 212–13, 1982). Intentionality as intensionality means that intentional systems involve the consideration of possible worlds as logically alternative scenarios (2006a, 21–24, 2006b, 557–58; see Dennett 1993, 122, 174–75). Thus, ToM depends

on the ability for counterfactual reasoning about contrastive states of affairs or courses of events (see also Buller 2005a, 194). Directedness is involved, but in an intensional sense: a mental state is directed toward something, but with the awareness that it might just as well be directed toward something else. In intentionality, as distinct from mere reactions, a course of action is selected from among multiple mentally represented alternatives.

If mind reading or ToM operates through a selection process with inhibition, understanding intentionality as intensionality fits it much better than the idea of intentionality as simply directedness. For example, understanding another's beliefs and desires in the false belief task involves precisely consideration of possible worlds as logically alternative scenarios. This entails understanding that the mind is a representational device and that there is no reason that all represented propositions should be true (Perner 1993). One has to be able to metarepresent another's beliefs, that is, to embed them in one's own beliefs, as in "I believe that 'she thinks that he is wrong.'" Valerie Stone (2005) argues that it is precisely the ability of metarepresentation that enables children to succeed at explicit false belief tasks: the child must be able to understand that various agents represent the location of the chocolate bar in various ways. According to Bering, it is around six or seven years of age that children become capable of understanding such third-order intentionality in which "he thinks that she believes that he wants," for example. At that point in their development, they are also able to regard random events as symbolic and declarative of a supernatural agent's mental states (Bering and Johnson 2005, 134–36).¹¹

Degrees or orders of intentionality is an idea of Dennett (1993, 243–46). It has been argued that normal adult humans are capable of fourth- or fifth-order intentionality (when no external memory stores are used) (Dunbar 2003, 170, 2006, 171–72). "I know¹² that John *wants* Mary to *understand* that Bill *believes* that Linda *loves* him" is an example of fourth-order intentionality. Such metarepresentation of intentional attitudes requires a capacity to understand recursive¹³ structures (Stone 2005; see Hintikka 1975). To the extent that we are interested in how ToM works in social cognition, only embedded *mental states* count, not mere embedded observed facts. For example, the sentence "John *says* Mary to *claim* that Bill *wrote* her that Linda *sent* him a love letter" is not a case of degrees of intentionality. Even an autistic person incapable of mind reading might be able to understand this sentence as reporting four interrelated facts.

The recursive embedding of metarepresentations seems to be an exclusively human capacity; apes, for example, do not understand recursive structures, though the evidence here is somewhat ambiguous (Dennett 2006, 111; Dunbar 2003, 170). Premack and Premack (2003, 149–53), however, claim that chimpanzees are only capable of first-order intention ascription and that even this requires much conscious effort from the chimp, whereas a ten- to twelve-month-old human infant does this spontaneously. Human social cognition starts where the chimpanzee's cognitive ability ends. As Tomasello and

Carpenter (2007, 122) observe, “from a very early age human infants are motivated to simply share interest and attention with others in a way that our nearest primate relatives are not.” Joint attention, which is one of the low-level inputs of ToM, starts to develop in human infants as early as nine months, when the learning child begins to look at an object and at the parent, trying to draw the adult’s attention to a shared object (see Moore and D’Entremot 2001; Tomasello and Carpenter 2007; Tomasello et al. 2005).

Sperber (1997) argues that recursively metarepresenting a belief within another belief can provide a validating context for the embedded belief. It is then accepted as true because of certain second-order beliefs about it (Boyer 1994b, 120; Sperber 1996, 69–70, 89–97). To the extent that these second-order beliefs seem compelling, the first-order belief embedded in them is also plausible.¹⁴ In, for example, “Jesus is our redeemer’ is true because it says so in the Bible,” “Jesus is our redeemer” derives its plausibility from the validating embedding “It says so in the Bible” (see Pyysiäinen 2003c).

Metarepresentation thus makes it possible to make inferences about propositions with an undecided truth-value—as in, for instance, “‘If it was John who stole my wallet,’ then he should be arrested.” Here the judgment of John being guilty must be temporally suspended in order to make a provisional inference (see Cosmides and Tooby 2000, 59–60; Pyysiäinen 2003c; Sperber 1996, 71–73). It is important to distinguish between what Sperber calls “half-understood” propositions and propositions with an undecided truth-value. It is possible to generate conditional inferences from premises with an undecided truth-value, but it is not possible to do so from incomprehensible premises. These can only be used as quotations that receive a meaning from the validating context in which they are metarepresented.

Failure to understand the nature of metarepresentation leads to accepting all beliefs as one’s own (see Sperber 1994, 2000a; Pyysiäinen 2003c). A psychiatrist may, for example, entertain the metarepresentation “Joe believes ‘I am Jesus.’” If she is not able to use the metarepresentation “Joe believes,” then “I am Jesus” becomes her own belief (see Corcoran et al. 1995; Frith 1995). Metarepresented beliefs are decoupled from reality, and their semantic relations are suspended, in the sense that one cannot directly reason from “Ann believes ‘John is a spy’” that John really is a spy (Cosmides and Tooby 2000). This would only follow in the case that the metarepresentational context is automatically validating. For a Christian believer, the metarepresentation “It says so in the Bible” often is such a validating context, for example (Pyysiäinen 2003c). It is beliefs about beliefs that validate beliefs.

1.3.3 *Teleofunctional Reasoning*

Humans are prone to see things as existing for a reason or purpose. This tendency to view entities as existing for a purpose may derive from children’s early

emerging ability to think that there are hidden intentions behind everything and to interpret agents' behavior as goal-directed (Johnson 2000, 188, 208; Kelemen and DiYanni 2005, 6). Children's understanding of how agents use objects as means to achieve goals may provoke a rudimentary teleofunctional view of entities: agents' intentions are understood as being intrinsic properties of the objects themselves (Kelemen 1999a,b; Kelemen and DiYanni 2005, 6). There is experimental evidence for a "promiscuous teleology": British and American elementary schoolchildren are prone to generating teleofunctional explanations of the origins and nature of living and nonliving natural entities and endorsing intelligent design as the source of both animals and artifacts (Evans 2000, 2001; Kelemen 2004; Kelemen and DiYanni 2005).

Kelemen and DiYanni (2005, 7) point out that this view about purpose does not necessarily derive from agentive intuitions. Teleo-functional reasoning may also be an independent characteristic of causal reasoning. However, Asher and Kemler Nelson (2008) provide evidence for the interpretation that three- and four-year-old children can adopt the intentional stance (Dennett 1993) and do understand the true functions of artifacts to be the *designed* functions. It thus is also possible that intuitions about purpose and design and intuitions about agency derive from a common source.

Lewis Wolpert (2007, 27–33, 67–82) argues that all causal reasoning in humans derives from the manufacturing and using of tools that typify the species, not from social interaction (cf. Dunbar 1993, 2002, 2003, 2006). The making of first tools created a new kind of selection pressure, to which causal thinking is an adaptation: the making of tools required an ability to understand the ideas of means and ends and of causality and intentionality. Although non-human primates can distinguish the animate from the inanimate, they do not view the world in terms of intermediate and often hidden underlying causes, reasons, intentions, and explanations. (They understand animacy but not mentality.)

Wolpert neither develops this idea carefully nor contrasts it in detail with competing accounts. It does not have to be incompatible with, for example, Dunbar's view of human sociality as having provided the selection pressure for the large neocortex of our species. Tool use and social interaction are not two mutually exclusive phenomena but seem to have developed together. Thus, we can see Wolpert's idea as an important addition or qualification to other arguments that may be able to help explain beliefs about supernatural agents.

The hyperactivity in HADD, HUI, and HTR means that they produce many false positives: people perceive agents where there are none, ascribe intentionality to events and structures that are purely mechanical, and use teleofunctional reasoning to explain the natural world. As natural-born tool users, humans see intelligent agency and design everywhere. As humans are the prototype of an agent, humans attribute at least some humanlike features also to the agents HADD postulates (Boyer 1996c; see Guthrie 1993; Richert and Barrett 2005).

However, supernatural agent concepts are not derived from these false positives. First, ambiguous perception and overextended attribution of intentionality and design in themselves do not contain enough information for forming a persisting agent concept. Second, the supernatural agent concepts of religious traditions are abstractions from a large number of individual mental representations that are communicated among persons. Therefore, we cannot explain these concepts only by referring to ambiguous individual perception; we have to explain why they have become widespread in populations (see Barrett 1998, 617, 2004b, 41, 43; Boyer 1994a,b).

The notions of HADD, HUI, and HTR can help explain why certain kinds of supernatural agent concepts are easier to adopt than others; this ease then explains, partly at least, why these concepts are found all over the world both in ancient and modern times. Because there is a natural place for these concepts in the human mind and in human practices, they are contagious and have become widespread (Barrett 2004b, 33–39). Mind and culture thus are not so much two different levels as the endpoints on a scale from individual to public and “shared” (see Sperber 2006). The folklore scholar Lauri Honko (1962, 93–99) tried to explain the actualization of traditional beliefs in casual encounters with the spirits by referring to the psychology of perception; we can now try to do much the same with the help of cognitive and developmental psychology and evolutionary theory. That effort may turn out to have far-reaching consequences for the study of religion and culture. The next thing for me to do with this purpose in mind is to provide some criteria for calling some concepts counterintuitive.

1.4 Out of the Ordinary

1.4.1 *Intuitive Ontology and Counterintuitiveness*

Animacy, mentality, and teleology are, of course, not all that the human mind processes intuitively. Some of the basic domains of intuition are enumerated in the list that follows (see Barrett 2000, 2008; Boyer 1994b; Keil 1979, 1996; Sommers 1959). We intuitively understand that things exist in *space* and *time*; some things also have *physicality* and *solidity*, and some of them belong to the folk-biological domain of *living kinds* (plants and animals). *Animacy* is a mediating category between folk biology and *mentality*, in the sense that it is a subcategory within folk biology and may be the basis on which mind reading has developed: the idea of beliefs and desires in a way completes the understanding that something is goal-directed. *Social position*, for its part, is the domain of such social relationships as “being someone’s father” or “being the king,” and so forth. Such phenomena as cooperation and reciprocal altruism require one to have intuitions about the kind of social relationship they entail. Finally, we must add the category of *time*—though temporality is a neglected theme

