The Modularity of Social Intelligence

Gerd Gigerenzer

In: Whiten, A. & Byrne, R. W. (eds.)
Machiavellian Intelligence II • Cambridge University Press, 1997

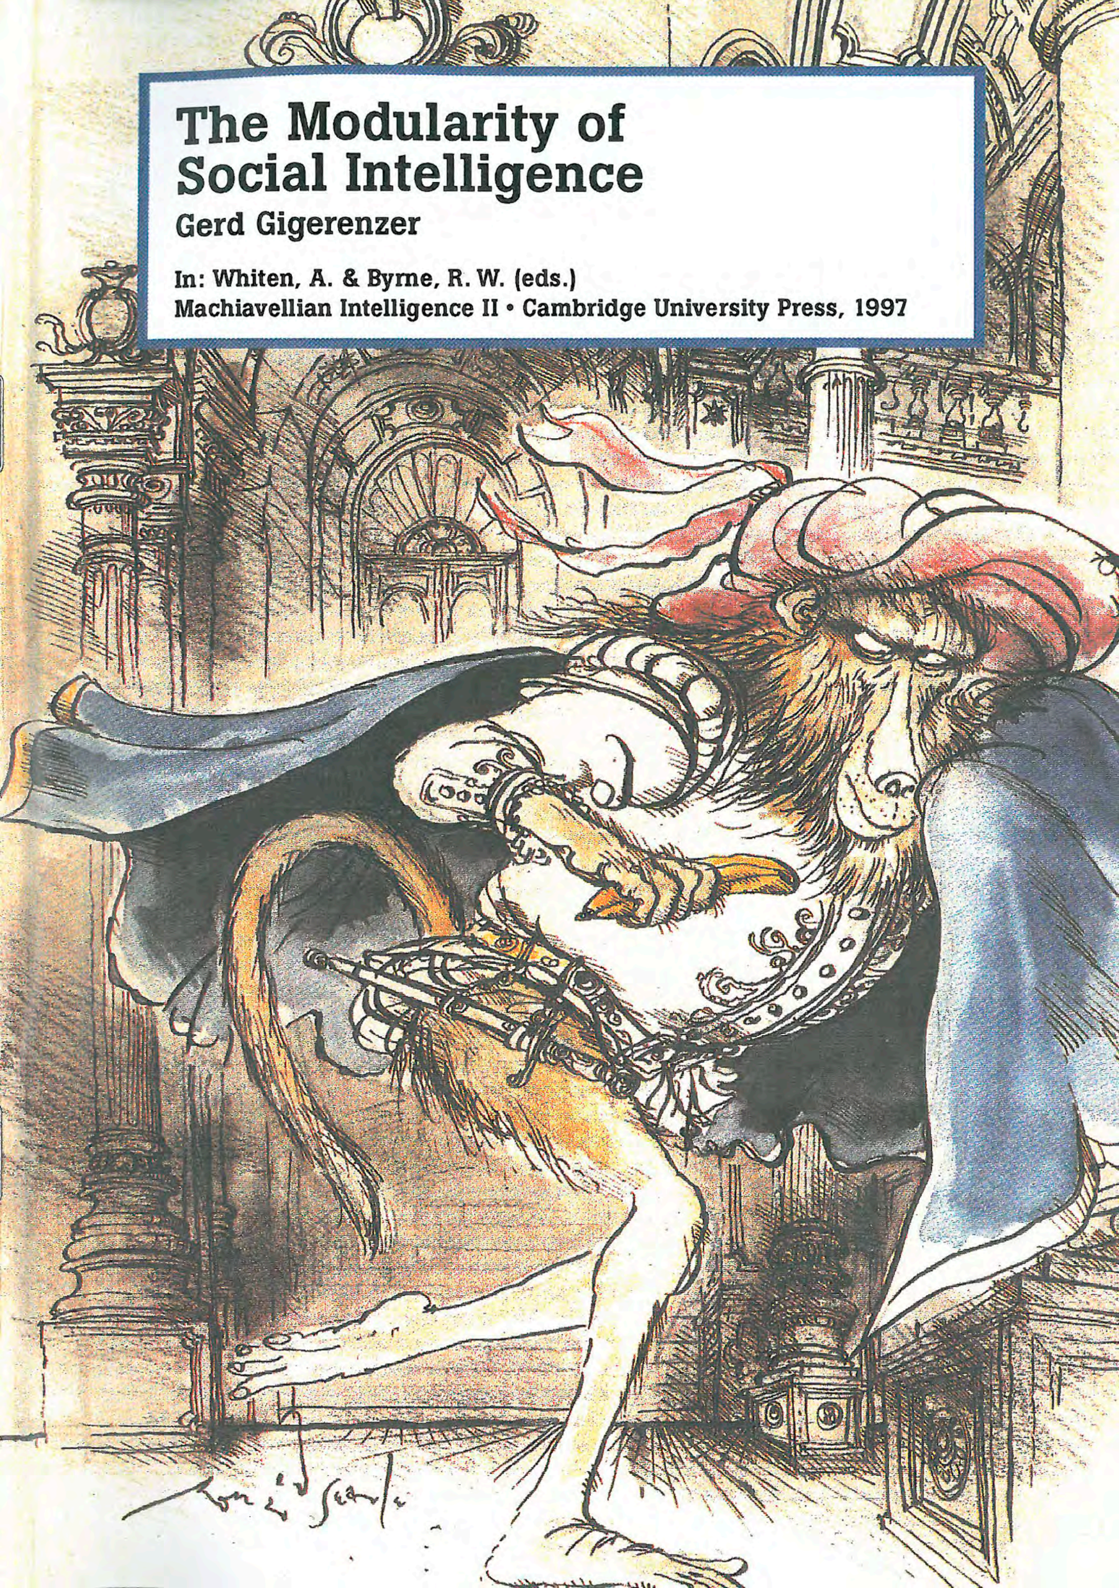

10 The modularity of social intelligence

GERD GIGERENZER

In a 'protected threat', a baboon induces a dominant member of its group to attack a third one. The baboon appeases the dominant member whom it uses as a tool to threaten the target and manoeuvres to prevent the target from doing the same (Kummer, 1988). This 'social tool use' is mastered by baboons at puberty, whereas chimpanzees are adult before they learn to use a stone as a tool for cracking hard nuts (Boesch & Boesch, 1984). Primates appear to manipulate social objects with more ease and sophistication than physical tools.

Observations such as these have suggested that primate intelligence is designed primarily for the social rather than the physical and have led to the *Machiavellian intelligence* hypothesis (Whiten & Byrne, 1988a) or *social intelligence* hypothesis (Kummer *et al.*, 1997). The term Machiavellian intelligence emphasises the besting of rivals for personal gain over co-operation, whereas the term social intelligence (which is the more general term) is neutral on the balance between exploitation and co-operation.

The social intelligence hypothesis is both stimulating and vague. It is stimulating because it reminds us that whenever psychologists study intelligence and learning in humans or animals, it is almost invariably about inanimate objects: symbols, sticks and bananas. It is vague because the nature of the intelligence it invokes is largely unclear, and as a consequence, the mechanisms of social intelligence have not yet been specified. This combination of exciting and imprecise should be an alarming signal. The social intelligence hypothesis has the seductive power of a political party with no precise programme that allows everyone unhappy with the established system to project his or her own values on it. At the risk of

Assumption 2 (and 3), that is, to assume that the selective pressures for this general intelligence stem from the social context rather than from our ancestors' use of tools (Wynn, 1988) or from foraging for hyper-dispersed and patchy foods in tropical forests (Milton, 1988). An illustration of this weaker form is the hypothesis that transitive inference is a universal cognitive process that originally evolved in the context of constructing a social dominance hierarchy: not all dyadic interactions need be observed directly; some can be inferred (Smith, 1988).

In this chapter, I will focus on the nature of social intelligence, ignore the issue of whether it preceded non-social intelligence (on this see Todd & Miller, 1995; Miller, 1997), and discuss some implications for its influence on non-social intelligence.

Is social intelligence qualitatively different from non-social intelligence?

The distinction between social and non-social intelligence is often motivated by the claim that social environments are qualitatively different from non-social environments. The social environment 'is special because it is reactive' (Chance & Mead, 1953; Whiten & Byrne, 1988a, p. 5); it 'contrasts with the physical world in that it is more challenging' (Whiten & Byrne, 1988a, p. 2); it is 'more intellectually demanding' (Humphrey, 1988, p. 15 (originally published in 1976)), and so on. I agree that social intelligence is special and cannot be subsumed with non-social intelligence under one general intelligence, but not for these reasons. These claims about the difference between the social and the physical world strike me as either untestable, false, or both. I propose to get rid of these mystical, vague justifications for a special social intelligence and to replace them with a more secure foundation: the modular view of intelligence.

Is the social environment more 'complex'?

Can we say that social environments are more 'complex' than non-social environments? Whiten & Byrne (1988b, p. 59) attempt to spell out what the vague term 'social complexity' could mean. Assume a primate distinguishes five conditions of social interaction with two possible states each, such as whether the interactant is of higher or lower rank and whether its own mother is present or not. The resulting number of possible social interactions (in our example $2^5 = 32$) is, according to Whiten and Byrne, a measure of social complexity. It is clear that this definition is one of the *perceived* complexity of a social environment, not of the environment itself that supposedly motivates a special social intelligence. Whatever the complexity of an environment *per se* is (if there is such a thing), not every species perceives and acts upon that complexity. For instance, termites vary individually, but the members of a termite colony do not seem to recognise this variation. The social intelligence of primates, in contrast, exploits individual variations among their conspecifics, although not of plants, predators and prey (Kummer *et al.*, 1997). Thus, even if one could measure the complexity of a social environment in one number, a hope for which we have no evidence, such a number would be irrelevant because what counts is the perceived complexity (or 'relational' complexity: the complexity in an environment to which an organism responds). The degree of perceived complexity, however, can hardly 'explain' why there is a particular level of social intelligence in a species, because the perceived complexity is itself dependent on, or even part of, social intelligence. Add to this the unresolved – and I believe, unresolvable – problem of finding a common denominator for measuring the complexity of social and non-social environments. Such a common denominator would be necessary to turn the statement that the social environment is more 'complex' than the physical environment into an empirically meaningful one. To illustrate this problem of comparability, can we say that triadic interactions in primates, which are said to be complex, are more complex than triadic interactions between heavenly bodies (the Three Body problem, which so far has been too complex for human minds to solve)? How could we? Complexity comparisons drive us into a conceptual cul-de-sac. I do not believe that the complexity argument is of any scientific promise for justifying a special social intelligence.

Is the social environment less predictable?

Complexity does not necessarily entail unpredictability and *vice versa*. Thus instead of complexity, one might consider the 'role of unpredictability in social interaction generating social intelligence' (Whiten & Byrne, 1988a, p. 8). Unpredictability is a particularly interesting feature of certain types of behaviour, such as Protean behaviour in predator–prey encounters (e.g. the rabbit that tries to escape the fox by running in a quasi-random pattern; Driver & Humphries, 1988). The evolution of adaptive unpredictability is an important and neglected topic (see Chapter 12). However, to propose that unpredictability is larger in interactions with the social than with the physical world seems to me empirically meaningless. What is the evidence that the social world has been less predictable than the rest of nature during primate and human evolution? I do not know of any common yardstick that would allow such a comparison. Thus, this line of justifying a special social intelligence seems to be, like the complexity argument, empirically empty, or possibly, false.

It may be false because the argument projects our 20th-century view of nature – tamed by technology, medicine and science – into the past. Until recently, nature used to be seen as fickle and unstable, willing to inflict unpredictable hazards: sudden death by plague, lightening, and flood; the permanent gamble of early infant death; unknown diseases killing livestock; famine caused by weather hazards, drought and millions of grasshoppers that surface without warning, among others. For instance, 16th-century insurers rejected the use of statistical information for calculating insurance premiums as impractical because they did not believe that the world was stable enough for the use of statistics (Daston, 1987). Why then do 20th-century researchers tend to see nature as more predictable than society? Nature may have indeed become more stable, or it may be our scientific idealisations that have made it appear to be so. Galileo's law of falling bodies is one illustration of the successful strategy of physicists to look for strong regularity and ignore the unpredictable parts of nature, such as air currents, surface area and a myriad of intervening factors. Any actual case of a falling leaf remains unpredictable by the law of falling bodies. More generally, 20th-century physics has replaced the deterministic world view of classical physics by a probabilistic view: Brownian motion and radioactive decay exemplify the unpredictability of individual particles, which become predictable only by statistical laws. Unpredictability and probability seems to be rooted not only in the human mind, but also in the constitution of the universe (Gigerenzer *et al.*, 1989).

To summarise, the claim that social interaction was once less predictable than encounters with nature, artefacts and the rest of the non-social world does not seem to be built on solid evidence. Therefore, it does not provide a foundation for a special social intelligence.

Is the social environment more intellectually challenging?

Humphrey (1988, p. 19) asserted that social systems create and maintain 'calculating beings' who 'must be able to calculate the consequence of their own behaviour, to calculate the likely behaviour of others, to calculate the balance of advantages and loss'. He concluded that 'here at last the intellectual faculties required are of the highest order'. Humphrey seems to assume that a complicated social system creates complicated intellectual operations ('creative intelligence', Humphrey, 1988, p. 16). This is not necessarily so. Herbert Simon (e.g. 1982) has argued that humans typically do not calculate the balance of advantages and losses in a complicated, 'optimal' way, such as prescribed by subjective expected utility theory (or its variants), but use simple 'satisficing' algorithms instead. Satisficing, a blend of sufficing and satisfying, is a word of Scottish origin that Herbert Simon introduced to characterise algorithms that successfully deal with conditions of limited time, knowledge, or computational capacities.

For instance, a person may accept the first mate who satisfies an aspiration level rather than going through the 'optimal' procedure of listing all candidates, listing all possible consequences associated with each candidate, determining the probabilities and utilities of each of these consequences, multiplying probabilities by utilities, calculating the expected utilities of all candidates and finally choosing the one with the maximum value. Even if one had the time to go through this 'optimal' procedure, the candidates may become insulted, tire of waiting, and no longer be available after all the candidates have been vetted. (There seem to be exceptions. Once I asked a rational-choice theorist from the University of Texas at Austin in which important decisions he had explicitly followed the expected utility procedure, and he answered when he chose his wife.) Nevertheless, from economics to cognitive science to optimal foraging theory, the fiction of 'optimal' computations is generally upheld. The fear is that satisficing algorithms would not be very intelligent. However, Dan Goldstein and I showed that this fear is unwarranted; simple satisficing strategies can make about as many accurate inferences about real-world environments as computationally costly rational calculations, and in less time and with less knowledge (Gigerenzer & Goldstein, 1996). If social situations can be mastered by fast and frugal satisficing algorithms, then there is little reason why more costly 'optimal' computations should have evolved instead.

Thus a system that looks complex may not demand those sophisticated (and mystical) cognitive processes that Humphrey termed 'creative intelligence'; simple processes may do the job as well. However, unless we specify these intelligent processes, complicated or simple, and actually test and simulate their performance on some adaptive task, we will not know. Therefore, a research programme on social intelligence needs to work out these satisficing algorithms that may be the building blocks of social intelligence.

To summarise, the possibility of a special social intelligence, different from technical, ecological, and other forms of intelligence, has been attributed to the alleged facts that social environments are more 'complex', 'unpredictable', or 'challenging'. Justifications such as these sound plausible but under closer inspection turn out to be either empirically empty (untestable and vague) or possibly false (they project 20th-century control over nature into the past). Last but not least, none of these supposed reasons for a special social intelligence has helped to specify the mechanisms of social intelligence, Machiavellian or otherwise.

I propose to drop this rhetoric and instead to start with a modular concept of intelligence. Modularity can justify why separate social intelligences (note the plural) are needed, and can suggest what the mechanisms actually could be. The notion of modularity is nothing new, but, as the introduction to a recent book on this topic illustrates, it means many things to many people (Hirschfeld & Gelman, 1994). There exist several important approaches to modularity (e.g. Millikan, 1984; Leslie, 1994; Pinker, 1994; Baron-Cohen, 1995). It will soon become clear which ones inspired my own speculations.

Modules for social intelligence

The modular version of the social intelligence hypothesis I propose assumes that social intelligences come in the plural, as do non-social intelligences. I will now describe the modular organisation of social intelligence. Reader be warned: from here on is speculation.

The thesis that social intelligence is modular is motivated by two reasons: the shortsightedness of natural selection and the combinatorial explosion of intelligent systems. Natural selection works without a big plan but results in specific adaptations accumulated over generations. Thus it seems unlikely that natural selection designed a general-purpose intelligence that embodies, say, the Piagetian formal operations or a Bayesian inference machine. Even if this had happened, such a general-purpose intelligence runs into the problem of combinatorial explosion, as evidenced by the frame problem in artificial intelligence: unless the infinite possibilities to combine elements and relations in a general-purpose system are drastically reduced by semantic constraints, an organism would be paralysed and unable to react in time. For instance, unless attention is constrained to specific types of interactors and interactions, and semantic structure is *a priori* built in that tells the organism what to learn, what to look for, and what to ignore, an intelligent organism would be unable to perform even the most elementary tasks: to detect predators, prey and mates, and be fast enough to survive and reproduce.

I propose the following assumptions about the nature of modular social intelligence:

1 A module for social intelligence is a faculty, not a general factor.

Intelligence is often assumed to be of one kind, one general ability that helps an organism cope with all situations – such as Francis Galton's 'natural ability,' Spearman's general intelligence factor g, and the numberless definitions that start with 'intelligence is the general ability to . . . .' The thesis that intelligence is a unified general ability has been invented only recently, in the mid-19th century, by Francis Galton, Herbert Spencer and Hippolyte Taine, among others (Daston, 1992). The idea of one general intelligence was motivated by Darwin's theory of evolution (Galton was Darwin's cousin) and seemed to provide the missing continuum between animals and humans, as well as between human races, and last but not least, between men and women.

Such a unified general ability was alien to the earlier *faculty* psychology, which dated back to Aristotle. Faculty psychology posited a collection of faculties and talents in the mind, such as imagination, memory and judgment. These faculties organised an intricate division of mental labour, and no single one nor their sum coincided with our concept of intelligence (Daston, 1992). Faculty psychology was revived, in the language of factor analysis, in the late 1930s when L. L. Thurstone claimed about seven primary mental abilities. In the second half of the 20th century, the mind has become again a crowded place. Evidence has been announced for dozens of factors of intelligence, and Guilford & Hoepfner (1971) even claimed the confirmation of some 98 factors of cognitive ability (see Carroll, 1982). Cognitive psychologists who use experiments rather than IQ tests also divide up cognition in terms of faculties (but you will not catch one using that term): deductive reasoning, inductive reasoning, problem-solving, memory, attention, judgement and decision-making, and so forth.

A modular organisation of intelligence assumes, similar to the earlier faculty psychology before Galton, several intelligences that have special designs, and not one general-purpose intelligence. The general thesis is that social intelligence is like a Swiss army knife with several special tools rather than one single general-purpose blade (Cosmides & Tooby, 1994a).

2 A module for social intelligence is domain-specific, not domain-general.

Intelligence modules, however, are not like Thurstone's primary mental abilities and faculties such as reasoning. I distinguish between two types of faculties: *domain-specific* and *domain-general*. Faculties such as deductive reasoning, memory and numerical ability (as well as such factors as 'fluid' and 'crystallised' intelligence) are assumed to treat any content identically, that is, to operate in a domain-general way. The laws of memory, for instance, in this view, are not about *what* is memorised; they are formulated without reference to content. Fodor (1983) called these domain-general faculties 'horizontal', as opposed to 'vertical' domain-specific faculties. The modularity of social intelligence, I propose, is vertical.

The doctrine of domain-general mechanisms flourished in Skinner's behaviourism, before it was generally rejected following experimental work by John Garcia and others (e.g. Garcia & Koelling, 1966). Skinner's laws of conditioning were designed to hold true for all stimuli and responses (the assumption of *equipotentiality* of stimuli). Garcia, however, showed that when the taste of flavoured water was repeatedly paired with an electric shock immediately after tasting, rats nevertheless had great difficulty learning to avoid the flavoured water. Yet in one single trial the rat can learn to avoid the flavoured water when it is followed by experimentally induced nausea, even if the nausea is 2 hours later (thus violating the law of contiguity). The rat seems to be 'prepared' by natural selection to form some associations rapidly but not others. Learning through imitation (rather than direct experience) is also reported to be domain-specific. Rhesus monkeys, for instance, reared in the laboratory, exhibit no fear toward venomous snakes. However, one will show fear if it sees another monkey emitting a fear reaction toward snakes. Yet the monkey does not become afraid of just any stimulus: if it sees another monkey emit a fear reaction toward a flower, it does not acquire a fear of flowers (Mineka & Cook, 1988; Cosmides & Tooby, 1994b). Learning by imitation of others, like learning by association, is simultaneously *enabled* and *constrained* by specific 'expectations' of what to avoid, what to fear, or more generally, what causal connections to establish. Without domain-specific mechanisms, an organism would not 'know' what to look for, nor which of the infinite possible causal connections to check. Such an organism would be paralysed by data analysis like the quantophrenic researcher who measures everything one can think of, computes correlation matrices of dinosaurian dimensions, and blindly searches for significant correlations. Despite the available evidence to the contrary, Skinner's ideal of domain generality has survived the cognitive revolution and is flourishing in present-day conceptions of the mind.

Domain generality is possibly the most influential and suspect idea in 20th-century psychology. Psychologists love to organise their field by horizontal faculties such as attention, memory, perception, problem-solving, and judgement and decision making. Terms such as these organise the chapter structure of textbooks, the specialties of scientific journals, the divisional structure in grant agencies, and the self-identity of numerous colleagues. Pick a psychologist at random at a conference and ask what she or he does. Most likely, you will get an answer such as 'I am a memory person' or 'I am in judgement and decision-making'. Psychologists tend to identify with horizontal faculties, not with domains.

I propose, in contrast, that modules for social intelligence are domain-specific. How should we think about these modules? Fodor (1983), a vehement proponent of modularity, has argued that modularity is restricted to input systems (the senses) and language, whereas central processes such as reasoning are domain-general. I term this the 'weak modularity thesis'. In his view, modules are specifically designed mechanisms for voice recognition in conspecifics, for face recognition in conspecifics, and colour perception, among others. I disagree with Fodor's opposition between modular sensory processes (and language) and general-purpose central processes. Social intelligence involves both perceptual processes and mechanisms for reasoning and inductive inference. For instance, assume there is a module for social contracts, that is, a module that enables co-operation between unrelated conspecifics for their mutual benefit. Such a module would need to incorporate both 'central' processes, such as cost–benefit computations and search algorithms for information that could reveal that one is being cheated, and sensory processes such as face recognition. Without both 'peripheral' and 'central' mechanisms, neither social contracts nor cheating detection would be possible.

What I call the 'strong modularity thesis' postulates that modules include central processes as well as sensory mechanisms (and language). The function of modules is not tied to 'peripheral' as opposed to 'central' processes. Rather, their function is to solve specific problems of adaptive significance, and to do this quickly. A problem of adaptive significance can be described as an evolutionarily recurrent problem whose solution promoted reproduction (Cosmides & Tooby, 1994b). Candidates include coalition forming and co-operation, foraging, predator avoidance, navigation and mate selection. To solve such problems, modules need to combine 'peripheral' and 'central' processes. Thus, the domains (more precisely, the 'proper' domains, see next section) of modules are important adaptive problems and not just perceptual (plus language) tasks.

3 A module for social intelligence has a proper and an actual domain.

Assume there is a social intelligence module designed for handling social contracts in a hunter–gatherer society. The *proper* domain of the module may have been the exchange of food for the mutual benefit of both parties involved in the contract (because food sharing is not too common among animals, an alternative hypothesis would be that the proper domain concerned social services such as alliance formation, see Harcourt & de Waal, 1992). Generations later, currency has been developed, and the module's representation of possible benefits and costs exchanged in a social contract needs to be expanded to tokens that can be exchanged for benefits. Soon economic systems will be invented where the exchange of hard currency is no longer the norm, and benefits and costs become more and more abstract. The *actual* domain of the module has shifted from the exchange of grain and meat to buying futures and options. The mechanisms of the module, however, perform largely the same task: a routine that leads individuals to enter into social contracts (sharing) when resources (such as meat) are highly variable and scarce; a representation of what the benefits and costs are for oneself and for one's kin; perceptual algorithms and a memory that allows identification of the partner by face, voice, or name recognition; and search processes that look for information that could reveal that one is being cheated.

Thus the *proper* domain of a module is that for which the module actually evolved; the *actual* domain is one to which the module is transferred or extended, following changes in the environment (Sperber, 1994). By being transferred to new domains, the mechanisms of the modules themselves may change; adaptations (modules) can become exaptations for new adaptations (Gould & Vrba, 1982). The distinction between proper and actual domain is a matter of degree rather than of kind, and the actual domain is most likely larger than the proper domain. Modules for social intelligence in humans seem to differ from those in primates in that they have larger actual domains, and in that the actual domains may have less overlap with the proper domain.

The actual domains of modules for human social intelligence can extend beyond the human. Anthropomorphism is social intelligence reaching out beyond *Homo sapiens*. Anthropomorphism has counted as a scientific sin since the 17th century, and earlier as a theological sin; nevertheless human intelligence cannot resist projecting human social categories, intentions and morals into non-humans (Mitchell *et al.*, 1997). Darwin himself practiced empathic anthropomorphism (but not anthropocentrism) in particular with respect to dogs; animal rights activists often invoke the same sentiment. Anthropomorphism of a less empathic nature extends to phylogenetically distant species: 'Rape' in scorpion flies and 'ant slavery' are examples. In opposition to all this, behaviourism values purifying scientific language, and a story told about the Columbia University philosopher Sidney Morgenbesser illustrates this value. After B. F. Skinner gave a talk at Columbia, Sidney stood up and said: 'Professor Skinner, I always tried to understand what the essence of behaviourism is. Now I think I know. Behaviorism is the denial of anthropomorphism for humans'.

Social intelligence can reach beyond animals and still create powerful metaphors. Dawkin's (1976) 'selfish gene' lives in a world of savage competition, exploitation and deceit. Physicists, chemists and astronomers certainly censor similar anthropomorphic descriptions in press, but in scientists' informal and private conversations, intentions are frequently attributed to particles and matter (Atran, 1990).

4 A module is activated by a triggering algorithm.

Assume there is a simple social organism with two modules for social intelligence: one deals with social contracts, the other with threats. Thus this organism knows only two ways to deal with conspecifics: to engage with them in the exchange of certain goods to their mutual benefit, and to threaten individuals to get one's way (and react when others do so). As simple as the social intelligence of this organism is, the organism needs to decide when to activate the social contract module and when the threat module. All modules cannot be activated at the same time because the very advantage of modularity is to focus attention and to prevent combinatorial explosion. For instance, the social contract module focuses attention on information that can reveal that the organism is being cheated, while this information is of no relevance for a threat module. A threat module needs to attend to information that can reveal, for instance, whether the other side is bluffing, or whether high status individuals are present who could be used for 'protected threat' (Kummer, 1988).

How is one of the two modules activated? I assume that there is a triggering algorithm that attends to a small set of cues whose presence signals either threat or social contract. These signals can include facial expressions, gestures, body movements and verbal statements. Assume the organisms do have language. A simple algorithm can quickly recognise whether a verbal statement of the type 'if you do *X*, then I do *Y*' is a threat or a social contract. If *Y* is a negative consequence for me, and follows *X* in time, then I am being threatened. If *Y* is a benefit for me, and the temporal sequence can be either way, then I am being offered a social contract. I call such simple rules 'triggering algorithms' because their function is to activate a module that can focus attention, emotion and behavioural responses so that fast reaction is possible.

Triggering algorithms can err, that is, not activate the appropriate module, such as mistaking a serious threat for pretend play. The likelihood of triggering the wrong module may increase when there are more than two modules, but redundancy in cues, such as verbal cues, facial cues and gestures may reduce errors.

5 Modules are hierarchically organised.

When a mind has not just two, but a large number of modules, a single triggering algorithm may be too slow to discriminate between all possibilities simultaneously. In such a socially more intelligent mind, modules can be hierarchically connected by triggering algorithms, as in a sequential decision tree. Hierarchical organisation corresponds to the idea that species started out with a few modules, to which more specialised modules were added later in phylogeny, and to Wimsatt's (1986) notion of generative entrenchment.

Assume I march through a forest at night. Visibility is poor, a storm is coming up, and I suddenly see the contours of a large dark object that seems slowly to move. A triggering algorithm needs to decide quickly whether the object is 'self-propelled' (animal or human) or not (plant or physical object) (Premack & Premack, 1994). According to the Premacks, this decision is based on the object's motion pattern. Recall the demonstrations by Fritz Heider where the motion patterns of two points in two-dimensional space make us 'see' the points as animate or inanimate, chasing, hunting, hurting, or supporting one another (e.g. Heider & Simmel, 1944). These are beautiful demonstrations, but they include no descriptions of the algorithms that make us see all these social behaviours. How could the first triggering algorithm work? A simple algorithm would analyse only external movements (such as the direction and acceleration of the object) and not internal movements (the relative movement of the body parts). For instance, if a motion pattern centres on my own position, such as an object that circles around me, or speeds up toward me, the algorithm infers a self-propelled object. Moreover, it infers a self-propelled object that takes some interest in myself. Motion patterns that centre around the object's own centre of gravity, in contrast, indicate that the object is a plant (e.g. a tree). Now, if the motion pattern indicates that the object is self-propelled, the triggering algorithm may activate a module for unrecognised self-propelled objects. This module will immediately set the organism into a state of physiological and emotional arousal, initiate behavioural routines such as stopping and preparing for running away, and activate a second, more specialised triggering algorithm whose task is to decide whether the self-propelled object is animal or human. Assume that this second triggering algorithm infers from shape and motion information that the object is human. A module for social encounters with unknown humans is subsequently activated, which initiates a search for individual recognition in memory and may initiate an appeal for voice contact in order to find out whether the other is friend or enemy, is going to threaten or help me, and so on. This is pure speculation, but one might work out the mechanisms of a hierarchical organisation along these lines.

Modules that are hierarchically organised can act quickly, as only a few branches of the combinatorial tree need to be travelled. For instance, if the first triggering algorithm had indicated that the unknown object was not self-propelled, then all subsequent information concerning whether it is human or animal, friend or enemy, or predator or prey, could have been ignored.

6 An intelligence module works with satisficing algorithms.

There are two views about the machinery of intelligent behaviour. The classical view is that the laws of probability theory or logic define intelligent processes: intelligent agents have rules such as Bayes's rule, the law of large numbers, transitive inference and consistency built in. This was the view of the Enlightenment mathematicians (Daston, 1988), to which Jean Piaget added an ontogenetic dimension, and it still is a dominant view in economics, cognitive psychology, artificial intelligence and optimal foraging theory. For instance, Bayes's rule has been proposed to model how animals infer the presence of predators and prey (Stephens & Krebs, 1986) as well as how humans reason, categorise and memorise (Anderson, 1990). The problem with this view is that in any sufficiently rich environment, Bayesian computations become mathematically so complex that one needs to assume that minds are 'Laplacean demons' that have unlimited computational power, time and knowledge. To find a way out of this problem, researchers often make unrealistic assumptions about the structure of natural environments, namely assumptions that reduce the Bayesian computations. Anderson (1990) for instance, found himself forced to make the false assumption that environmental features would be generally independent, in order to save the fiction of the Bayesian homunculus in the mind (Gigerenzer, 1991a). Despite their psychological implausibility, Laplacean demons are bustling in contemporary theories of the mind; as models of choice, categorisation, estimation and inference, among others (Gigerenzer & Murray, 1987). The rationale seems to be this: cognition is rational, Bayes's theorem defines rationality, *ergo* cognition works with Bayes's theorem. The same rationale seems to hold for other statistical tools that turned into theories of mind, such as multiple regression, Neyman–Pearson decision techniques and analysis of variance (Gigerenzer, 1991b).

The second view, the one I propose here, is that modules of social intelligence, including the triggering algorithms, work with simple 'satisficing' algorithms instead of the costly 'optimal' algorithms. There are several reasons that favour simple and specifically designed principles rather than expensive and general ones. Firstly, there is, in fact, no single method of inference – statistical or logical – that works best in all real-world contexts (Gigerenzer *et al.*, 1989). Secondly, as mentioned before, in real-world situations, 'optimal' computations can become quickly so complex that one is forced to make highly simplifying assumptions about the environment. Thirdly, algorithms for social intelligence need to work under constraints of limited time and knowledge – for instance, one may not have the time to search for further information. Fourthly, the means and ends of social intelligence are broader than consistency (coherence) and accuracy – the accepted norms of logic and statistics. Social intelligence can involve being inconsistent (e.g. adaptive unpredictability may be optimal in competitive situations: the opponent will be unable to predict one's behaviour), taking high risks in trying to come out first (that is, options with low probabilities rather than those that maximise expected value), responding quickly rather than accurately (e.g. to make too long a pause in a conversation in order to think of the best answer can be embarrassing and seen as impolite), and making decisions that one can justify or defend afterward (Gigerenzer, 1996).

The key argument against satisficing strategies is that their simplicity raises the suspicion that they are really bad. But this need not be the case: simple principles can be quick-and-clean rather than quick-and-dirty. Consider a simple algorithm for choice called 'Take The Best,' which Dan Goldstein and I have studied in some detail (Gigerenzer & Goldstein, 1996). Take The Best infers which of two alternatives scores higher on some criterion, such as which of two food items is more dangerous, or which of two cities has a higher population. Take The Best bases decisions on one single reason, namely on the first good reason on which two alternatives differ. The first good reason can be simply that the individual does not recognise (has never heard of) one of the two alternatives. This 'recognition principle' seems to operate in domains where recognition is correlated with the variable that needs to be inferred. For instance, rats who can choose between food that they recognise and food that is new to them do not accept the new food unless they have smelled it on the breath of a fellow rat (Gallistel *et al.*, 1991).

But how good is Take The Best? We set up a contest between Take The Best and various linear integration algorithms, including multiple regression (Gigerenzer & Goldstein, 1996). The task was to infer which of two cities has the larger population. We used as a test environment all German cities with more than 100 000 inhabitants and 10 binary predictors of population (such as whether or not a city has a soccer team in the major league). Take The Best simply goes through the predictors one by one, in the order of their validity, until it finds one in which one city has a positive value (e.g. a soccer team in the major league) and the other has not (no team or unknown). Then the inference is made that the city with the positive value is the larger, and no further predictor values are searched – it takes the best predictor that makes a difference and ignores the rest. Thus, Take The Best violates the classical standards of rationality: it does not search for all evidence (here the values of two objects on all ten predictors), and it does not integrate any evidence, and variants of Take The Best can violate transitivity. The competitors – such as multiple regression – in contrast, always used all predictors and always integrated this information. The competition was performed with millions of simulated subjects with varying degrees of limited knowledge about predictor values. Take The Best matched or outperformed all competitors in accuracy, including multiple regression, and outperformed all in speed. This surprising result demonstrates that simple psychological principles can be as accurate as costly statistical algorithms of the Laplacean demon type.

If short-sighted evolution has equipped us with simple satisficing algorithms rather than with the collected works of logic and probability theory, this result indicates that we need not necessarily worry about human rationality. The challenges are to understand what these satisficing algorithms are and to describe the structure of the environments in which they can perform well, and where they cannot. My proposal is that both the triggering algorithms and the mechanisms of the module are satisficing algorithms.

7 Satisficing algorithms combine cognitive, emotional and motivational tools to guide inference and behaviour.

In each module, various cognitive, emotional, behavioural and motivational processes are wired together. A social contract module, for instance, includes perceptual machinery to recognise different individuals; a long-term memory that stores the history of past exchanges with other individuals in order to know when to co-operate, when to defect and when to punish for defection; knowledge about what constitutes a benefit and what a cost for oneself; and emotional reactions such as anger that signal to others that one will go ruthlessly after cheaters (Cosmides & Tooby, 1992; Gigerenzer, 1995). Satisficing algorithms can use cognitive, emotional, or motivational processes vicariously as means to achieve a goal. For instance, emotional reactions (such as disgust in matters of food) can substitute for learning from experience when events are too rare or too deadly for individual learning.

The very challenge of a modular concept of social intelligence is that it crosses the established boundaries of horizontal faculties. It opens up a conception of intelligence that integrates cognition with, rather than setting it apart from, other adaptive functions such as motivations and emotions.

Experimental tests

This sketch of a modular social intelligence is speculative and there is room for improvement. Does such a theory produce testable predictions? It does, and I give one example: the thesis that there is a module for social contracts in humans. Humans are among the very few species that engage in co-operation between unrelated individuals, such as social contracts. But co-operation (reciprocal altruism) cannot evolve in the first place unless one can detect cheaters (Trivers, 1971). The thesis is that humans are equipped with a module for social contracts, which has special mechanisms, such as weighing benefits against costs, recognising individuals, and detecting when other individuals cheat. The theory was proposed first by Cosmides (1989) and has been further elaborated and experimentally tested (e.g. Cosmides & Tooby, 1989, 1992; Gigerenzer & Hug, 1992; Gigerenzer, 1995).

One testable implication of a social contract module is that if an interaction is classified as a social contract, the module will search for information that could reveal that the other party is cheating. In a social contract of the type 'If you take the benefit, then you must pay the costs', cheating is defined as 'benefit taken and costs not paid'.

Let us apply this hypothesis to a standard reasoning task that seems to have absolutely nothing to do with social intelligence. The task is known as the four card problem (or the Wason selection task). There is a conditional rule of the type 'if *P* then *Q*' such as 'if there is a "4" on one side of the card, then there is a "*B*" on the other'. There are four cards that have a number on one side and a letter on the other. The participant can read only one side, and the cards read '4', '5', '*A*', and '*B*'. The participant is asked which of these cards must be turned around in order to find out whether the rule has been violated. Peter Wason (1966), who introduced this task, and his followers believed that the correct answer is to turn around the '4' and the '*A*' cards, because in propositional logic, a conditional statement 'if *P* then *Q*' is violated if and only if '*P* and not-*Q*' is obtained. An avalanche of experimental studies showed that people sometimes reasoned according to propositional logic, sometimes not, depending on the content of the *P*s and *Q*s. For instance, only about 10% of participants picked the '*P*' and the 'not-*Q*' cards in rules about numbers and letters, as above, whereas this number skyrocketed to 75% when the rule was 'if an employee works on the weekend, then that person gets a day off during the week' (Gigerenzer & Hug, 1992). In the latter experiment, the four cards represented information about four previous employees who had worked (had not worked) on the weekend and did get (did not get) a day off. For more than two decades psychologists puzzled about why people sometimes reason logically and sometimes not. But no convincing answer was found. Propositional logic, which many researchers (mis)take as the yardstick of sound reasoning did not tell – it is pure syntax and mute on semantics.

A modular theory of intelligence, however, postulates that even in the cases where judgements correspond to the logical '*P* and not-*Q*,' people are not necessarily performing logical operations. If the rule is a social contract, such as the day-off rule, then a social contract module is activated that searches for information that could reveal that the other party is cheating. The 75% who chose the 'worked on the weekend' and 'did not get a day off' cards did not perform a logical abstraction, that is, they did not infer that 'if *P* then *Q*' is violated if and only if *P* and not-*Q* was obtained. These participants chose the '*P*' and 'not-*Q*' card for a different reason: '*P*' and 'not-*Q*' happens to coincide with 'benefit taken' and 'costs not paid', what the module is attentive to. One way to test these competing explanations – modular social intelligence versus general-purpose logic – is to switch perspectives: to tell half of the participants that they are an employee, and the other half that they are the employer, leaving everything else as it was. If people reason by propositional logic, this should make no difference, because logic knows no perspective. If people, however, reason by a social contract module, then the perspective is the very essence. The experimental results were consistent with the postulate of a social contract module: 75% of the employees choose the cards 'worked on the weekend' and 'did not get a day off' but only 2% of the employers. The majority of employers selected the 'did not work on the weekend' and 'did get a day off' cards, which looks like the logical 'not-*P* and *Q*'. Thus, if one believes in logic as a general-purpose norm for good reasoning, then people's performance seems to be inconsistent: sometimes they reason logically and sometimes they do not. If one assumes a social contract module, however, then people are quite consistent: whether employee or employer, they search for 'benefit taken' and 'costs not paid' (Gigerenzer & Hug, 1992).

Note that these results can not be explained away by saying that the conditional in a social contract is a logical biconditional. If people treat the statement as a logical biconditional, then they would have to turn around all four cards (because both '*P* and not-*Q*' and 'not-*P* and *Q*' violate a biconditional). Not one single employee did so, and very few of the employers. Note that participants were not looking for cheaters *per se*, but only for being cheated themselves. Our participants were not reasoning by a Kantian moral, but by a Machiavellian intelligence.

This and other experiments illustrate how social intelligence can explain apparently contradictory data in a standard laboratory task that has been assumed to reveal human irrationality. The thesis that the evolutionary theory of reciprocal altruism can enlighten us about cognitive psychology has stirred up a virulent controversy (e.g. Cheng & Holyoak, 1989; Pollard, 1990). The thought that social intelligence might exist in the syntactical world of research on reasoning has been treated by many like a virus infection. In this world where the content of reasoning is often arbitrary and irrelevant for the theory, the debate is still about which syntactical model is the right one: is it mental logic (e.g. Rips, 1994) or mental models (e.g. Johnson-Laird, 1983)? But the data, such as on the 'Machiavellian' effect of perspective on reasoning, has now been replicated several times, and some researchers have begun to adjust their theories to these data (e.g. Holyoak & Cheng, 1995a,b; Cummins, 1996).

Summing up

In 1976, Nick Humphrey wrote a stimulating essay on the social function of intellect that provided 'the single most important seed' (Byrne & Whiten, 1988a, p. 1) for *Machiavellian Intelligence*. Humphrey's paper contains stimulating ideas about social intelligence, but also the seed of a sterile research methodology. In the postscript we are told that his 'central thesis' demands 'that there should be a positive correlation across species between "social complexity" and "individual intelligence" ' (Humphrey, 1988, p. 26 (originally published in 1976)). A laboratory test of social skill, Humphrey proposes, is urgently needed, which 'ought, if I am right, to double as a test of "high-level intelligence" ' (p. 26). This proposal, as innocent as it looks, steers toward repeating a grave error that has swamped research on (non-social) intelligence tests. The error is to start with no theory but seductive everyday concepts such as social skill and creative intelligence, then go on to design tests to measure these vague concepts, and to pray to heaven that these tests miraculously transform loose thinking into precise mechanisms of intelligence.[1] Despite prayers that were backed up by statistics, some 90 years of factor analysing and correlating IQ tests has not noticeably increased our understanding of the mechanisms of human intelligence. I fear that Humphrey's proposal to look for correlations between some tests for social complexity, social skill and individual intelligence will be doomed to the same failure.

The alternative is to start boldly and theoretically. The challenge is to design the possible mechanisms of social intelligence, and to test these by means of experiment, observation and simulation. The mechanisms of a modular social intelligence I have outlined, as speculative as they are, can serve as a start. We will have to take up some hard questions. What are the domains (proper and actual) for a given species? What is the mechanism of a module? What is the algorithm that triggers a module? If we join forces, we can do it.

Acknowledgements

I would like to thank Arnold Davidson, Ulrich Hoffrage, Geoffrey Miller, Anita Todd, Frans de Waal, Bill Wimsatt and Andy Whiten for their helpful comments on earlier drafts of this chapter.

References

Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Lawrence Erlbaum Associastes.

Atran, S. (1990). *Cognitive Foundations of Natural History: Towards an Anthropology of Science*. Cambridge: Cambridge University Press.

Baron-Cohen, S. (1995). *Mindblindness*. Cambridge, MA: MIT Press.

Binet, A. & Simon, T. (1914). *Mentally Defective Children*. London: Edward Arnold.

---

[1] The notion of test intelligence, left undefined in its content, has had many faces, including the moral and the social (Daston, 1992). The creators of the first intelligence tests, Binet & Simon (1914), for instance, asked questions concerning social skills, such as 'Why should one judge a person by his acts rather than by his words?' Louis Terman (1916), in the first edition of the Stanford–Binet Test, expressed the intimate link between lack of intelligence and of morally inappropriate behaviour as follows: 'Every feeble-minded woman is a prostitute'. In the 1937 revision of the text (with M. A. Merrill), this sentence was deleted. Piece by piece, IQ tests became pure and puritan.

Boesch, C. & Boesch, H. (1984). Mental map in wild chimpanzees: An analysis of hammer transports for nut cracking. *Primates*, 25, 160–70.

Byrne, R. & Whiten, A. (eds) (1988). *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*. Oxford: Clarendon Press.

Carroll, J. B. (1982). The measurement of intelligence. In *Handbook of Human Intelligence* ed. R. J. Sternberg, pp. 29–120. Cambridge: Cambridge University Press.

Chance, M. R. A. & Mead, A. P. (1953). Social behavior and primate evolution. *Symposia of the Society for Experimental Biology, Evolution*, 7, 395–439.

Cheng, P. W. & Holyoak, K. J. (1989). On the natural selection of reasoning theories. *Cognition*, 33, 285–313.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? *Cognition*, 31, 187–276.

Cosmides, L. & Tooby, J. (1989). Evolutionary psychology and the generation of culture, Part II. Case study: A computational theory of social exchange. *Ethology and Sociobiology*, 10, 51–97.

Cosmides, L. & Tooby, J. (1992). Cognitive adaptations for social exchange. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, ed. J. H. Barkow, L. Cosmides & J. Tooby, pp. 163–228. Oxford: Oxford University Press.

Cosmides, L. & Tooby, J. (1994a). Beyond intuition and instinct blindness: Towards an evolutionary rigorous cognitive science. *Cognition*, 50, 41–77.

Cosmides, L. & Tooby, J. (1994b). Origins of domain specificity: The evolution of functional organization. In *Mapping the Mind: Domain Specificity in Cognition and Culture*, ed. L. A. Hirschfeld & S. A. Gelman, pp. 85–116. Cambridge: Cambridge University Press.

Cummins, D. D. (1996). Evidence for the innateness of deontic reasoning. *Mind and Language*, 11, 160–90.

Daston, L. (1987). The domestication of risk: Mathematical probability and insurance 1650–1830. In *The Probabilistic Revolution: Volume 1. Ideas in History*, ed. L. Krüger, L. J. Daston & M. Heidelberger, pp. 237–60. Cambridge, MA: MIT Press.

Daston, L. (1988). *Classical Probability in the Enlightenment*. Princeton, NJ: Princeton University Press.

Daston, L. (1992). The naturalized female intellect. *Science in Context*, 5, 209–35.

Dawkins, R. (1976). *The Selfish Gene*. Oxford: Oxford University Press.

Driver, P. M. & Humphries, D. A. (1988). *Protean Behavior: The Biology of Unpredictability*. Oxford: Oxford University Press.

Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.

Gallistel, C. R., Brown, A. L., Carey, S., Gelman, R. & Keil, F. C. (1991). Lessons from animal learning for the study of cognitive development. In *The Epigenesis of Mind: Essays on Biology and Cognition*, ed. S. Carey & R. Gelman, pp. 3–36. Hillsdale, NJ: Lawrence Erlbaum Associates.

Garcia, J. & Koelling, R. A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science*, 4, 123–4.

Gigerenzer, G. (1991a). Does the environment have the same structure as Bayes' theorem? *Behavioral and Brain Sciences*, 14, 495.

Gigerenzer, G. (1991b). From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological Review*, 98, 254–67.

Gigerenzer, G. (1995). The taming of content: Some thoughts about domains and modules. *Thinking and Reasoning*, 1, 324–33.

Gigerenzer, G. (1996). Rationality: Why social context matters. In *Interactive Minds: Life-span Perspectives on the Social Foundation of Cognition*, ed. P. Baltes & U. M. Staudinger, pp. 319–46. Cambridge: Cambridge University Press.

Gigerenzer, G. & Goldstein, D. (1996). Reasoning the fast and frugal way: Models for bounded rationality. *Psychological Review*, 103, 650–69.

Gigerenzer, G. & Hug, K. (1992). Domain-specific reasoning: Social contracts, cheating, and perspective change. *Cognition*, 43, 127–71.

Gigerenzer, G. & Murray, D. J. (1987). *Cognition as Intuitive Statistics*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J. & Krüger, L. (1989). *The Empire of Chance: How Probability Changed Science and Everyday Life*. Cambridge: Cambridge University Press.

Gould, S. J. & Vrba, E. S. (1982). Exaptation – a missing term in the science of form. *Paleobiology*, 8, 4–15.

Guilford, J. P. & Hoepfner, R. (1971). *The Analysis of Intelligence*. New York: McGraw-Hill.

Harcourt, A. H. & de Waal, F. B. M. (eds) (1992). *Coalitions and Alliances in Humans and Other Animals*. Oxford: Oxford University Press.

Heider, F. & Simmel, M. (1944). An experimental study of apparent behavior. *American Journal of Psychology*, 57, 243–59.

Hirschfeld, L. A. & Gelman, S. A. (1994). Toward a topography of mind: An introduction to domain specificity. In *Mapping the Mind: Domain Specificity in Cognition and Culture*, ed. L. A. Hirschfeld & S. A. Gelman, pp. 3–35. Cambridge: Cambridge University Press.

Holyoak, K. J. & Cheng, P. W. (1995a). Pragmatic reasoning with a point of view. *Thinking and Reasoning*, 1, 289–313.

Holyoak, K. J. & Cheng, P. W. (1995b). Pragmatic reasoning from multiple points of view: A response. *Thinking and Reasoning*, 1, 373–88.

Humphrey, N. K. (1988). The social function of intellect. In *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*, ed. R. W. Byrne & A. Whiten, pp. 13–26. Oxford: Clarendon Press. (Original work published in 1976).

Johnson-Laird, P. N. (1983). *Mental Models*. Cambridge, MA: Harvard University Press.

Kummer, H. (1988). Tripartite relations in hamadryas baboons. In *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*, ed. R. W. Byrne & A. Whiten, pp. 113–21. Oxford: Clarendon Press.

Kummer, H., Daston, L., Gigerenzer, G. & Silk, J. (1997). The social intelligence hypothesis. In *Human by Nature: Between Biology and the Social Sciences*, ed. P. Weingart, P. Richerson, S. D. Mitchell & S. Maasen, pp. 157–79. Hillsdale, NJ: Lawrence Erlbaum Associates.

Leslie, A. (1994). ToMM, ToBY, and Agency: Core architecture and domain specificity. In *Mapping the Mind: Domain Specificity in Cognition and Culture*, ed. L. A. Hirschfeld & S. A. Gelman, pp. 119–48. Cambridge: Cambridge University Press.

Miller, G. F. (1997). A review of sexual selection and human evolution: How mate choice shaped human nature. In *Evolution and Human Behavior: Ideas, Issues, and Applications*, ed. C. Crawford & D. Krebs. Hillsdale, NJ: Lawrence Erlbaum Associates. (In press).

Millikan, R. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.

Milton, K. (1988). Foraging behaviour and the evolution of primate intelligence. In *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*, ed. R. W. Byrne & A. Whiten, pp. 285–305. Oxford: Clarendon Press.

Mineka, S. & Cook, M. (1988). Social learning and the acquisition of snake fear in monkeys. In *Social learning: Psychological and Biological Perspectives*, ed. T. R. Zentall & B. G. Galef, pp. 51–73. Hillsdale, NJ: Lawrence Erlbaum Associates.

Mitchell, S. D., Daston, L., Sesardic, N., Sloep, P. & Gigerenzer, G. (1997). In the service of pluralism and interdisciplinarity. In *Human by Nature: Between Biology and the Social Sciences*, ed. P. Weingart, P. Richerson, S. D. Mitchell & S. Maasen, pp. 103–50. Hillsdale, NJ: Lawrence Erlbaum Associates.

Pinker, S. (1994). *The Language Instinct*. London: Penguin Press.

Pollard, P. (1990). Natural selection for the selection task: Limits to social exchange theory. *Cognition*, 36, 195–204.

Premack, D. & Premack, A. J. (1994). Moral belief: Form versus content. In *Mapping the Mind: Domain Specificity in Cognition and Culture*, ed. L. A. Hirschfeld & S. A. Gelman, pp. 149–68. Cambridge: Cambridge University Press.

Rips, L. J. (1994). *The Psychology of Proof: Deductive Reasoning in Human Thinking*. Cambridge, MA: MIT Press.

Simon, H. A. (1982). *Models of Bounded Rationality*. Cambridge, MA: MIT Press.

Smith, P. K. (1988). The cognitive demands of children's social interactions with peers. In *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*, ed. R. W. Byrne & A. Whiten, pp. 94–109). Oxford: Clarendon Press.

Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In *Mapping the Mind: Domain Specificity in Cognition and Culture*, ed. L. A. Hirschfeld & S. A. Gelman, pp. 39–67. Cambridge: Cambridge University Press.

Stephens, D. W. & Krebs, J. R. (1986). *Foraging Theory*. Princeton, NJ: Princeton University Press.

Terman, L. M. (1916). *The Measurement of Intelligence*. Boston, MA: Houghton Mifflin Co.

Terman, L. M. & Merrilll, M. A. (1937). *Measuring Intelligence*. Boston, MA: Houghton Mifflin Co.

Todd, P. M. & Miller, G. F. (1995). The role of mate choice in biocomputation: Sexual selection as a process of search, optimization, and diversification. In *Evolution and Biocomputation: Computational Models of Evolution. Lecture Notes in Computer Science no. 899*, ed. W. Banzof & F. H. Eeckman, pp. 169–204. New York: Springer-Verlag.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46, 35–57.

Wason, P. C. (1966). Reasoning. In *New Horizons in Psychology*, ed. B. M. Foss, pp. 135–51. Harmondsworth: Penguin.

Whiten, A. & Byrne, R. W. (1988a). The Machiavellian intelligence hypotheses. In *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*, ed. R. W. Byrne & A. Whiten, pp. 1–9. Oxford: Clarendon Press.

Whiten, A. & Byrne, R. W. (1988b). Taking (Machiavellian) intelligence apart. In *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*, ed. R. W. Byrne & A. Whiten, pp. 50–65. Oxford: Clarendon Press.

Wimsatt, W. C. (1986). Developmental constraints, generative entrenchment, and the innate-acquired distinction. In *Integrating Scientific Disciplines*, ed. P. W. Bechtel, pp. 185–208. Dordrecht: Martinus-Nijhoff.

Wynn, T. (1988). *The Evolution of Spatial Competence*. Urbana, IL: University of Illinois Press.