

An Evolutionary Perspective on Testimony and Argumentation

Dan Sperber
Institut Jean Nicod, Paris

INTRODUCTION

In the preface of his seminal *Knowledge in a Social World*, Alvin Goldman writes:

Traditional epistemology has long preserved the Cartesian image of inquiry as an activity of isolated thinkers, each pursuing truth in a spirit of individualism and pure self-reliance. This image ignores the interpersonal and institutional contexts in which most knowledge endeavors are actually undertaken. Epistemology must come to grips with the social interactions that both brighten and threaten the prospect for knowledge.¹

In chapters 4 and 5, he discusses two generic social practices, testimony, i.e., the transmission of observed (or allegedly observed) information from one person to others, and argumentation, i.e., the defense of some conclusion by appeal to a set of premises that provide support for it. In discussing these practices, Goldman has many important things to say about the way they brighten the prospect for knowledge, and very little about the way they threaten it. I would like to slightly redress the balance and put a touch of gray in Goldman's rosy picture by considering testimony and argumentation in the light of some evolutionary considerations.

My main claim will be that a significant proportion of socially acquired beliefs are likely to be false beliefs, and this not just as a result of the malfunctioning, but also of the proper functioning of social communication. I will argue in particular that the cognitive manipulation of others is one of the effects that makes the practices of testimony and argumentation adaptive. This contributes to explaining why these practices have evolved and stabilized among humans. To highlight the claim, I start by contrasting social with individual mechanisms of belief production, arguing that individual mechanisms are, under normal conditions and in the absence of social interference, reliable sources of true beliefs. Humans, being permanently immersed in society and culture, are, even when on their own, the locus of ongoing cultural processes, and therefore never good examples of truly individual systems of belief production in the intended sense. The contrast I draw is not, therefore, between human individual and social cognition; it is between ideal types. Since, moreover, I mention individual cognition just for the sake of this contrast, I will not spend time defending or hedging the evolutionary psychology approach I adopt on the topic.

Cognitive systems found in individual organisms are biological adaptations. Adaptations are traits that have evolved and stabilized because, by producing some characteristic effect, they have contributed to the fitness of the organisms endowed with them. Producing this fitness-enhancing effect can be described as the function of the adaptation, and I will be using “function” in this sense.² Roughly, the function of a cognitive system is to provide the organism with information about itself and its environment and thus guide its behavior. There may be cases and situations where it is adaptive for a cognitive system to introduce systematic biases, for instance, of excessive caution or on the contrary of overconfidence,³ but such cases are, I believe, marginal. We should generally expect the beliefs produced by an evolved cognitive system to be true. In other terms, cognitive systems are basically producers of knowledge. Of course, their function is not to produce knowledge per se, let alone scientific knowledge. It is to produce knowledge relevant to the organism’s welfare. They do so reliably in the kind of environment in which they have evolved. Put in a different type of environment, whether by historical accident or by experimental design, stimulated by phenomena the representation of which is irrelevant to the organism’s welfare, cognitive systems may well become quite unreliable. For instance, perceptual illusions, which are very rare in a natural and familiar environment, may be common in artificially devised settings.

A normative evaluation of evolved cognitive systems will find them to be, in the performance of their function, at least good enough to make them advantageous to the organisms endowed with them (or else they would have been selected out). Given the high risks involved in moving around (in contrast with the plant strategy of staying put and letting things happen), the

cognitive systems on which self-mobile organisms rely must be quite good at producing genuine information rather than errors. However, the function of these systems can be performed by means of an articulation of task- and domain-specific subsystems.⁴ Natural individual cognition is likely therefore to produce true beliefs of a very limited variety and import, nothing to wax epistemological about. Beliefs and belief-producing systems worth a philosopher's evaluation come with communication, language, and culture.

Communication might be seen as a wonderful extension of individual cognition, a kind of "cognition by proxy." A communicating organism is not limited to information derived from its own perceptions and inferences. It can benefit from the perceptions and inferences of others. Of course, it risks suffering from the others' mistakes, but to the extent that individual cognition is reliable, so should communication be, or so the story goes. An early defender of such a view is Thomas Reid—approvingly quoted by Goldman⁵—who maintained:

The wise and beneficent Author of nature, who intended that we should be social creatures, and that we should receive the greatest and most important part of our knowledge by the information of others, hath, for these purposes implanted in our natures two principles that tally with each other. The first of these principles is a propensity to speak truth. . . . [The second principle] is a disposition to confide in the veracity of others, and to believe what they tell us.⁶

In stark contrast to this view, Dawkins and Krebs, in their famous article "Animal Signals: Information or Manipulation,"⁷ have argued that the prime function of communication is not information but manipulation of others. They were focusing on the interests of the signaler as a driving force in the evolution of signals. These interests are generally different from the receiver's. Both Reid's view and Dawkins and Krebs's original view were too extreme. In a later paper, "Animals Signals: Mind-reading and Manipulation,"⁸ Krebs and Dawkins argued for taking into account both the signaler's and the receiver's perspective. This is obviously correct and should apply also for the study of human communication.

For communication to stabilize within a species, as it has among humans, both the production and the reception of messages should be advantageous. If communication were on the whole beneficial to producers of messages (by contributing to their fitness) at the expense of receivers, or beneficial to receivers at the expense of producers, one of the two behaviors would be likely to have been selected out, and the other behavior would have collapsed by the same token (incidentally, there are exceptions, in particular in inter-species communication). In other words, for communication to evolve, it must be a positive-sum game where, in the long run at least, both communicators and receivers stand to gain. For this, it is not necessary

that the interests of the two parties coincide, it is enough that they overlap. The way these interests match and differ influences the manner in which communication evolves and works. Let us look, then, at testimony and argumentation as two communicative practices from the points of view of both communicators and receivers.

TESTIMONY

Unlike argumentation, which is specifically human, testimony (in the sense of “the transmission of observed information”)⁹ is found also in other species. A paradigmatic nonhuman example is the bee dance: one worker bee, having found a source of food, communicates to other worker bees the direction and distance at which it is to be found. At the end of the process, the receiver bees are, presumably, in the same cognitive state with respect to the source of food as they would have been had they found it themselves. This indeed can be described as cognition by proxy. In the human case, testimony does not have quite the same effects as direct perception. When Mary tells John that there is beer in the refrigerator, John is not exactly in the same cognitive state regarding the whereabouts of the beer as he would be had he seen it there himself. To begin with, had John seen the beer, his knowledge of its location would be more detailed and more vivid than when he is just told. More importantly, understanding what one is told involves recognizing a speaker’s meaning, which need not be automatically accepted as true.¹⁰

From the point of view of receivers, communication, and testimony in particular, is beneficial only to the extent that it is a source of genuine (and of course relevant) information. Just as in the case of individual cognition, there may be cases where biases in communicated information are beneficial (think of exaggerated encouragement or warnings, for instance), but these cases are marginal and I will ignore them.

From the point of view of producers of messages, what makes communication, and testimony in particular, beneficial is that it allows them to have desirable effects on the receivers’ attitudes and behavior. By communicating, one can cause others to do what one wants them to do and to take specific attitudes to people, objects, and so on. To achieve these effects the communicator must cause the audience to accept as true messages that in turn will cause the adoption of the intended behaviors or attitudes. Often, these behaviors or attitudes are best brought about by messages that are indeed truthful. In other cases, however, they are best brought about by messages that are not. It is common and often practically (if not morally) appropriate for communicators to achieve the goals they pursue through

communication by misleading or deceiving their audience to some small or large degree. Deception is found in nonhuman animals, but there, just as communication in general, it is quite limited in contents, and highly stereotyped. Humans, thanks to their cognitive abilities and in particular to their metarepresentational capacity to represent the mental states of others, are unique in their ability to engage in creative and elaborate distortion and deception, and also in their ability to question in a reasoned manner the honesty of communicators. Except in marginal cases, it is not in the audience's interest to be deceived or in the communicator's interest to be disbelieved. Dishonest communicators go against the interest of their audience, and distrusting addressees thwart the intentions of communicators.

Looking first at single communication events, we can sketch the possible payoffs of communicators and addressees in game theoretical terms:

		addressee	
		trusting	distrusting
communicator	truthful	<i>gain/gain</i>	<i>loss/no gain</i>
	untruthful	<i>gain/loss</i>	<i>loss/no gain</i>

Communicators can be truthful or untruthful. Addressees can be trusting or distrusting. From the point of view of the communicator, the payoff of communication depends not on her own truthfulness or untruthfulness but solely on the addressee's trust or distrust. The communicator, whatever she chooses to communicate, is better off if the addressee is trusting and worse off (in that her effort is thwarted) if he is distrusting. From the point of view of the addressee, the payoff of communication depends both on his own trust and on the communicator's truthfulness. When he is trusting, the addressee is better off if the communicator is truthful and worst off if she is untruthful. When the addressee is distrusting, he neither gains nor loses from communication (apart from missing a possible gain if the communicator was in fact truthful).

Even though a communication event where the communicator is truthful and the addressee is trusting is advantageous to both, there is no stable solution to the game. The optimal strategy varies for each party with the circumstances. Communicators do not gain just from having any message believed by the addressee. They gain from having the addressee believe a message that brings about effects beneficial to the communicator. Communicators, accordingly, do not choose between being truthful and being untruthful, they choose between expressing and withholding a message,

whether truthful or not, that, if believed by the addressee, should have the desired effects. Addressees, on their side, know that it is not systematically in the interest of the communicator to speak truly, and therefore it is not in their interest to be systematically trusting.

How trusting should addressees be? If, in order never to be misled, addressees were to decide to be systematically distrusting, they would miss all the potential benefit of testimony. After all, it is not at all the case that the communicator's interests are always best served by misinforming the audience. Quite often, true testimony is the best or even the only way to have the intended effect on the addressee, who then stands to gain from accepting the testimony as truthful. On the other hand, if addressees were to decide to be systematically trusting, they would often be deceived (especially by communicators who, having realized that the addressees were systematically trusting, would not be hindered in the use of distortion and deception by the fear of detection). So, it is in the interest of addressees to calibrate their trust as closely as possible to match the trustworthiness of the communicator in the situation. However, there is no failsafe way for addressees to reap all the benefits of communication without incurring the cost of being at times deceived. Still, if communication has stabilized among humans, it must be that there are ways to calibrate one's confidence in communicated information so as to have the expected benefits be greater than the expected costs. More about this later.

So far I have argued that part of the function of communication—the part having to do with the communicators' interest—is optimally fulfilled by the production of messages likely to have certain effects on the audience, irrespective of their truth. It is the causing of desirable effects on the audience that makes communication advantageous to the communicator; without these effects, communication would not have evolved and stabilized. In other words, it is naive to think, in the general case, and in particular in the human case, of the communicator as acting as proxy for the addressee's cognitive needs. The false beliefs spread by communication are not just due to communicators transmitting their own false beliefs, or to their sometimes misusing communication to serve a purpose that goes against its function. Communication produces some amount of misinformation in the performance of its function, more specifically in the performance of those aspects of its function that are beneficial to the communicator.

There is, though, a possible Reidian objection. The game of communication is iterated again and again among the same parties, who moreover alternate in the roles of communicator and addressees. Just as the iteration of prisoner's dilemma games may cause the parties to converge on cooperation,¹¹ the iteration of acts of communication might stabilize a strategy of truthfulness on the part of communicators and of trust on the part of addressees. The evolution of communication is just, so the argument goes,

a particular case of the evolution of cooperation, and reciprocal altruism is possible in that area just as it is in others.

The argument could be fleshed out as follows. In iterated communication, a communicator achieves each time both short-term effects and long-term effects. The short-term effects typically consist in modifying the addressee's beliefs and, indirectly, the addressee's attitude and behavior toward the things communication was about. The long-term effects achieved by an act of communication have to do with the opinion that the audience takes of the communicator's reliability as a source of information, and also, more generally, of her as a person who is helpful or unhelpful, modest or arrogant, considerate or inconsiderate, and so on. The authority, respect, etc., granted to a communicator affect the effectiveness of her future communications.

Whereas it is easy to see that, in many cases, one's short-term goals as a communicator may be best served by departing from the truth, it might seem that one's long-term goals of establishing or maintaining credibility—and therefore one's future short-term goals as a communicator—should always best be served by being truthful. Often indeed, one's overall interest is best served by forsaking the benefits that would come from misleading one's audience, in order to maintain or increase one's ability to influence this audience in the future. Often yes, but not always. There are situations where communicators might better serve not just their short-term but also their long-term goals by being untruthful. Let me mention two considerations in passing. First, credibility is not always best served by truthfulness: some lies are more credible than some truths. Second, credibility is not the only virtue addressees look for in communicators. For instance, it is often advantageous to flatter powerful people, who care as much about loyalty as about credibility, even if it means misleading them. More to the point here is a third consideration, the fact that, in the communication game, long-term effects don't always trump short-term effects. To understand why, we must shift from the communicator's perspective to the audience's.

It is a matter of common observation that people are willing to believe, not everybody, but many other people—relatives, spouses, friends, colleagues, or politicians—even when they know that these have occasionally lied to them. Why shouldn't some version of the tit-for-tat strategy¹² prevail in the iterated game of communication (you lie to me, I cease believing you)? In games of the prisoner's dilemma type, it is always advantageous to defect, provided there is no sanction for defection. By the same token, in iterated prisoner's dilemma games, it is rational to sanction any defection of the other party by a refusal to cooperate for at least several turns. In the communication game, only sometimes is it advantageous to deceive; often, it so happens that the communicator stands to benefit most by sharing true information. Therefore, the fact that a communicator has lied in specific circumstances is no evidence that he or she would in other circumstances.

The communication game is different also from the market situation, where the same goods can be obtained from many different sellers, and where the buyer can and should turn away from dishonest sellers. Each communicator has information that no one else possesses, be it just information about him or herself. So, deciding to systematically ignore the messages of a particular communicator, especially a communicator with whom one stands in an ongoing relationship, can be very costly. The best choice for addressees is to adjust trust not just to each communicator, but to every combination of communicator, situation, and topic. Given this, it is not particularly advantageous for a communicator to adopt a (morally correct) policy of systematic truthfulness. It will generally not be sufficient to cause total trust in the audience, and occasional departures from truthfulness, even if discovered, are unlikely to cause radical distrust.

ARGUMENTATION

The effectiveness of nonhuman animal communication depends on receivers automatically accepting most signals. There may be forms of human unintentional communication where acceptance of a signal is also automatic, as in crowd panic. In the case of human intentional communication, acceptance of a testimony is dependent on trust in the communicator's truthfulness. Testimony, however, is not the only mode of communication of facts, and the effectiveness of human communication is not entirely dependent on the audience's trust. Given the cognitive capacities of humans, and in particular their metarepresentational abilities, human communicators are not limited to testifying to the truth of what they want their audience to accept. They can give reasons as to why the addressee should accept their assertion, and addressees can inspect these reasons and recognize their force, even if they have no confidence at all in the communicator. To take an extreme example, a recognized liar whose testimony would never be accepted on anything could nevertheless convince his audience of a logical or mathematical truth by providing a clear proof of it.

The capacity to reason is in evidence both in individual reflection and in dialogical argumentation. However, it is typically viewed as first and foremost a property of the individual Cartesian thinker. Its function, beside practical reasoning, is seen as that of allowing the individual to go beyond perception-based beliefs and to discover facts with which it happens not to have had perceptual acquaintance, and, more importantly, theoretical facts with which there is no way to be perceptually acquainted. On this view, reasoning is a higher-level form of individual cognition, a superior tool for the pursuit of knowledge.

From an evolutionary psychology perspective, there is something implausible in this view of reasoning. The expectation is that there should be domain- and task-specific inferential mechanisms corresponding to problems and opportunities met in the environment in which a species has evolved. But, or so the argument goes, it is unclear that there would have been much pressure for the evolution of a general reasoning ability that would presumably be slow and costly and would perform less well than specialized mechanisms in their domains. At best such an ability would handle—and not too well—issues and data the processing of which had not been preempted by more effective specialized devices. Some evolutionary psychologists have concluded that there is no general “logical” ability in the human psychological make-up. I have argued differently that there are evolutionary reasons to expect a kind of seemingly general reasoning mechanism in humans, but one that is, in fact, specialized for processing communicated or to-be-communicated information.¹³ The function of this mechanism is linked to communication rather than to individual cognition. It is to help audiences decide what messages to accept, and to help communicators produce messages that will be accepted. It is an evaluation and persuasion mechanism, not, or at least not directly, a knowledge production mechanism.

As I noted earlier, for communication to have stabilized among humans, audiences must have developed ways of calibrating their confidence in incoming information effectively enough for the benefits of communication to remain well above the costs. Actually, the potential benefits are so important, and the risks of deception so serious, that all available means of calibration of trust may well have evolved. Three such means come to mind. One is paying attention to behavioral signs of sincerity or insincerity (but these can be faked to some extent).¹⁴ A second, more important means consists in trusting more or less as a function of the known degree of benevolence of the communicator (thus trusting relatives more than strangers, friends more than enemies, and so on—this may seem obvious, but note that it is very different from the Reidian idea—that Reid himself went on to qualify—of a “disposition to confide in the veracity of others, and to believe what they tell us”). There is a third means, which is to pay attention both to the internal coherence of the message and to its external coherence with what is already believed otherwise. It seems plausible that all three of these means have indeed evolved among humans. I want to focus on the third, coherence checking. (I will be using “coherence” and “incoherence” to refer to logical relationships and to evidential relationships of support and undermining.)

A problem well known to anybody who has ever tried to lie and to stick to one’s lie over time, is that it is increasingly hard to keep it coherent with what is otherwise known to the audience without embellishing it, and

increasingly difficult to embellish it without introducing internal inconsistencies. A sincere but false claim is also likely to encounter coherence problems. A useful method, then, to detect misinformation and in particular deception is to check for the internal and external coherence of messages.

Coherence checking should be useful for detecting all false beliefs, whether derived from communication or from individual cognition, so why, if it exists at all, shouldn't it have evolved primarily as a tool of individual cognition? Here is the answer. Coherence checking involves a high processing cost, it cannot be done on a large scale for it would lead to a computational explosion, and is itself fallible. Individual mechanisms of perception and inference, though not perfect, are probably reliable enough to make, on balance, checking the coherence of their outputs superfluous or even disadvantageous. We would be surprised to discover that coherence checking occurs in a nonhuman species (and if we did, we should look for peculiarity of the informational environment of the species that would make the procedure beneficial).

My first suggestion is this: coherence checking—which involves metarepresentational attention to logical and evidential relationships between representations—evolved as a means of reaping the benefits of communication while limiting its costs. It originated as a defense against the risks of deception. This, however, was just the first step in an evolutionary arms race between communicators and audiences (who are of course the same people, but playing—and relying more or less on—two different roles).

The next move in the evaluation-persuasion arms race was from the communicators' side and consisted in displaying the very coherence the audience might look for before accepting the message, a kind of "honest display" with many well-known analogs in animal interaction. Testimony can be given by a mere concatenation of descriptive sentences. Displaying coherence requires an argumentative form, the use of logical terms such as "if," "and," "or," and "unless," and of words indicating inferential relationships such as "therefore," "since," "but," and "nevertheless." It is generally taken for granted that the logical and inferential vocabulary is—and presumably emerged—as tools for reflection and reasoning. From an evolutionary point of view, this is not particularly plausible. The hypothesis that they emerged as tools of persuasion may be easier to defend.

The next steps in the arms race are for the audience to develop skills at examining arguments of the communicator, and for communicators at improving their argumentative skills. Thus emerges an argumentation mechanism of rhetorical construction and epistemic evaluation of messages. This mechanism processes representations that have been, or that are to be, communicated. These are a very special kind of object in the world. Moreover, the mechanism pays attention only to some properties of these representations, their logical and evidential relationships. In other words this metarep-

representational device is highly domain- and task-specific. On the other hand, the communicative representations it processes can themselves be about anything. So, by causing these representations to be accepted, the argumentation mechanism contributes to the production of beliefs in all domains, and possesses in this sense a kind of virtual domain-generalness.

Whether or not you accept the evolutionary hypothesis I just suggested, the fact remains that argumentation has a different function for communicator and audience. To communicators, it is a means of persuasion; to audiences, it is a means to critically evaluate a message before accepting it. It might seem that, by displaying the coherence of their message, communicators handicap their ability to deceive their audience. This is true—honest argumentation is harder to fake than honest testimony—but only to a certain extent. As already noted, coherence checking cannot be thorough (especially when it carried out at the speed of speech at the same time as the already effort-demanding comprehension process), and it is fallible. Real arguments are at best enthymematic, and often just allude to the existence and structure of a complete demonstration. Effective argumentation, from the point of view of the persuader, is argumentation that can sustain the degree of checking that the audience is likely to submit it to. From the point of view of the audience, cost and benefit considerations are in play: the cost of risking being misled has to be weighed against both the cost of refusing some genuine and relevant information (as in the case of testimony) and the processing cost of checking the coherence of the argument. This latter cost can be modulated by checking more or less thoroughly, and clearly it would not be advantageous to check every argument quite thoroughly. This leaves room for a deceptive use of argumentation, as is well known since the Sophists.

From a logical and epistemological normative point of view, sophistry is a perversion of argumentation, a practice that goes against its very *raison d'être*. From the evolutionary perspective that I have just sketched, sophistry is a way to use the “honest display” strategy of argumentation in a dishonest way and thereby make it more advantageous for the communicator. In other words, sophistry contributes to making argumentation adaptive.

CONCLUSION

Wherever you have a function, you can normatively evaluate the degree to which the function is effectively fulfilled. In this limited sense, function entails norm. Since the function of communication presents itself differently for communicator and audience, one can evaluate to what extent a communicative practice allows communicators to achieve intended effects on the audience, and to what extent it provides to the audience genuine and relevant

information. One can work out a global evaluation of the extent to which a communicative practice provides to both parties sufficient overall benefits so as to cause them to perpetuate the practice.

It might seem then that to approach communicative practices from a veristic epistemological point of view, as does Goldman, is to espouse the point of view of the audience, to impose on it a norm that may be perfectly justifiable on ethical or pragmatic grounds, but that is not the norm that effectively governs the communicative process. However, the true situation is interestingly more complex. Communicators present themselves as honest, whether or not they are, and whether or not it is in their interest to be. Without presenting themselves as truthful, liars could not even begin to lie. Without presenting themselves as rational arguers, sophists could not even begin to persuade. In other terms, the point of view of the audience determines a norm that is implicit in and intrinsic to all communication, a norm of truthfulness in testimony and of rationality in argumentation. These norms are not imposed from without, or from just one point of view in the communication process. They are overtly accepted by both parties to the process.

The point of my paper could be rephrased by saying that the epistemic norms implicit in the process of communication¹⁵ are to a limited but interesting extent at odds with the very function of communication. Therefore the prospect for knowledge discussed by Goldman is not just brightened—immensely so—it is also significantly threatened by well-functioning communicative practices.

NOTES

1. Alvin I. Goldman, *Knowledge in a Social World* (Oxford: Clarendon Press, 1999), vii.
2. For elaborations and variations of this notion of function, see C. Allen, M. Bekoff, and G. Lauder (eds.), *Nature's Purposes: Analyses of Function and Design in Biology* (Cambridge, Mass.: MIT Press, 1998).
3. See Stephen Stich, *The Fragmentation of Reason* (Cambridge, Mass.: MIT Press, 1990).
4. In fact, it is not obvious that it could be performed in any other way; see L. Cosmides and J. Tooby, "Origins of Domain Specificity: The Evolution of Functional Organization," in L. A. Hirschfeld and S. A. Gelman (eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture* (New York: Cambridge University Press, 1994), 85–116; and Dan Sperber, "The Modularity of Thought and the Epidemiology of Representations," in L. A. Hirschfeld and S. A. Gelman (eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture* (New York: Cambridge University Press, 1994), 39–67.
5. Goldman, *Knowledge in a Social World*, 106, 129.
6. Thomas Reid, *An Inquiry into the Human Mind*, ed. T. Duggan (Chicago: University of Chicago Press, 1970), 238–40.
7. R. Dawkins and J. R. Krebs, "Animal Signals: Information or Manipulation?" in J. R. Krebs and N. B. Davies (eds.), *Behavioural Ecology* (Oxford: Basil Blackwell Scientific Publications, 1978), 282–309.

8. J. R. Krebs and R. Dawkins, "Animal Signals: Mind-reading and Manipulation," in J. R. Krebs and N. B. Davies (eds.), *Behavioural Ecology* (Sunderland, Mass.: Sinauer Associates, 1984), 380–402.
9. Goldman, *Knowledge in a Social World*, 103.
10. Millikan disagrees and defends the view that human communication also is cognition by proxy; Ruth Millikan, *Language, Thought and Other Biological Categories* (Cambridge, Mass.: MIT Press, 1984). For a discussion, see G. Origgi and D. Sperber, "Evolution, Communication, and the Proper Function of Language," in P. Carruthers and A. Chamberlain (eds.), *Evolution and the Human Mind: Language, Modularity and Social Cognition* (Cambridge: Cambridge University Press, 2000)
11. See R. Axelrod, *The Evolution of Cooperation* (New York: Basic Books, 1984); and P. Kitcher, "The Evolution of Human Altruism," *Journal of Philosophy* 90 (1993): 497–516.
12. Axelrod, *The Evolution of Cooperation*.
13. Dan Sperber, "Metarepresentations in an Evolutionary Perspective," in Dan Sperber (ed.), *Metarepresentations: A Multidisciplinary Perspective* (New York: Oxford University Press, 2000), 117–37.
14. P. Ekman, *Telling Lies* (New York: Norton, 1985).
15. D. Wilson and Dan Sperber, "Truthfulness and Relevance," *Mind* (in press).