

SPEAKERS ARE HONEST BECAUSE HEARERS ARE VIGILANT REPLY TO KOURKEN MICHAELIAN¹

DAN SPERBER
dan.sperber@gmail.com

ABSTRACT

In ‘The Evolution of Testimony: Receiver Vigilance, Speaker Honesty and the Reliability of Communication,’ Kourken Michaelian questions the basic tenets of our article ‘Epistemic Vigilance’ (Sperber et al. 2010). Here I defend against Michaelian’s criticisms the view that epistemic vigilance plays a major role in explaining the evolutionary stability of communication and that the honesty of speakers and the reliability of their testimony are, to a large extent, an effect of hearers’ vigilance.

Human communication is without equivalent in the animal kingdom both because of the wealth of contents it conveys and because of its pervasive role in individual cognitive development, in social interaction, and in cultural transmission. This makes understanding human communication from an evolutionary point of view quite challenging.

First, instead of a small repertoire of simple genetically transmitted signals with rudimentary meanings adapted to quite definite situations, humans make use of acquired languages that can generate sentences of unlimited syntactic and semantic complexity. Moreover, the interpretation of linguistic utterances goes well beyond mere decoding and involves using rich contextual evidence to attribute a communicative intention to speakers. Deirdre Wilson and I have developed an account of the inferential basis of human communication, relevance theory, that I believe helps address the challenge to evolutionary thinking presented by this extraordinary semantic richness of human communication (Sperber and Wilson 1995; Wilson and Sperber 2012).

Secondly, human communication takes place not just among closely related individuals and not just on a few issues of common interest but among strangers as well as kin, and on all topics on which humans are capable of having thoughts. Just for contrast, try imagining a non-human animal walking up to a non-related conspecific and asking for directions. This omnipresence and multipotency of communication in human life raises a variety of important questions. In particular this: Communication involves an interaction between two parties, a sender and a receiver, whose interests may diverge. Generally speaking, senders stand to gain from communication by influencing receivers, that is, by modifying their beliefs and preferences; receivers stand to gain from communication by acquiring true and relevant information from senders. Depending on their specific goal, senders can best exercise the influence they want by sometimes truly informing receivers and sometimes misinforming them. Given the massive dependence of humans on communication from others,

1 I am grateful to Fabrice Clément, Christophe Heintz, Olivier Mascaro, Hugo Mercier, Gloria Origgi, and Deirdre Wilson for their suggestions and comments.

the risk for receivers of being misinformed is quite serious. It cannot but compromise to some degree the overall benefits receivers get from communication.

The probability of a biological trait evolving is contingent on its costs-benefits balance. Only if the benefits are greater than the costs, is it likely to evolve at all, and it is likely to evolve in a manner that, within the local range of possibilities, optimizes this balance. For communication to evolve, the cost-benefit balance must be favourable both for senders and receivers. Given the difference of interests between senders and receivers, the dispositions involved in each are likely to have distinct evolutionary trajectories, with the potential of an evolutionary arms race between senders' and receivers' dispositions, senders becoming better at influencing receivers and receivers becoming better at discriminating information from misinformation.

What I have said so far should hardly be controversial. It just amounts to listing considerations that are potentially relevant to the study of the evolution of communication. Whether they are actually relevant is another matter. There may be interfering factors that make these considerations not so relevant after all. Also, it might be impossible, in the current state of our understanding, to go beyond these general considerations and to speculate usefully – that is, with serious arguments and evidence—on what mechanisms have actually evolved.

In *Epistemic Vigilance* (Sperber et al. 2010), we tried to show that useful speculation is possible by engaging in it. We formulated a conjecture based on the considerations just evoked. Humans, we suggested, 'have a suite of cognitive mechanisms for epistemic vigilance, targeted at the risk of being misinformed by others'. A suite rather than a single mechanism because, while there is no failsafe algorithm to avoid being misinformed, there are several different and complementary possible heuristics that may each contribute in a cost-effective way to approximating this goal. Given the stakes (the benefits and risks of communication for receivers), we assume that those heuristics that could evolve and be cost effective probably did. I have argued that the role of selective pressures in cultural evolution is more complex (Sperber 1996, Sperber and Claidière 2006), but there too, I would expect efficient social mechanisms of epistemic vigilance to have evolved and to still be evolving today (e.g. Heintz 2006; Tennie et al. 2010).

To give plausibility and substance to this conjecture, we surveyed a variety of issues, research and theories in different domains of philosophy, linguistics, cognitive psychology and the social sciences. Our further goals in doing this were (1) to highlight the relevance to the study of epistemic vigilance of various research paradigms that have been pursued in isolation from one another; (2) to encourage new research in this area; and (3) to start an interdisciplinary conversation on the topic. It is gratifying to be able to contribute to this conversation in *Episteme* by responding to Kourken Michaelian's insightful discussion of our article. Michaelian discusses relevant evidence and makes interesting suggestions, but, let me say at the outset, I do not see this evidence as weighing against our conjecture, nor these suggestions as outlining a better supported alternative.

IS EPISTEMIC VIGILANCE INDISPENSABLE?

Part of Michaelian's strategy is to propose an array of stronger or weaker possible interpretations of our claims and to argue at best for the weakest. He starts by suggesting that our evolutionary conjecture can be given two readings with importantly different

implications: according to a stronger reading, epistemic vigilance is indispensable for communication to endure among humans. According to a weaker reading, vigilance is dispensable but still adaptive. We did write, ‘It is because of the risk of deception that epistemic vigilance may be not merely advantageous but indispensable if communication itself is to remain advantageous’, which suggests that we recognize these two possible readings without committing to either. In the light of Michaelian’s discussion, I now believe that we could and should have been even less definite. Given the variety of forms of communication and of types of partners involved, neither reading would be plausible if taken to apply to human communication as a whole. Communication among people, in particular close kin, on matters where their interests are aligned does not require epistemic vigilance towards the risk of deception in order to be advantageous to all participants. On the other hand, several typical forms of human communication such as bargaining or arbitration could hardly be practiced in a manner advantageous to all in the long run without vigilance.

Michaelian’s ‘two readings’ involve treating human communication as a static given and asking whether epistemic vigilance is indispensable for it as a whole. From an evolutionary point of view, it would make better sense to pose the question in a more fine-grained and dynamic manner. In the six millions of years or so since the separation of the ancestors of modern humans from the ancestors of their most closely related great apes, chimpanzees and bonobos, the semantic and social scope of human communication has hugely broadened, with presumably an ever increasing acceleration of this broadening in the past 40,000 years. Together with the huge expansion of human cognitive capacities and of human cooperation — the three developments being obviously related — this is the most striking aspect of hominin evolution. In other animals, we expect the evolution of communication to have been severely limited by, among other factors, differences of interests between senders and receivers (Dawkins and Krebs 1978; Krebs and Dawkins 1984). Why not in humans? Our suggestion is that humans evolved cognitive capacities and cultural institutions that (1) made it possible to communicate much richer contents, and (2) also made possible a degree of epistemic vigilance that reduced the risks of misinformation and made communication advantageous on a uniquely large scale. The benefits of vigilance may be negligible in some communicative interactions and essential in other interactions. All I feel confident to say is that, without vigilance, human communication would be a very different and probably much more restricted affair. Anthropological studies in societies where forms and levels of vigilance differ would certainly throw some light on the issue (see for instance Gilsenan 1976).

IS EPISTEMIC VIGILANCE EFFICIENT?

With the same strategy of suggesting stronger and weaker interpretations of our claims, Michaelian writes:

When Sperber et al. argue that humans are epistemically vigilant, they might, at the most general level, be making either of two claims:

1. *Bare vigilance*: Recipients are vigilant in the sense that they monitor for (are on the lookout for) deception on the part of communicators, whether or not they often succeed in detecting it when it is present.

2. *Effective vigilance*: Recipients are vigilant in the sense that they monitor for deception on the part of communicators (they exercise bare vigilance), and this monitoring is effective, i.e., they often succeed in detecting deception when it is present.

Needless to say, we conjecture that the stronger claim is true: people are generally vigilant towards the risk of deception (in the form of a low-key monitoring, not of active suspicion) and they are good enough at detecting it to make it a cost-effective investment of mental energy. If people were uselessly vigilant, this vigilance — or should we then call it ‘paranoid attitude’? — far from explaining anything, would be another and quite puzzling thing to explain.

That epistemic vigilance is a permanent dimension of our communicative exchange is not received wisdom. Many philosophers believe that people are generally not just trustful, but trustful as a matter of course, without any permanent even low-level attention to the risk of misinformation. Thomas Reid famously argued that we have been endowed by God with a disposition to speak the truth and a disposition to accept what other people tell us as true. Modern philosophers such as Burge, Davidson, Lewis or Millikan have defended similar views not because of any specific empirical evidence or argument, but because basic truthfulness and trust play a central role in their philosophy of language (for a detailed discussion of Millikan’s views in this respect see Origg and Sperber 2000). The psychologist Daniel Gilbert has produced experimental evidence in favour of such a Reidian (or, as he calls it, Spinozan) approach (e.g. Gilbert et al. 1993). In our article we discussed these philosophical arguments and Gilbert’s experimental evidence and argued that they were not particularly compelling. We surveyed various sources of empirical evidence suggesting that people, including children, are as a matter of course paying a modicum of attention to the trustworthiness of communicators and to the plausibility of what they assert. I will readily grant that the import of this evidence is open to discussion. Still, it is strong enough to justify doing further empirical research on the topic rather than assuming that the facts of the matter can be decided on a priori or common sense grounds, or on the basis of just a few clever experiments.

Considering both the degree to which vigilance might be practiced and its degree of efficiency, Michaelian distinguishes six possibilities. In three of them, vigilance plays no role, either because people are not vigilant, or because vigilance is inefficient. He focuses on the three possibilities where vigilance is exercised and is not wholly without efficacy:

Strongly effective vigilance: Recipients exercise bare vigilance, and, due primarily to this, they usually avoid being deceived.

Moderately effective vigilance: Recipients exercise bare vigilance, and they usually avoid being deceived; both vigilance and some other factor make significant contributions.

Weakly effective vigilance: Recipients exercise bare vigilance, and they usually avoid being deceived, but this is due primarily to some other factor.

Michaelian concludes in favour of weakly effective vigilance. For my part, on theoretical grounds I would tend to think that vigilance is likely to be moderately effective, but I don’t believe that we have, or even are close to having, clear empirical evidence to help us decide the issue. Michaelian is more confident, but this is because he redefines ‘epistemic vigilance’ in a narrow way that does not correspond to what we were talking about, and that, more importantly, is not particularly useful for addressing the puzzle presented by the unique scope of human communication.

Michaelian seems to attribute to us the view that ‘epistemic vigilance is a matter of processes devoted to screening out incoming false information on the basis of available behavioural cues’. Showing that vigilance in this narrow sense is not efficient would, he holds, be quite damaging to our conjecture. This is a misunderstanding.

Here is how this misunderstanding might have arisen. To suggest that a concern for reliability of communicated information is a permanent feature of human interaction, we cited experiments (Willis and Todorov 2006; see also Origgi 2012) suggesting that looking for signs of trustworthiness is the first thing people do when seeing a new face, even when they are presented with the picture of a face for a mere 100 milliseconds. We did not assume that people were correct in forming judgments on that basis. We mentioned this research as evidence of people’s concern, not to say obsession, with trustworthiness. In the same spirit, we talked briefly about the rich lie detection literature that suggests (even if there are problems of ecological validity with the experiments) that people are much less good at detecting lies from behavioural cues than they think they are, and are in fact close to random in their judgments in the matter.

Michaelian goes at length over a wider sample of this same literature defending its negative conclusion (which we do not dispute) as if it was a source of objection against the idea that epistemic vigilance is pervasive and for the most part efficient. Nowhere did we argue or imply that detection of lies through behavioural cues was important to the success of vigilance. Indeed, if people could not help but provide their audience with behavioural evidence of their bad faith when they lie, deception would be nearly impossible and the whole issue we are discussing would hardly arise. But if, as we suggested, there has been an evolutionary arms race between our abilities as senders and our abilities as receivers, being able to lie without behaviourally betraying that one is doing so should be an early evolving adaptation for the sender in that arms race, pushing the receiver to look for less easily suppressed evidence. So, I agree with Michaelian that there is no evidence that this kind of vigilance would be effective. I would surmise that very little energy is invested in this form of vigilance in ordinary conditions. Of course, if someone we have independent reasons to suspect of lying starts acting nervous, we notice and interpret this as possible evidence of bad faith. If someone talking to us starts stuttering and sweating profusely, we look for an explanation and moral unease caused by lying is one distinct possibility. On the other hand, we do not monitor for such kinds of evidence every time someone is talking to us. Indeed, it would not be cost-effective at all.

When we argued that there is a *suite* of abilities contributing to epistemic vigilance, we did not see lie detection on the basis of behavioural cues as a particularly important item in that suite. In any case, remember why we introduced the notion of epistemic vigilance: to explain how receivers can limit the risk of being misinformed and in particular deceived. Our point was that there had been selective pressure for *any* cost-effective relevant ability. Abilities likely to be involved include drawing inferences from remembered past communicative performance, understanding the sender’s immediate purpose in communicating, or taking into account the sender’s reputation. Michaelian seems to think these could not be bona fide items in the epistemic vigilance suite. Here is an example of his reservations. We evoked a series of experiments by means of which Olivier Mascaró investigated the development of epistemic vigilance among pre-schoolers (see Mascaró and Sperber 2009). In one of the experiments, children were told that an informant (a puppet) was mean and did not want them to find a sweet. Five-year-olds—but not

four-year-olds – inferred that the informant was lying about the location of the sweet (see also Vanderbilt et al. 2011). This, Michaelian comments, ‘tells us nothing about children’s ability to detect deception where they do not already have reason to take it to be present’. Right, if by ‘detection’ we understand ‘detection based on behavioural cues’, but what these results do tell us is that five-year-olds pay attention to such reasons and draw appropriate conclusions from them. This is a perfect instance of epistemic vigilance as we defined it (not arbitrarily but on theoretical grounds).

There are obvious costs in remembering someone’s past communications, evaluating them in the light of later confirming or confuting evidence, tracking the reputation of people who might want to inform or misinform us, interpreting not just what a person says, but her motivations in saying it. Paying these costs makes sense if they help us evaluate the reliability of present and future communication (and sometimes in revising our evaluation of past communications that still influence us). Michaelian himself writes:

Though dishonesty is rarely detected at the time of the utterance, it can often be detected after the fact, by means including physical evidence and third-party testimony. (The capacity for such after-the-fact deception detection does not amount to a form of vigilance, in the relevant sense, but only to a normal sensitivity to evidence: as we saw in section 2, what is at issue here is our capacity for real-time detection of deception.)

Again, I don’t see why only what he calls ‘real-time detection’ should be involved in a sensible, theoretically useful notion of epistemic vigilance. Finding out after the fact that one has been deceived may limit the damages of this deception and, going beyond mere updating of one’s beliefs, prevent future deception, that is, contribute to fulfilling the function of epistemic vigilance.

Also, what Michaelian calls ‘normal sensitivity to evidence’ needs unpacking. Every state of affairs, every event is bursting with evidence that we do not pay attention to because of its irrelevance to us. We are ‘normally sensitive’ to potentially relevant evidence, and not to evidence in general. Our sensitivity to evidence potentially relevant to the epistemic evaluation of communicated information cannot be explained by appealing to some unspecified notion of ‘normal sensitivity’. It can be explained by appealing to the hypothesis that we have an evolved disposition to be epistemically vigilant.

In our article we distinguished vigilance towards incompetence (which has been richly studied, in particular from a developmental point of view, see e.g. Clément et al. 2004; Koenig and Harris 2007) and vigilance towards deception and argued that only the latter posed a major evolutionary problem. While the selective pressure for the two kinds of vigilance may differ and be stronger for vigilance towards deception, the psychological mechanism involved may be in part the same in both. In particular, we distinguished vigilance towards the source of information – who to believe? – and vigilance towards the content communicated – what to believe? Vigilance towards the content attends to the coherence of information communicated with background information, and to its internal coherence (Mercier and Sperber 2011). We argue that some evidence on coherence is available as a by-product of a relevance-guided comprehension procedure, so that, up to a point, vigilance towards the content can be a low-cost, low-attention affair. This vigilance towards the content is equally relevant to the detection of deception as it is to the detection of incompetence. As any liar knows, the coherence that normally goes with speaking truly takes some effort to achieve and maintain in deception. Michaelian

however ignores this vigilance towards the content, as if it was irrelevant to our conjecture, but it is not.

EXPLAINING HUMAN COMMUNICATION

Michaelian concludes his discussion of epistemic vigilance (understood as real-time detection on the basis of behavioural cues) with the claim that humans have a ‘slight sensitivity to dishonesty [that] stems from the occasional transparency of liars’.

[T]here remains, [he writes] the question of how to account for the reliability of communication given that we are largely insensitive to deception. Sperber et al. focus on the potential role of vigilance in ensuring the reliability of communication. But if the base rate of dishonesty is sufficiently low, then vigilance need not be invoked to explain the reliability of communication.

So, what we took to be a serious evolutionary problem seems not to arise.

What does it mean to say that ‘the base rate of deception is low’? There are a few problems with such a statement. Michaelian cites a number of interesting studies but what they document is the rate of daily deception in given (contemporary, Western) populations as indicated by self-reports, or the rate of deception in controlled experiments, with obvious problems of extrapolation or generalisation. In what sense are these rates ‘base rates’? The expression would make sense if we could measure the frequency of deception in a communication system where receivers were gullible. If that base rate were sufficiently low, epistemic vigilance might indeed be superfluous. But there is no way to obtain such a measure. Otherwise what is aimed at is measuring the rate of deception in the world as it is, and since it is certainly not proven that there is no epistemic vigilance at work, a low rate might be to a large extent an effect of this vigilance.

Moreover, available data does not even vaguely indicate the aggregated expected utility of deception for receivers. Vigilance would be much more important if people lied once in a while on some momentous matter than if they lied ten times a day on irrelevant details. Mere frequency of deception without a sense of its potential cost for receivers is not really relevant to the issue.

Still, do the studies cited by Michaelian at least suggest that there is little need for receivers to be vigilant towards the risk of deception? Well, not really. Here is an example. Levine et al. (2010), cited by Michaelian, report three experiments investigating whether people would lie with or without a motive to do so. The authors write:

As predicted, when honesty was sufficient to meet situational demands, honest messages were selected, generated, and observed 98.5% to 100% of the time. Alternatively, deception was observed 60.0% to 64.3% of the time when variations in the same situations made the truth problematic. It is concluded that people usually deceive for a reason, that motives producing deception are usually the same that guide honesty, and that people usually do not lie when goals are attainable through honest means. (Levine et al. 2010: 271)

Even without extrapolating from this experiment, I don’t see how a 60% rate of deception when it is in the interest of the communicator to deceive can be construed as evidence that epistemic vigilance is superfluous.

Michaelian agrees that available evidence on the base rate of deception ‘by itself is too limited to be conclusive’. He offers however another argument, a deduction actually, for assuming that this rate must be low:

[G]iven that we are largely insensitive to deception, the base rate must be low: otherwise, communication would collapse, but communication appears to be evolutionarily stable. In short, we can infer that the base rate of deception is low from the fact that receivers are largely insensitive to deception.

I agree with the premise that we are largely insensitive to behavioural cues of deception, but not with the premise that therefore unless the base rate were low, communication would collapse. What is simply ignored in this argument is the possibility that other mechanisms for epistemic vigilance stabilise communication. By ignoring the possibility of epistemic vigilance and the effect it would have on the rate of deception, you ‘prove’ that the base rate of deception is low enough to make vigilance quasi superfluous.

To explain this alleged low rate of deception, Michaelian offers some considerations that have relevant implications for our disagreement. He writes:

More importantly, lying entails *cognitive costs*. [...]:

- Formulating a lie may be more cognitively demanding, since coherence cannot be taken for granted
- Liars generally devote more resources to monitoring and controlling their own behaviour, since they are less likely to take their credibility for granted
- Liars generally devote more resources to monitoring receivers, for the same reason Liars may have to remind themselves to role-play

I agree – and we mentioned the point ourselves – that lying involves cognitive costs that may make it preferable to tell the truth rather than lie in the absence of reasons to do so (see, however, Shalvi et al. 2012). But the second and third costs invoked in the above quotation make sense only if receivers exercise epistemic vigilance. Why otherwise should liars worry about their credibility? Why should they monitor the receiver and remind themselves to role-play? All this extra effort is clearly useful when addressing a vigilant audience. Not so if receivers are gullible.

Michaelian also evokes emotional cost:

Lying will, in normally socialized agents, have an *emotional cost*, since they have internalized norms that forbid lying, except under special circumstances. Violation of such internalized norms constitutes a disincentive to lying. Though they do not play the major role, I note that emotional costs based on violation of social norms against lying can be invoked in the sort of evolutionary explanation developed here, since there is no requirement that, in such an explanation, the low base rate of deception be explained entirely in terms of built-in factors internal to communicators – socially imposed deterrents can play a role, and internal factors derived from social pressures can also play a role.

I don’t disagree with this conjecture (even though it should be noted that, unlike energy costs, emotional costs are not necessarily detrimental to fitness and that therefore invoking them in an evolutionary explanation is not self-evident). However, wouldn’t the ‘socially

imposed deterrents' that Michaelian invokes require widespread epistemic vigilance to have any efficacy? The phenomena that he sees as making epistemic vigilance superfluous are best explained, I would claim, by the prevalence of epistemic vigilance.

From an evolutionary point of view, the fact that humans can engage in communication the way they do, expect others to be rather honest and indeed are rather honest themselves begs for an explanation. Our explanation is that communicators are generally honest *because* they are addressing a vigilant audience. Quasi-universal vigilance makes dishonesty less likely to be beneficial in the short run and more likely to be costly in the long run: falsehoods may be disbelieved, and dishonesty may have reputational costs.

There is however an alternative possible explanation that deserves discussion. It could be that the deterrence Michaelian invokes is just an application of whatever practices contribute to the imposition of morality in general rather than being the effect of dedicated epistemic vigilance mechanisms. More generally, one way in which our conjecture about epistemic vigilance might be false or at least off the mark would be if the evolution of honesty in communication was just an automatic consequence of the evolution of cooperation and morality.

The evolution of human cooperation and the related evolution of human morality have no doubt played a major role in the extraordinary broadening of human communication. It might even be thought that the difference of interests between senders and receivers is prevented from compromising the stability of communication by the evolution of moral dispositions. These dispositions would make senders quite generally honest; the risk of misinformation would be so low as to make it more cost-effective for receivers to be quite generally trustful. So might the argument go. Note however that theoreticians of the evolution of human cooperation and morality all assume that some form of what could be called 'moral vigilance' is essential. Those who favour group-selection accounts of the evolution of morality typically give a central role to punishment (e.g. Gintis et al. 2003); those who favour a mutualistic approach give a central role to reputation (e.g., Baumard et al. 2013; Sperber and Baumard, 2012). In other terms, no one assumes that people's own individual moral sense is sufficient to cause them always to behave morally. What is generally assumed on the contrary is that the moral vigilance of others plays an important role in everybody's relatively moral behaviour.

So, could honesty in communication result from, or be an aspect of, a more general moral sense bolstered by moral vigilance? Could what we have described as epistemic vigilance be just an effect of moral vigilance? This is conceivable. It might be investigated by looking for instance for dissociations between epistemic and other forms of moral vigilance both in development and in attention and inference. If some such dissociation were found, this would suggest that epistemic vigilance is *sui generis*; if none were found, it would suggest that epistemic vigilance is just moral vigilance applied to communication.

According to the approach we suggested however, there is not one single mechanism of epistemic vigilance; there are several mechanisms turned towards either the trustworthiness of the source or the believability of the content. For similar reasons, moral vigilance might also better be viewed as a suite of mechanisms attuned to different types of moral behaviour and different types of evidence of morality, rather than as a single evaluator of the moral conformity of people's behaviour. I would expect all these more specific mechanisms to have distinct evolutionary and/or developmental trajectories, and multiple kinds

of dissociation to occur. Again, these are conjectures that I believe are sensible and worth exploring, not strong claims.

In this speculative context, I have found Michaelian's discussion and arguments, if not compelling, at least stimulating and useful, and I am grateful for them.

REFERENCES

- Baumard, Nicolas, André, Jean-Baptiste, and Sperber, Dan. 2013. 'A Mutualistic Approach to Morality.' *Behavioral and Brain Sciences*, 36: 59–122.
- Clément, Fabrice, Koenig, Melissa, and Harris, Paul. 2004. 'The Ontogeny of Trust.' *Mind and Language*, 19: 360–79.
- Dawkins, Richard, and Krebs, John. 1978. 'Animal Signals: Information or Manipulation.' In John Krebs and N. B. Davies (eds), *Behavioural Ecology*, pp. 282–309, Oxford: Blackwell.
- Gilbert, Daniel, Tafarodi, Romin, and Malone, Patrick. 1993. 'You Can't Not Believe Everything you Read.' *Journal of Personality and Social Psychology*, 65: 221–33.
- Gilsenan, Michael. 1976. 'Lying, Honor and Contradiction.' In B. Kapferer (ed.), *Transaction and Meaning*, pp. 191–219. Philadelphia: Institute for the Study of Human Issues.
- Gintis, Herbert, Bowles, Samuel, Boyd, Rob, and Fehr, Ernst. 2003. 'Explaining Altruistic Behavior in Humans.' *Evolution and Human Behavior*, 24(3): 153–72.
- Heintz, Christophe. 2006. 'Web Search Engines and Distributed Assessment Systems.' *Pragmatics and Cognition*, 14(2): 387–409.
- Koenig, Melissa, and Harris, Paul. 2007. 'The Basis of Epistemic Trust: Reliable Testimony or Reliable Sources?' *Episteme*, 4: 264–84.
- Krebs, John, and Dawkins, Richard. 1984. 'Animal Signals: Mind-Reading and Manipulation.' In John Krebs and N. B. Davies (eds), *Behavioural Ecology: An Evolutionary Approach*, pp. 380–402. Oxford: Blackwell.
- Levine, Timothy, Kim, Rachel, and Hamel, Lauren. 2010. 'People Lie for a Reason: Three Experiments Documenting the Principle of Veracity.' *Communication Research Reports* 27(4): 271–85.
- Mascaro, Olivier, and Sperber, Dan. 2009. 'The Moral, Epistemic, and Mindreading Components of Children's Vigilance Towards Deception.' *Cognition* 112(3): 367–80.
- Mercier, Hugo, and Sperber, Dan. 2011. 'Why do Humans Reason? Arguments for an Argumentative Theory.' *Behavioral and Brain Sciences*, 34: 57–11.
- Origgi, Gloria. 2012. 'Epistemic Injustice and Epistemic Trust.' *Social Epistemology*. 26(2): 221–35.
- , and Sperber, Dan. 2000. 'Evolution, Communication and the Proper Function of Language.' In Peter Carruthers and Andrew Chamberlain (eds), *Evolution and the Human Mind: Modularity, Language and Meta-Cognition*. Cambridge: Cambridge University Press.
- Shalvi, S., Eldar, O., and Bereby-Meyer, Y. 2012. 'Honesty Requires Time (and Lack of Justifications).' *Psychological Science*. 23(10): 1264–70.
- Sperber, Dan. 1996. *Explaining Culture: A Naturalistic Approach*. Oxford: Blackwell.
- , and Baumard, Nicolas. 2012. 'Moral Reputation: An Evolutionary and Cognitive Approach.' *Mind and Language*, 27(5): 485–518.
- , and Claidière, Nicolas. 2006. 'Why Modeling Cultural Evolution is Still Such a Challenge.' *Biological Theory*, 1(1): 20–2.
- , and Wilson, Deirdre. 1995. *Relevance: Communication and Cognition*, 2nd edn. Oxford: Blackwell (1st edn 1986).
- , Clément, Fabrice, Heintz, Christophe, Mascaro, Olivier, Mercier, Hugo, Origgi, Gloria, and Wilson, Deirdre. 2010. 'Epistemic Vigilance.' *Mind and Language* 25(4): 359–93.
- Tennie, Claudio, Frith, Uta, and Frith, Chris. 2010. 'Reputation Management in the Age of the World-Wide Web.' *Trends in Cognitive Sciences*, 14(11): 482–8.

- Willis, Janine, and Todorov, Alexander. 2006. 'First Impressions: Making up your Mind After a 100-ms Exposure to a Face.' *Psychological Science*, 17: 592–8.
- Wilson, Deirdre, and Sperber, Dan. 2012. *Meaning and Relevance*. Cambridge: Cambridge University Press.
- Vanderbilt, Kimberly, Lio, David, and Heyman, Gail. 2011. 'The Development of Distrust.' *Child Development*, 82(5): 1372–80.

DAN SPERBER is professor in the departments of philosophy and of cognitive science at the Central European University in Budapest and researcher at the Institut Jean Nicod in Paris. He has done work in anthropology, linguistics, philosophy and psychology. His latest book, co-authored with Deirdre Wilson, is *Meaning and Relevance* (Cambridge University Press, 2012).
