



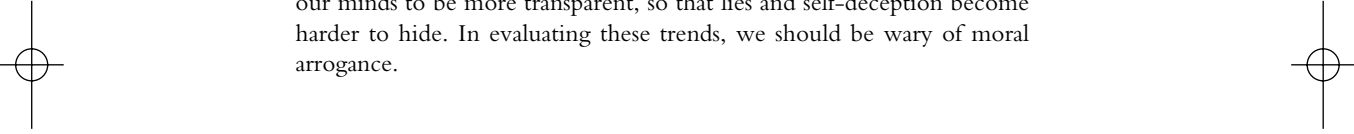
17

Enhancing Our Truth Orientation

*Robin Hanson**

Abstract

Humans lie and deceive themselves, and often choose beliefs for reasons other than how closely those beliefs approximate truth. This is mainly why we disagree. Three future trends may reduce these epistemic vices. First, increased documentation and surveillance should make it harder to lie and self-deceive about the patterns of our lives. Second, speculative markets can create a relatively unbiased consensus on most debated topics in science, business, and policy. Third, brain modifications may allow our minds to be more transparent, so that lies and self-deception become harder to hide. In evaluating these trends, we should be wary of moral arrogance.



Self-deception has done us proud. Without it . . . we might still be running naked through the forest. . . . Self-deception was a splendid adaptation in a world populated by nomadic bands armed with sticks and stones. It is no longer such a good option in a world stocked with nuclear and biological weapons. . . . The most dangerous forms of self-deception are the collective ones. Patriotism, moral crusades, and religious fervor sweep across nations like plagues. (Smith, 2004)

People will hold an opinion because they want to keep the company of others who share the opinion, or because they think it is the respectable opinion, or because they have publicly expressed the opinion in the past and would be embarrassed by a “U-turn,” or because the world would suit them better if the opinion were true. (Whyte, 2004)

For comments I thank Colleen Berndt, Nick Bostrom, Tyler Cowen, Hal Finney, Peter Kramer, Brett Paatsch, Julian Savulescu, Peter Singer, James Watson, Dan Wikler, and participants of the Transvision 2004 and the Oxford Uehiro Enhancing Humans 2004 conferences. I thank the Oxford Uehiro Center, the Center for Study of Public Choice, the Mercatus Center, and the International Foundation for Research in Experimental Economics for financial support.

* Department of Economics, George Mason University, November 2004 (revised June 2007).



Moral . . . opinions are often the result of self-interest, self-deception, historical and cultural accident, hidden class bias, and so on. (Daniels, 1979)

Introduction

Humans today have many epistemic virtues. We are clever animals who have discovered a vast division of labor, enabling our unprecedented and rapidly increasing power and understanding of science, industry, and more. But we also have many epistemic vices, such as using our knowledge to enable cruelty.

In particular, we often choose beliefs that are not the closest feasible approximation to the truth. By this I do not mean that we fail to spend all possible resources on obtaining or analyzing information. Instead, I mean that our beliefs are biased because we have motivations other than that of approximating truth while conserving resources (Mele, 2001). For example, because we like to think well of ourselves and our groups, we tend to overestimate our abilities and morality, exaggerate our influence, and take credit for our successes while blaming external factors for our failures (Giannetti, 1997).

This sort of self-deception is arguably the main reason humans disagree. It seems that honest truth-seeking agents should not knowingly disagree. Specifically, such agents should not be able to predict how others' future opinions will disagree with their own opinion about which of many possible worlds is the real one. Yet humans clearly do so disagree, especially on abstract topics in business, science, policy, and morality.

Apparently, beliefs are like clothes. In a harsh environment, we choose our clothes mainly to be functional, i.e., to keep us safe and comfortable. But when the weather is mild, we choose our clothes mainly for their appearance, i.e., to show our figure, our creativity, and our allegiances. Similarly, when the stakes are high we may mainly want accurate beliefs to help us make good decisions. But when a belief has few direct personal consequences, we in effect mainly care about the image it helps us to project.

Now admittedly, bias and self-deception do produce many personal and social benefits. For example, by over-estimating our ability we not only attract social allies, we also raise our self-esteem and happiness (Taylor, 1989), and motivate ourselves to excel (Kitcher, 1990). Depressed and

mentally-ill people tend to be less self-deceived than others. Nevertheless, bias and self-deception seem to be morally wrong overall and especially dangerous in a world with technologies as powerful as nuclear weapons and genetic engineering.

We humans have been enhancing ourselves for many thousands of years, through nutrition, education, gadgets, and social institutions. In the future we may also directly and substantially change and augment our brains. What are the prospects that these changes will increase our truth-orientation, and decrease our bias and self-deception?

In this chapter, I review the basics of disagreement and self-deception, and then consider three general types of enhancements which may increase our truth-orientation: documentation, idea futures, and mental transparency. Finally, I warn against the danger of moral arrogance when evaluating these trends.

Disagreement

People disagree all the time, and always have. While many specific disagreements do not seem entirely honest, most people believe honest disagreements are both possible and common.

In his most cited paper, Nobel Prize winner Robert Aumann proved in 1976 a remarkable result on “agreeing to disagree.” In technical terms, he showed that Bayesians with a common prior and common knowledge of their respective estimates of a random variable must have exactly the same estimates (Aumann, 1976). This result was provocative, but its assumptions seemed too strong to reflect much on human disagreement.

Aumann’s result has since been greatly generalized, however. For example, agents need not know, and know that they know, and know that they know that they know, and so on forever with perfect confidence. Instead, agents need only be reasonably sure that they are mutually aware of their disagreement (Monderer and Samet, 1989). Instead of knowing the exact estimates they would each say next, the agents need only be aware of whether one of them thinks another’s estimate will be higher or lower at some particular future date (Hanson, 2002).

In place of Bayesians, we need only agents who aspire to be Bayesian, even if they suffer severe computational limitations (Hanson, 2003b). And

in place of common priors, the agents need only believe that the process that created their belief tendencies treated them similarly (Hanson, 2006). This analysis applies to any dispute that can be described in terms of which of many possible worlds is the real world. So agents can be uncertain about logical truths, via impossible possible worlds (Hintikka, 1975), and this analysis can include the morality of actions, especially when moral intuitions are used as evidence.

With these generalizations, this Aumann-initiated literature does now seem relevant to real human disagreements. In fact, it suggests that honest truth-seeking agents, i.e., who seek to minimize differences between their stated beliefs and the truth, should not knowingly disagree on business, policy, science, or morality. Of course the attentive reader need not yet be persuaded of this claim; space has limited me to quickly summarizing results described elsewhere (Cowen and Hanson, 2002), and which I will elaborate in an upcoming book.

If honest truth-seeking agents would not knowingly disagree, then how is it that humans do knowingly disagree, even when they are confident of their honesty? The obvious explanation is that humans do not seek only to believe the truth; they have other motivations influencing the choice of their beliefs, such as hope and self-respect (otherwise known as wishful-thinking and overconfidence).

Note that humans are not only surprisingly unaware of these other belief influences; they are often downright hostile to the suggestion that they are under such influences. Thus humans seem to be “self-deceived” about their non-truth-seeking belief motivations, and this fact can plausibly explain most human disagreement. That is, even though we usually consciously see ourselves as just trying to believe and say what is true, motivated biases keep us both from listening to others as we should, and from seeing this failing in ourselves.

Why self-deception?

To better understand what we are up against, let us first review the standard account of why humans evolved to be biased and self-deceived. This account could well contain errors, but it seems the place to start.

If we could easily believe one thing while we say another, there would be little need for self-deception. We do not, however, find it easy to say one thing and consciously believe another. It turns out that because of the way our brains are structured, we give off many unconscious clues about what we are consciously thinking. These include facial expressions, tone of voice, reaction time, and much more. These clues make it hard to deceive others about what we are thinking.

It seems that our brains evolved to deal this problem by unconsciously biasing our beliefs (Lockard and Paulhaus, 1988; Trivers, 2000). Instead of just pretending to think well of ourselves, in order to convince others to think well of us, we *honestly* think well of ourselves, even in the face of contrary evidence.

Of course it is not quite that simple. You can't lie about just one thing; you also have to lie about lying, and so on. So we also seem to have evolved to not much notice our tendencies to self-deceive, even though we are well aware of such tendencies in others. Even though we are intuitively skilled at social maneuvering, we find it hard to think explicitly about social strategizing, making social science a latecomer to the sciences. We also avoid being very obviously biased on topics that others might too easily check and challenge us on. Yet we maintain strong biases on topics, like patriotism, where others are likely to challenge us for not being biased enough.

How can we not know and yet be so skilled? It may be that parts of our unconscious mind are well aware of many of our deceptions, and it is only our conscious mind that is fooled. Our conscious minds may even be like the public relations department of a corporation, whose main purpose is to present a certain coherent image to the outside, and not to make key corporate policy decisions.

In any case, there are limits to our ability to self-deceive; we cannot just make ourselves believe anything we would like to believe. For example, we cannot very easily self-deceive about our recent overt actions or our direct perceptions; these are just too obvious to question. If I just punched you, I would have to admit it. But this still leaves much room for bias; I might still tell myself that "you asked for it."

Now let us consider three trends which might reduce this tendency toward self-deception.

Documentation

Humans have long worked to document their lives, inventing gadgets to aid in writing and recording, concepts and conventions to make what we say meaningful and comparable, and social institutions to let us coordinate in monitoring and verifying our documentation. It is harder to lie, and so to self-deceive, about documented events.

Once upon a time, you could believe that you were a faster runner than most, by focusing on that one time when you blazed through the woods. Standardized sporting races made it harder, but you could still think you've never seen anyone go faster than in that race you won. With the invention of the stopwatch, even that deception got harder. In general, standardized tests and contests make it harder for people to lie and self-deceive about their knowledge and skills.

Similarly, standardized accounting makes it harder for corporate officers to steal corporate assets, and for households to avoid paying taxes. Lab and field notebooks, and standardized procedures, make it harder for researchers to lie about their observations. Standardized formal models make it harder for theorists to hide the biased assumptions that lead to biased conclusions. And standardized statistical procedures make it harder to bias the inferences we draw from all this data.

The memory required for these records is tiny compared to the memory of a single human brain. Similarly, the computation needed to implement these procedures is tiny compared to the computational ability of a single human brain. Even so, these records and procedures have had a great influence, because of their relative "objectivity." That is, these records were stored in a form that was relatively hard for us to consciously or unconsciously manipulate.

The price of such external memory and computation has been falling rapidly over the last half century, roughly following "Moore's Law." It should continue to fall at a similar pace for many decades to come.¹ As such prices fall, we will be able to document far more of our lives, and to make far more standardized inferences.

¹ Moore's Law is roughly that computing power and memory become twice as cheap every two years.

Many lament, and some celebrate (Brin, 1998), a coming “surveillance society.” Most web pages and email are already archived, and it is now feasible and cheap for individuals to make audio recordings of their entire lives. It will soon be feasible to make full video recordings as well. Add to this recordings by security cameras in stores and business, and most physical actions in public spaces may soon be a matter of public record. Private spaces will similarly be a matter of at least private record.²

Face recognition software will soon be sufficient to translate these recordings into a history of who was where when, and with whom. Speech recognition software will create searchable transcripts of everything said in public spaces. Cameras trained on eyes can also allow a record of who was looking where when. Simple visual processing should even be able to put our actions into standard categories such as sleeping, walking, talking, eating, etc.

Even more is possible if we regularly take standardized tests, or augment our records with indications of our subjective feelings. We might mark how happy we feel every hour, or how much we like each person we meet. Records will sometimes be faked, but redundant and independent sources should make faking hard. Even subjective records can be checked against body measurements, and cryptographic aids can prevent back-dating of records.

In addition, we will soon be able to fully record each person’s DNA, and will learn more and more about how to interpret such records to predict individual behavior.

The net result is that it will be much harder to lie about, or to self-deceive about, who said what specific words, who spent how much time with whom, what percent of each person’s time was spent in standardized categories, who is genetically similar to whom, and many other documentable facts.

Idea futures

Even massive documentation, however, does not obviously make it much harder to lie and self-deceive about abstract reasoning and conclusions. If

² Some hope that cryptography will let people work and play in complete anonymity by using virtual reality from secure rooms. Anonymous work would violate tax laws, however, and any correlation

you are biased in your reasons for preferring a political candidate, or for believing in a scientific theory, it seems hard to catch that bias just by looking at where you have been or what you have said.

For disputes that can eventually be resolved, it turns out that we have a social institution, feasible today, that is capable of eliminating much of this bias, producing relatively accurate unbiased estimates. That institution is idea futures, i.e., betting markets, also known as prediction markets or information markets (Wolfers and Zitzewitz, 2004; Hanson, 1999, 1995).

The basic idea is that people who have to put their money where their mouth is form beliefs in the equivalent of harsher weather, where they prefer more functional clothes. Those who know that they don't know tend to shut up, and those who don't know that they don't know lose their money and then shut up. Those who remain tend to know, and know that they know.

Decades of research on financial markets have shown that it is hard to find biases in market estimates. The few direct comparisons made so far have found markets to be at least as accurate as other institutions. Orange Juice futures improve on National Weather Service forecasts (Roll, 1984), horse race markets beat horse race experts (Figlewski, 1979), Oscar markets beat columnist forecasts (Pennock, Giles, and Nielsen, 2001), gas demand markets beat gas demand experts (Spencer, 2004), stock markets beat the official NASA panel at fingering the guilty company in the Challenger accident (Maloney and Mulherin, 2003), markets on economic statistics beat surveys of professional forecasters (Grkaynak and Wolfers, 2006), election markets beat national opinion polls (Berg, Nelson, and Rietz, 2001), and corporate sales markets beat official corporate forecasts (Chen and Plott, 1998).

Intentionally biasing these market estimates via trades turns out to be very difficult. Field observations have found that such attempts consistently fail, and controlled lab experiments have found no lasting effect on average prices or price accuracy from giving some traders such a manipulation incentive (Hanson, Oprea, and Porter, 2007). In fact, theoretical models suggest that manipulators are fundamentally no different from other noise

between your public and private persona would expose you, such as from your rhythm of speech or when you were sick or your connection was down.

traders, and so having more of them should on average increase price accuracy (Hanson and Oprea, 2004).

Today there are relatively few such markets. But this is largely because such markets are usually banned as illegal gambling. The same rapidly falling prices that are making documentation cheaper are also making these markets cheaper. In fact, it is possible to cheaply sustain billions of markets per person who trades in them (Hanson, 2003a). It is not entirely clear just how many topics are suitable for such market consensus, but it may be possible to apply similar methods to topics where no independent final judgment is possible (Prelec, 2004).

Once legal, or cheap enough to be widespread even if illegal, these markets can create unbiased consensus estimates on the observable policy consequences, on future scientific and moral consensus, and on much more. Having idea futures markets on most disputed topics could well change the character of conversation. Mass media might typically defer to market consensus on most disputed issues, just as business news now defers to stock prices. And pundits who dispute a market consensus might be expected to put their money where their mouth is. These markets could thus raise the cost of bias and self-deception in such debate, and offer the public precise continuously-updated consensus estimates, estimates that are consistent across a wide range of topics, and are relatively free of self-deception distortions.

Self-deception would remain

Even with far more documentation and market consensus, however, there would remain many topics where we can be biased and self-deceived.

Consider, for example, our level of interest and attachment in romantic and business relationships. Early in such relationships we tend to over-estimate how much we like our partners, and the chance we will remain together (Tennov, 1979). Late in such relationships, however, we tend to under-estimate these things. Both of these self-deceptions follow our biological interests; early on we want to convince others to invest in us, while later on we want them to think we will leave if not treated fairly.

Today we signal our interest and attachment in many ways, including in how we spend our time, attention, and money. Our ability to gauge

levels of interest and attachment of partners and potential partners should improve with more documentation about their activities. Market estimates of relationship duration given various conditions should also help. Even so, a great deal of uncertainty would likely remain, and hence our desire to perceive and deceive should also remain.

Mental Transparency

What will happen if we can enhance our minds, not just via education, but also via more direct modification and augmentation of our brains? As we get smarter, more things will become obvious to us, and so there will be more topics on which we find it hard to self-deceive. But we will also have opinions on more topics, creating more opportunities for self-deception.

One tempting enhancement will be to better control the various subconscious clues we give off, so that we can lie more convincingly. We might put our facial expressions under more conscious control, for example, and use hidden internal speed improvements to help us fake reaction times.

Deception generally helps those who are trying to look like something they are not, at the expense of those who really are what others pretend to be. Such deception also generally hurts observers, who are trying to tell the difference between these other two types of people. Of these three groups, two have a common interest in preventing deception.

Those who really are what others pretend to be, and want it to show on their face, would try to convince others that they are not controlling facial expressions. So they might try to use brain scanners or brain modifications to give others the ability to see more directly inside their minds, in order to directly confirm their honesty (Glover, 1984). Once such mental transparency becomes possible, everyone might feel pressured to have and use it, if they did not want others to presume they were hiding something.

This situation would be similar to the way we now prefer to meet someone we need to trust in person rather than via teleconferencing, telecommuting, phone, or email; we subconsciously reveal more about ourselves in person. Also similar is the way corporations hire outside accounting firms to monitor their internal processes, and nations allow outside inspectors to monitor their elections and weapons.

It should be relatively easy to show people that our brains are not attached to hidden electronic devices. But it may be hard to let others see inside our brains, at least in a way that they could make much sense of what they see; our brains were not designed with this sort of transparency in mind. And even if we can let others make enough sense of things to figure out our current beliefs, that might just push our self-deception back to the process that produced those current beliefs.

To deal with self-deception in belief production, we might want to provide audit trails, giving more transparency about the origins of our beliefs. However this feature may be even harder to allow via small modifications of our evolved brains. There may thus be natural limits on how far our brain structure may allow this internal transparency trend to go.

On the other hand, these considerations might encourage efforts to more substantially reorganize our mental architecture. We might eventually expect such reorganization to also be attempted for other reasons, such as to allow a closer integration with other minds and computer aids.

On moral arrogance

While we may abstractly approve of the epistemic virtues that these changes would produce, we might also be uncomfortable with their sheer magnitude. How human are creatures with transparent minds, documented lives, and market induced agreement on most abstract topics? Could it really be good to create such creatures instead of creatures more like us? And if the end result of these trends seems repugnant, should we not enlist our moral fervor to push public policy to regulate such behavior?

It is hard to escape reflecting on and acting on our simple moral intuitions about such issues, at least regarding our own actions (such as what children we create). It is also tempting to use these same conclusions to regulate the behavior of others, including if possible that of our distant descendants (such as by trying to prevent the development of certain technologies). There are often good reasons to resist such temptations, however.

One reason often mentioned is a value of autonomy; lives seem more worth living when people make their own choices, independent of what actual choices they make. There is, however, a more important but rarely mentioned consideration: the risk of moral arrogance.

People ordinarily consider moral issues when choosing their actions, and typically estimate their own actions to be moral. So if you, based on your moral analysis, estimate that other people's behavior should be regulated and limited for moral reasons, then it seems that you estimate your moral judgment to be better than theirs. If so, the literature on the rationality of disagreement should make you pause.

Of course you need not pause if you have good reason to think they do not believe they are acting morally. In this case there is no disagreement; there is just your good fighting their evil. The same applies when you would likely also do what they would do. For example, if taxes were voluntary, then you might well not pay your taxes, even though you think society is better off when all are forced to pay taxes. These are largely questions of coordination, not moral disagreement.

Setting such cases aside, however, we are left with cases where you disapprove of their moral judgments. If it is possible for you to advise them before they act, and if they reject your advice, then you and they will soon find yourself in a position where you knowingly disagree with each other. And if you both see yourselves as just trying to believe moral truth, then if honest truth-seeking agents would not knowingly disagree, we must conclude that you or they are self-deceived about your moral analysis.

You might be tempted to believe that your moral analysis is better because you have more information. You may, after all, be a respected and learned academic who took years to consider such issues while consulting with similar colleagues. However, when a rational (i.e., truth-seeking) agent with little information disagrees with a self-deceived agent with much information, it is the rational agent who is more likely to be right.

The question is not who started the conversation knowing more, but rather who better learns by listening to others. Yes, by thinking too highly of their ability to analyze moral questions, people may take too little of your moral advice. But by thinking too highly of your own moral ability, you may offer too much moral advice. It is not obvious which bias is more severe.

You might argue that you are less likely to be self-deceived because you are a distant observer, with little direct interest in the actions in question. But such distance could tempt you to indulge general fantasies, instead of attending to important details. And what if the people you disagree with rely on the advice of other distant observers who disagree with you?

Perhaps your moral advice is based on the belief that most everyone, including yourself, is subject to a particular bias. For example, you might think most everyone is biased to quit school early, and so should be forced to stay in school. If true, this might be a reason to regulate behavior. But the question is why you are better positioned than they to evaluate such a claim. Those who reject your claim may, for example, see you as being in a position that makes you especially likely to be biased in favor of such a belief. You might, for example, be a teacher, or be beholden to a teacher union.

In general, you may point to particular indications suggesting that others are more biased than you, indications which do in fact correlate with who is more biased. But remember that those you disagree with can and do play the same game; humans are skilled at selectively finding reasons to dismiss others as more biased than they. Perhaps the safest rule is to avoid extreme positions, relative to the existing distribution of opinions, as these seem more likely to be the result of self-deception. But again we should beware of the human tendency to selectively apply such rules.

The warning to beware of our self-deception regarding our moral abilities would seem to apply with a special force to those who argue the virtues of self-deception. After all, does not the pro-self-deception side in a debate seem more likely to be self-deceived in this matter?

We should also be especially wary of moral arrogance regarding the moral behavior of our distant descendants, as those descendants will have a clear information advantage over us; we cannot listen to them as we could when arguing with a contemporary. Our descendants will know of our advice, and also of many other things we do not know. In addition, this chapter has suggested they may well have a stronger truth orientation than us.

Conclusion

I have briefly surveyed three trends which may cause our descendants to have a stronger truth-orientation than we do: more recording and standardized statistics on our lives, speculative markets in major disputed topics, and minds enhanced to allow more transparency. These three trends should reduce the epistemic vice of human self-deception, a vice that seems

particularly dangerous in a world that differs so much from the environment to which those self-deceptions were adapted. A world with so much less self-deception and disagreement seems jarringly different from our world. Of course we are vastly different from our ancestors a hundred thousand years ago—our minds have been radically modified by education, culture, and civilization.

Even so, some may be tempted to reject such a truth-oriented future, regulating behavior to discourage it. Others may be tempted to regulate behavior to encourage a truth-orientation. Those who are so tempted should consider the risk of moral arrogance. How sure can they be that their moral judgment is better than the judgments of those they would regulate, especially the judgments of our better informed and more-truth-oriented descendants?

References

- Aumann, R. (1976). 'Agreeing to Disagree', *The Annals of Statistics*, 4 (6), 1236–9.
- Berg, J., Nelson, F., and Rietz, T. (2001). 'Accuracy and Forecast Standard Error of Prediction Markets'. Tech. rep., University of Iowa, College of Business Administration.
- Brin, D. (1998). *Transparent Society*. Reading, Massachusetts: Addison-Wesley.
- Chen, K.-Y., and Plott, C. R. (1998). 'Prediction markets and information aggregation mechanism: Experiments and application'. Tech. rep., California Institute of Technology.
- Cowen, T., and Hanson, R. (2002). 'Are Disagreements Honest?'. Tech. rep., George Mason University Economics.
- Daniels, N. (1979). 'Wide Reflective Equilibrium and Theory Acceptance in Ethics', *Journal of Philosophy*, 76 (5), 256–82.
- Figlewski, S. (1979). 'Subjective Information and Market Efficiency in a Betting Market', *Journal of Political Economy*, 87 (1), 75–88.
- Giannetti, E. (1997). *Lies We Live By—The Art of Self-Deception*. London: Bloomsbury.
- Glover, J. (1984). *What Sort of People Should There Be?* New York: Penguin Books.
- Grkaynak, R. S., and Wolfers, J. (2006). 'Macroeconomic Derivatives: An Initial Analysis of Market-Based Macro Forecasts, Uncertainty and Risk'. Tech. rep. 11929, NBER.
- Hanson, R. (1995). 'Could Gambling Save Science? Encouraging an Honest Consensus', *Social Epistemology*, 9 (1), 3–33.

- ____ (1999). 'Decision Markets', *IEEE Intelligent Systems*, 14 (3), 16–19.
- ____ (2002). 'Disagreement is Unpredictable', *Economics Letters*, 77, 365–9.
- ____ (2003a). 'Combinatorial Information Market Design', *Information Systems Frontiers* 5 (1), 105–19.
- ____ (2003b). 'For Savvy Bayesian Wannabes, Are Disagreements Not About Information?', *Theory and Decision*, 54 (2), 105–23.
- ____ (2006). 'Uncommon Priors Require Origin Disputes', *Theory and Decision*, 61 (4), 318–28.
- ____ and Oprea, R. (2008). 'Manipulators Increase Information Market Accuracy', *Economica*, forthcoming.
- ____ Oprea, R., and Porter, D. (2006). 'Information Aggregation and Manipulation in an Experimental Market', *Journal of Economic Behavior and Organization*, 60(4): 449–59.
- Hintikka, J. (1975). 'Impossible Possible Worlds Vindicated', *Journal of Philosophical Logic*, 4, 475–84.
- Kitcher, P. (1990). 'The Division of Cognitive Labor', *Journal of Philosophy*, 87 (1), 5–22.
- Lockard, J., and Paulhaus, D. (eds.), (1988). *Self-Deception: An Adaptive Mechanism?* Englewood Cliffs, N. J.: Prentice Hall.
- Maloney, M. T., and Mulherin, J. H. (2003). 'The complexity of price discovery in an efficient market: the stock market reaction to the Challenger crash', *Journal of Corporate Finance*, 9 (4), 453–79.
- Mele, A. R. (2001). *Self-Deception Unmasked*. Princeton: Princeton University Press.
- Monderer, D., and Samet, D. (1989). 'Approximating Common Knowledge with Common Beliefs', *Games and Economic Behavior*, 1, 170–90.
- Pennock, D. M., Giles, C. L., and Nielsen, F. A. (2001). 'The Real Power of Artificial Markets', *Science*, 291, 987–8.
- Prelec, D. (2004). 'A Bayesian truth serum for subjective data', *Science*, 306, 462–6.
- Roll, R. (1984). 'Orange Juice and Weather', *American Economic Review*, 74 (5), 861–80.
- Smith, D. L. (2004). *Why We Lie—The Evolutionary Roots of Deception and the Unconscious Mind*. New York: St. Martin's Press.
- Spencer, J. (2004). 'New ICAP-Nymex Derivatives Have U.S. Gas Market's Number', *Wall Street Journal*, 4 Aug., p.1.
- Taylor, S. (1989). *Positive Illusions: Creative Self-Deception and the Healthy Mind*. New York: Basic Books.
- Tennov, D. (1979). *Love and Limerence*. Chelsea, MI: Scarborough House.

- Trivers, R. (2000). 'The Elements of a Scientific Theory of Self-Deception'. In LeCroy, D., and Moller, P. (eds.), *Evolutionary Perspectives on Human Reproductive Behavior*, vol. 907, *Annals of the New York Academy of Sciences*.
- Whyte, J. (2004). *Crimes Against Logic*. New York: McGraw Hill.
- Wolfers, J., and Zitzewitz, E. (2004). 'Prediction Markets', *Journal of Economic Perspectives*, 18 (2), 107–26.