Irrationality: Philosophical Aspects

Richard Samuels, The Ohio State University, Columbus, OH, USA
Stephen P Stich, Rutgers, The State University of New Jersey, New Brunswick, NJ, USA

Abstract

Over the past four decades, human reason and rationality has been among the most intensely investigated topics in psychology, cognitive science, and economics. At the heart of this debate is a view of human rationality, often associated with the *Heuristics and Biases* tradition, on which much of our reasoning and decision making is normatively problematic because it relies on heuristics and biases rather than rational principles. In this article we describe briefly some of the evidence that has been invoked in support of this contention and consider some of the more prominent critical responses, especially from evolutionary psychology and dual-process theorists.

The Aristotelian Tradition vs the Heuristics and Biases Program

There is a philosophical tradition, reaching at least as far back as Aristotle, which maintains that man is a rational animal. According to this Aristotelian traditional, human beings are *intrinsically* rational, though of course even the most ardent defender of this view would admit that people's decisions and thought processes are sometimes very irrational indeed. When people are tired, or drunk, or in the grip of rage, they can think and act in ways that no account of rationality would condone. To accommodate this obvious fact, contemporary defenders of human rationality often appeal to a distinction between *competence* and *performance* similar to the one invoked by contemporary linguists (Chomsky, 1965; Samuels et al., 1999). The claim that human beings are rational is interpreted to mean that normal humans posses a rational reasoning competence – a tacit knowledge of normatively appropriate principles for reasoning and decision making. Under ordinary circumstances these principles are reflected in people's thinking and their actions, though sometimes, for various reasons, our performance can fail to reflect this underlying rational competence.

In the early 1970s, however, Amos Tversky, Daniel Kahneman, and a number of other psychologists began reporting findings suggesting that the Aristotelian view of human rationality is seriously mistaken. What these studies demonstrated is that even under quite ordinary circumstances where fatigue, drugs, and strong emotions are not factors, people reason and make decisions in ways that appear systematically to violate familiar canons of rationality on a broad array of problems. Those first surprising studies sparked the growth of an enormously influential research program whose impact has been felt in a wide range of disciplines including psychology, economics, political theory, and medicine. The expression *Heuristics and Biases* is often used as a label for this tradition, since it holds that much of our reasoning and decision making relies on Heuristics and Biases rather than rational principles.

In the sections to follow we will offer some examples of the sort of experimental findings that many have interpreted as showing that the reasoning and decision-making competence of normal humans is *far* from rational. (For more comprehensive reviews of the literature see Dawes, 1988; Baron, 2008; Gilovich et al., 2002.) But first we need to say what standards of rationality are being applied by those who think the experimental findings demonstrate widespread irrationality. Though writers in the Heuristics and Biases tradition rarely offer an explicit and general theory of rationality, the most plausible reading of their work is that they are assuming some version of what Edward Stein (1996) calls the *Standard Picture*. On this view, to be rational is to reason and make decisions in accordance with principles that are based on the rules of logic, probability theory, and decision theory. Sometimes, of course, it is far from clear how these theories are to be applied – a problem that we will return to in Section Do These Studies Have 'Bleak Implications' for the Extent of Human Rationality?. For the moment, however, we follow the lead of those who work in the Heuristics and Biases tradition and assume that the irrationality manifested in the experiments we describe is intuitively obvious.

Irrationality in Reasoning

The Conjunction Fallacy

In what has become perhaps the most famous experiment in the Heuristics and Biases tradition, Tversky and Kahneman (1982) presented people with the following task.

> Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.
>
> Please rank the following statements by their probability, using 1 for the most probable and 8 for the least probable.
>
> (1) Linda is a teacher in elementary school.
> (2) Linda works in a bookstore and takes Yoga classes.
> (3) Linda is active in the feminist movement.
> (4) Linda is a psychiatric social worker.
> (5) Linda is a member of the League of Women Voters.
> (6) Linda is a bank teller.
> (7) Linda is an insurance sales person.
> (8) Linda is a bank teller and is active in the feminist movement.

In a group of naive subjects with no background in probability and statistics, 89% judged that statement (8) was more

International Encyclopedia of the Social & Behavioral Sciences, Second Edition, 2015, 719–723

probable than statement (6) despite the obvious fact that one cannot be a *feminist* bank teller unless one is a *bank teller*. When the same question was presented to statistically sophisticated subjects – graduate students in the decision science program of the Stanford Business School – 85% made the same judgment! Results of this sort, in which subjects judge that a compound event or state of affairs is more probable than one of the components of the compound, have been found repeatedly since Tversky and Kahneman's pioneering studies, and they are remarkably robust. This pattern of reasoning has been labeled *the conjunction fallacy*.

Base-Rate Neglect

On the familiar Bayesian account, the probability of a hypothesis on a given body of evidence depends, in part, on the prior probability of the hypothesis. However, in a series of experiments, Kahneman and Tversky (1973) showed that subjects often seriously undervalue the importance of prior probabilities. One of these experiments presented half of the subjects with the following 'cover story'.

> A panel of psychologists have interviewed and administered personality tests to 30 engineers and 70 lawyers, all successful in their respective fields. On the basis of this information, thumbnail descriptions of the 30 engineers and 70 lawyers have been written. You will find on your forms five descriptions, chosen at random from the 100 available descriptions. For each description, please indicate your probability that the person described is an engineer, on a scale from 0 to 100.

The other half of the subjects were presented with the same text, except the 'base rates' were reversed. They were told that the personality tests had been administered to 70 engineers and 30 lawyers. Some of the descriptions that were provided were designed to be compatible with the subjects' stereotypes of engineers, though not with their stereotypes of lawyers. Others were designed to fit the lawyer stereotype, but not the engineer stereotype. And one was intended to be quite neutral, giving subjects no information at all that would be of use in making their decision. Here are two examples, the first intended to sound like an engineer, the second intended to sound neutral:

> Jack is a 45-year-old man. He is married and has four children. He is generally conservative, careful, and ambitious. He shows no interest in political and social issues and spends most of his free time on his many hobbies which include home carpentry, sailing, and mathematical puzzles.
>
> Dick is a 30-year-old man. He is married with no children. A man of high ability and high motivation, he promises to be quite successful in his field. He is well liked by his colleagues.

As expected, subjects in both groups thought that the probability that Jack is an engineer is quite high. Moreover, in what seems to be a clear violation of Bayesian principles, the difference in cover stories between the two groups of subjects had almost no effect at all. The neglect of base-rate information was even more striking in the case of Dick. That description was constructed to be totally uninformative with regard to Dick's profession. Thus the *only* useful information that subjects had was the base-rate information provided in the cover story. But that information was entirely ignored. The median probability estimate in both groups of subjects was 50%.

Before leaving the topic of base-rate neglect, we want to offer one further example illustrating the way in which the phenomenon might well have serious practical consequences. Here is a problem that Casscells et al. (1978) presented to a group of faculty, staff, and fourth-year students at Harvard Medical School.

> If a test to detect a disease whose prevalence is 1/1000 has a false positive rate of 5%, what is the chance that a person found to have a positive result actually has the disease, assuming that you know nothing about the person's symptoms or signs? ____%

Under the most plausible interpretation of the problem, the correct Bayesian answer is 2%. But only 18% of the Harvard audience gave an answer close to 2%. Forty-five percent of this distinguished group completely ignored the base-rate information and said that the answer was 95%.

Irrationality in Decision Making

Framing

In a study that is widely believed to illustrate a deeply irrational feature of human decision making, Tversky and Kahneman (1981) presented a group of subjects with the following problem:

> Imagine that the US is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the programs are as follows:
>
> If Program A is adopted, 200 people will be saved.
>
> If Program B is adopted, there is a 1/3 probability that 600 people will be saved, and a 2/3 probability that no people will be saved.

A second group of subjects was given an identical problem, except that the programs were described as follows:

> If Program C is adopted, 400 people will die.
>
> If Program D is adopted, there is a 1/3 probability that nobody will die and a 2/3 probability that 600 people will die.

On the first version of the problem most subjects chose Program A. But on the second version most chose Program D, despite the fact that the outcome described in A is identical to the one described in C.

The Ultimatum Game

Consider a game that works as follows: There are two players, an 'Allocator' and a 'Recipient' whose identities are unknown to each other. At the beginning of the game the Allocator is given a sum of money (say $10) and asked to divide it between himself and the other player. The only restriction is that the Allocator must offer *something*. So, if the smallest unit of money available is $1, then the Allocator must offer at least this much to the Recipient. Once the Allocator has divided the money, the

Recipient gets to decide whether to accept or reject the amount offered. If the Recipient accepts, each player gets to keep the sum specified by the Allocator. If the Recipient rejects the offer, then both players receive *nothing*. The game is not repeated, and the players know from the outset that they will play only once.

On one widely accepted analysis of the problem, it is never rational for a Recipient to reject an offer, since it is better to get *some* money rather than none. If this is right, then it is always rational for the Allocator to offer the Recipient the smallest amount of money possible, since rational Recipients will always accept the offer, and this will maximize the amount of money the Allocator receives. However, in many studies in both developed and developing countries, subjects very rarely make these choices. Indeed, in these studies it is often the case that the modal offer is to divide the money equally. Moreover, Recipients frequently reject offers of less than 25% of the total, preferring instead to go home empty-handed. (For an overview of the literature see Thaler, 1992.)

The Public Goods Game and Costly Punishment

The Public Goods Game is another game that has been widely used to study economic decision making. In one version of this game, four subjects are each given a sum of money (say $10) and are offered the opportunity of contributing some or all of this money to a group fund. Decisions are made without discussion or negotiation, and subjects are told that they will play the game only once with this set of subjects. After all the subjects have had an opportunity to contribute to the fund, the money in the fund is multiplied by a fixed amount (50% for example) and then distributed equally among all the players, regardless of their contribution. On one analysis, a rational agent concerned to maximize his economic gain should contribute nothing to the group fund, since that will enable him to keep his initial $10 and to get 25% of whatever the other subjects have contributed to the fund. However, in many studies it has been shown that this is rarely what subjects do. The majority of subjects contribute something to the group fund, and many subjects contribute more than 50% of the sum they are initially given.

In one frequently used alternative version of the Public Goods Game, subjects are offered the opportunity to 'punish' other players whose contribution to the group fund displeases them. Specifically, after the game described above is completed, subjects are told that they can spend some portion of their game money to take money from another player, even though they, themselves, will not be allowed to keep the money. They might, for example, be permitted to spend $1 to take $3 from another player, under the proviso that the $4 will be kept by the experimenter. Obviously, under such circumstances, both players are worse off financially. Moreover, the game is designed to minimize the collateral benefits to punishers: decisions about punishment are made privately, and subjects are aware that they will not play the game again with the same set of people. Nevertheless, in a manner that runs quite contrary to standard economic analyses, subjects often chose to spend their money to punish other players who have made smaller than average contributions to the group fund (Yamagishi, 1986).

Do These Studies Have 'Bleak Implications' for the Extent of Human Rationality?

The results we have sketched are just a tiny sample of an enormous and growing literature on human reasoning and decision making. But what conclusions about human rationality – and irrationality – should we draw from these findings? On this question opinions are sharply divided. In one often quoted comment, two leading investigators in the Heuristics and Biases tradition claimed that the experimental results have 'bleak implications' for human rationality (Nisbett and Borgida, 1975). Other researchers have concluded that "individuals are generally affected by systematic deviations from rationality" (Bazerman and Neale, 1986). And still others have suggested that the fault may be in our genes, or at least in our evolutionary history – "that people lack the correct programs for many important judgmental tasks…[because we]… have not had the opportunity to evolve an intellect capable of dealing conceptually with uncertainty" (Slovic et al., 1976). But many other theorists have argued that conclusions like these are unwarranted. In what follows we discuss four of the more prominent responses to the 'bleak implications' view and the evidence adduced in its support. (For an extended review of the debate see Samuels et al., 2002.)

Problems with Specific Experiments

There is a substantial literature arguing that one or another experiment in the Heuristics and Biases tradition is flawed and thus does not really demonstrate that the subjects' performance is irrational. Some of the putative flaws are methodological. One charge that obviously has some merit is that it is often less than clear that subjects understand the problem in the way that experimenters want them to. For example, principles of conversational implicature might lead subjects in the 'feminist bank teller' experiment to assume that choice (6) (Linda is a bank teller) means that she is a bank teller who is *not* active in the feminist movement (Schwarz, 1996). Other putative flaws turn on the way that the experimenters propose to apply the normative principles to particular problems. For instance, Gigerenzer (2000) argues that there are typically several different ways in which the principles of statistics and probability can be applied to a given problem and that these different analyses of the problem lead to different answers – or in some cases to no answer at all. If this is correct, then obviously we cannot conclude that subjects are being irrational simply because they do not give the answer that the experimenters prefer.

Darwinian Responses

A second prominent response to the claim that the human mind is prone to 'systematic deviations from rationality' comes from work by evolutionary psychologists. These theorists argue that if we are adequately to assess the extent of human rationality, we need to pay due attention to issues about the evolved structure and function of our cognitive systems. Further, they maintain that attention to such matters provides grounds for rejecting the 'bleak implications' view.

Informational Format

One especially prominent instance of this sort of critique arises in research regarding our ability to reason about probabilities. Evolutionary psychologists proceed by noting that the effective use of probabilistic information would have been highly advantageous to our hominid forebears, and therefore, they maintain, it would be surprising if we had not evolved 'an intellect capable of dealing conceptually with uncertainty'. But, they continue, we should expect this capacity to manifest itself most clearly when probabilistic information is presented in a format that would have been common in ancestral environments. In these environments, probabilistic information would almost invariably have been presented in the form of information about frequencies as opposed to single events. Thus, evolutionary psychologists argue, we should expect people to do much better on probabilistic reasoning tasks if the problems are presented in a frequency format. And it appears that they do. In the 'feminist bank teller' problem, for example, if the description of Linda is followed by a question like this one:

There are 100 people who fit the description above. How many of them are
…
(6) bank tellers?
…
(8) bank tellers and active in the feminist movement?
…

the number of subjects who commit the conjunction fallacy drops from over 90% to only about 10%.

Further evidence comes from Cosmides and Tooby's (1996) systematic exploration of the 'Harvard Medical School problem' discussed in Section Base-Rate Neglect. They showed that subjects find it much easier to use base-rate information rationally on 'frequentist' versions of the problem, like the one that follows.

1 out of every 1000 Americans has disease X. A test has been developed to detect when a person has disease X. Every time the test is given to a person who has the disease, the test comes out positive. But sometimes the test also comes out positive when it is given to a person who is completely healthy. Specifically, out of every 1000 people who are perfectly healthy, 50 of them test positive for the disease.
Imagine that we have assembled a random sample of 1000 Americans. They were selected by lottery. Those who conducted the lottery had no information about the health status of any of these people.
Given the information above:
on average,
How many people who test positive for the disease will *actually* have the disease? _____ out of _____.

On this version of the problem, 76% of subjects – as opposed to the mere 18% in Casscells et al.'s original study – gave the correct Bayesian answer.

Preference Structure

Evolutionary psychologists have also proposed a related response to the claim that subjects who make substantial offers in the Ultimatum Game contribute to the group fund in the Public Goods Game or pay to punish low contributors in the Public Good Game are being irrational. They concede that it may be irrational to make these decisions if one's only goal in these games is to maximize one's financial gain. But, they note, money is not the only thing that people value. Many people also value fairness and thus they assign positive utility both to being fair in their own offers and contributions, and to punishing others whom they judge to have behaved unfairly in these games. Further, since people in many cultures make fair offers in the Ultimatum Game and pay to punish below average contributions in Public Goods Games, some evolutionary psychologists have suggested that this preference for fairness may be innate and pan-cultural (Baumard et al., 2013).

The suggestion that there is a pan-cultural preference for fairness has not gone unchallenged, however. For example, Henrich et al. (2004) conducted Ultimatum Game experiments in 15 small-scale societies, and found that while people in some of these groups made offers that were similar to those made by subjects in previously studied groups, people in other groups exhibited a wide range of behaviors (for instance, among the Machiguenga, a group in the Peruvian Amazon, the modal Allocator offers only 15% of the money provided by the experimenter, and low offers are almost always accepted). Even more surprising are the results of Gächter and colleagues (Herrmann et al., 2008) who found that in a number of societies scattered around the globe subjects in Public Goods Games with Costly Punishment would pay to punish other players who had contributed *more* than average to the group fund. Ongoing research is aimed at explaining these decisions, which appear to be blatantly irrational and to pose a prima facie challenge to the hypothesis that humans have an innate preference for fairness.

Ecological Rationality

In light of findings such as those just mentioned, some evolutionary psychologists have maintained that our normative standards of rationality require modification. Most influentially, Gigerenzer and collaborators have proposed that for many purposes a conception of rationality that they call *ecological rationality* may be more appropriate as a criterion for normative assessment than the Standard Picture typically assumed by Heuristics and Biases researchers (Gigerenzer et al., 1999). On this view, an inferential or judgmental strategy is ecologically rational if it is accurate and efficient on the sorts of tasks that were important in the environments in which we evolved. If we interpret Aristotle's thesis as the claim that *man is an ecologically rational animal*, evolutionary psychologists maintain that there is a growing body of evidence indicating that he was right (Gigerenzer and Gaissmaier, 2011).

Dual Process Accounts of Human Reasoning

Another response to the 'bleak implications' view that has become quite influential in recent years – even among early advocates of the Heuristics and Biases program – comes from dual-process theories of reasoning (Evans, 2008; Kahneman, 2011; Sloman, 1996; Stanovich, 1999). Dual-process theorists accept that people's reasoning and decision making often deviate systematically from appropriate norms of rationality; but they also maintain that there are many contexts in which people perform quite well. More significantly, they advocate a kind of explanation of this pattern that appears to chart a 'middle way' between those who maintain that normatively problematic heuristics are the *only* reasoning resources people possess, and

those evolutionary psychologists who suggest that *all* reasoning is subserved by normatively unproblematic cognitive mechanisms.

Though dual-process theories differ from one another in many details (Evans and Frankish, 2009), they all propose that reasoning and decision making are subserved by two quite different sorts of system. One kind of system is typically characterized as relatively fast, holistic, automatic, largely unconscious, and requires relatively little cognitive capacity. The other is relatively slow, rule based, more readily controlled, and requires significantly more cognitive capacity. Further, systems of the former kind are widely assumed to have emerged relatively early in human evolution and, as evolutionary psychologists suggest, to have been shaped by natural selection to do a good job on the problems that would have been important to our hominid forebears (Stanovich, 1999, 2004). The other kind of system, by contrast, is assumed to have evolved more recently, is more heavily influenced by culture and formal education, and is often more adept at dealing with many of the problems posed by a modern, technologically advanced and highly bureaucratized society (Stanovich, 1999, 2004).

Since the new system requires more cognitive capacity, is more influenced by culture and education, and does not get used automatically, Stanovich hypothesized that there might be significant individual differences in people's ability and inclination to use it. More specifically, he reasoned, people with higher cognitive capacity, as measured by instruments like the Scholastic Test (SAT), and with a cognitive style that emphasizes 'epistemic self-regulation' should do better on tasks that the old system was not designed to handle. Stanovich agrees with evolutionary psychologists that many of the tasks studied in the Heuristics and Biases tradition fall into this category. In extensive studies of these tasks he has shown that, while the average performance on these tasks is indeed quite poor, there are *some* subjects who give the answer that the Standard Picture suggests is normatively correct on *many* of these problems, and these subjects typically have significantly higher SAT scores and score higher on tests designed to detect cognitive styles that include epistemic self-regulation.

If Stanovich and other dual-process theorists are on the right track, then the unbridled optimism sometimes suggested by evolutionary psychologists is unwarranted, since most untutored people do indeed lack the capacity to deal with a wide range of problems that are important in a technological society. But the glum pessimism often associated with the Heuristics and Biases tradition is not warranted either. Since the fast, automatic, and evolutionarily older system requires little cognitive capacity, everyone has the capacity to deal rationally with many reasoning and decision-making problems that were important in the environment in which we evolved. Moreover, since the new, slow, rule-based system can be significantly affected by education, there is reason to hope that better educational strategies will improve people's performance on those problems that the old system was not designed to deal with.

See also: Bayesian Theory, History, Applications, and Contemporary Directions; Decision Making, Psychology of; Decision Making: Nonrational Theories; Decision and Choice: Bounded Rationality; Decision and Choice: Economic Psychology; Decision and Choice: Heuristics; Moral Reasoning in Psychology; Practical Reasoning: Philosophical Aspects; Probability Theory: Formal; Probability Theory: Interpretations.

Bibliography

Baron, J., 2008. Thinking and Deciding. Cambridge University Press, Cambridge.
Baumard, N., André, J., Sperber, D., 2013. A mutualistic approach to morality: the evolution of fairness by partner choice. Behavioral and Brain Sciences 36, 59–122.
Bazerman, M., Neale, M., 1986. Heuristics in negotiation. In: Arkes, H., Hammond, K. (Eds.), Judgment and Decision Making: An Interdisciplinary Reader. Cambridge University Press, Cambridge.
Casscells, W., Schoenberger, A., Grayboys, T., 1978. Interpretation by physicians of clinical laboratory results. New England Journal of Medicine 299, 999–1000.
Chomsky, N., 1965. Aspects of the Theory of Syntax. MIT Press, Cambridge, MA.
Cosmides, L., Tooby, J., 1996. Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. Cognition 58, 1–73.
Dawes, R., 1988. Rational Choice in an Uncertain World. Harcourt Brace Jovanovich, Orlando, FL.
Evans, J.St.B.T., 2008. Dual-processing accounts of reasoning, judgment and social cognition. Annual Review of Psychology 59, 255–278.
Evans, J.St.B.T., Frankish, K., 2009. In Two Minds: Dual Processes and Beyond. Oxford University Press, Oxford, UK.
Gigerenzer, G., 2000. Adaptive Thinking: Rationality in the Real World. Oxford University Press, New York.
Gigerenzer, G., Gaissmaier, W., 2011. Heuristic decision making. Annual Review of Psychology 62, 451–482.
Gigerenzer, G., Todd, P., The ABC Research Group, 1999. Simple Heuristics that Make Us Smart. Oxford University Press, New York.
Gilovich, T., Griffin, D., Kahneman, D., 2002. Heuristics and Biases: The Psychology of Intuitive Judgment. Cambridge University Press, Cambridge, UK.
Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., 2004. Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies. Oxford University Press, Oxford, UK.
Herrmann, B., Thöni, C., Gächter, S., 2008. Antisocial punishment across societies. Science 319, 1362–1367.
Kahneman, D., 2011. Thinking, Fast and Slow. Farrar, Straus and Giroux, New York, NY.
Kahneman, D., Tversky, A., 1973. On the psychology of prediction. Psychological Review 80, 237–251.
Nisbett, R., Borgida, E., 1975. Attribution and the social psychology of prediction. Journal of Personality and Social Psychology 32, 932–943.
Samuels, R., Stich, S., Tremoulet, P., 1999. Rethinking rationality: from bleak implications to Darwinian modules. In: LePore, E., Pylyshyn, Z. (Eds.), What Is Cognitive Science. Blackwell, Oxford.
Samuels, R., Stich, S., Bishop, M., 2002. Ending the rationality wars: how to make disputes about human rationality disappear. In: Elio, R. (Ed.), Common Sense, Reasoning and Rationality. Oxford University Press, Oxford, UK.
Schwarz, N., 1996. Cognition and Communication: Judgmental Biases, Research Methods and the Logic of Conversation. Erlbaum, Hillsdale, NJ.
Sloman, S.A., 1996. The empirical case for two systems of reasoning. Psychological Bulletin 119, 3–22.
Slovic, P., Fischhoff, B., Lichtenstein, S., 1976. Cognitive processes and societal risk taking. In: Carol, J.S., Payne, J.W. (Eds.), Cognition and Social Behavior. Erlbaum, Hillsdale, NJ.
Stanovich, K.E., 1999. Who is Rational? Studies of Individual Differences in Reasoning. Erlbaum, Mahwah, NJ.
Stanovich, K.E., 2004. The Robot's Rebellion: Finding Meaning the Age of Darwin. University of Chicago Press, Chicago, IL.
Stein, E., 1996. Without Good Reason. Clarendon Press, Oxford.
Thaler, R., 1992. The Winner's Curse: Paradoxes and Anomalies of Economic Life. Princeton University Press, Princeton, NJ.
Tversky, A., Kahneman, D., 1981. The framing of decisions and the psychology of choice. Science 211, 453–458.
Tversky, A., Kahneman, D., 1982. In: Kahneman, Slovic, Tversky (Eds.), Judgments of and by Representativeness.
Yamagishi, T., 1986. The provision of a sanctioning system as a public good. Journal of Personality and Social Psychology 51 (1), 110–116.