TOPiCS

TOPICS IN COGNITIVE SCIENCE

http://onlinelibrary.wiley.com/journal/10.1111/(ISSN)1756-8765

Moral Satisficing: Rethinking Moral Behavior as Bounded Rationality

Gerd Gigerenzer

Max Planck Institute for Human Development, Berlin

Abstract

What is the nature of moral behavior? According to the study of bounded rationality, it results not from character traits or rational deliberation alone, but from the interplay between mind and environment. In this view, moral behavior is based on pragmatic social heuristics rather than moral rules or maximization principles. These social heuristics are not good or bad per se, but solely in relation to the environments in which they are used. This has methodological implications for the study of morality: Behavior needs to be studied in social groups as well as in isolation, in natural environments as well as in labs. It also has implications for moral policy: Only by accepting the fact that behavior is a function of both mind and environmental structures can realistic prescriptive means of achieving moral goals be developed.

Keywords: Moral behavior; Social heuristics; Bounded rationality

1. Introduction

What is the nature of moral behavior? I will try to answer this question by analogy with another big question: What is the nature of rational behavior? One can ask whether morality and rationality have much to do with one another, and an entire tradition of moral philosophers, including Hume and Smith, would doubt this. Others, since at least the ancient Greeks and Romans, have seen morality and rationality as two sides of the same coin, albeit with varying meanings. As Cicero (De finibus 3, 75–76) explained, once reason has taught the ideal Stoic—the wise man—that moral goodness is the only thing of real value, he is happy forever and the freest of men, since his mind is not enslaved by desires. Here, reason makes humans moral. During the Enlightenment, the theory of probability emerged and with it a new vision of rationality, once again tied to morality, which later evolved into various forms

Correspondence should be sent to Gerd Gigerenzer, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany. E-mail: gigerenzer@mpib-berlin.mpg.de

of consequentialism in ethics. In the 20th century, the notion of bounded rationality arose in reaction to the Enlightenment theory of (unbounded) rationality and its modern versions. In this essay, I ask: What vision of moral behavior emerges from the perspective of bounded rationality?

I will use the term *moral behavior* as short for behavior in morally significant situations, subsuming actions evaluated as moral or immoral. The study of bounded rationality (Gigerenzer, 2008a; Gigerenzer & Selten, 2001a; Simon, 1990) examines how people actually make decisions in an uncertain world with limited time and information. Following Herbert A. Simon, I will analyze moral behavior as a result of the match between mind and environment, as opposed to an internal perspective of character or rational reflection. My project is to use a structure I know well—the study of bounded rationality—and ask how it would apply to understanding moral behavior. I argue that *much* (not all) of moral behavior is based on heuristics. A heuristic is a mental process that ignores part of the available information and does not optimize, meaning that it does not involve the computation of a maximum or minimum. Relying on heuristics in place of optimizing is called *satisficing*. To prefigure my answer to the above question, the analogy between bounded rationality and morality leads to five propositions:

1. *Moral behavior is based on satisficing, rarely on maximizing.* Maximizing (finding the provably best course of action) is possible in ''small worlds'' (Savage, 1954) where all alternatives, consequences, and probabilities are known with certainty, but not in ''large worlds'' where not all is known and surprises can happen. Given that the certainty of small worlds is rare, normative theories that propose maximization can seldom guide moral behavior. But can maximizing at least serve as a normative goal? The next proposition provides two reasons why this may not be so.

2. *Satisficing can reach better results than maximizing.* There are two possible cases. First, even if maximizing is feasible, relying on heuristics can lead to better (or worse) outcomes than when relying on a maximization calculus, depending on the structure of the environment (Gigerenzer & Brighton, 2009). This result contradicts a view in moral philosophy that satisficing is a strategy whose outcome is or is expected to be second-best rather than optimal. ''Everyone writing about satisficing seems to agree on at least that much'' (Byron, 2004, p. 192). Second, if maximizing is not possible, trying to approximate it by fulfilling more of its conditions does not imply coming closer to the best solution, as the theory of the second-best proves (Lipsey, 1956). Together, these two results challenge the normative ideal that maximizing can generally define how people ought to behave.

3. *Satisficing operates typically with social heuristics rather than exclusively moral rules.* The heuristics underlying moral behavior are often the same as those that coordinate social behavior in general. This proposition contrasts with the moral rules postulated by rule consequentialism, as well as the view that humans have a specially ''hardwired'' moral grammar with rules such as ''don't kill.''

4. *Moral behavior is a function of both mind and the environment.* Moral behavior results from the match (or mismatch) of the mental processes with the structure of the social

environment. It is not the consequence of mental states or processes alone, such as character, moral reasoning, or intuition.

5. *Moral design.* To improve moral behavior towards a given end, changing environments can be a more successful policy than trying to change beliefs or inner virtues.

This essay should be read as an invitation to discuss morality in terms of bounded rationality and is by no means a fully fledged theory of moral satisficing. Following Hume rather than Kant, my aim is not to provide a normative theory that tells us how we ought to behave, but a descriptive theory with prescriptive consequences, such as how to design environments that help people to reach their own goals. Following Kant rather than Hume, moral philosophers have often insisted that the facts about human psychology should not constrain ethical reflection. I believe that this poses a risk of missing essential insights. For instance, Doris (2002) argued that the conception of character in moral philosophy is deeply problematic, because it ignores the evidence amassed by social psychologists that moral behavior is not simply a function of character, but of the situation or environment as well (e.g., Mischel, 1968). A normative theory that is uninformed of the workings of the mind or impossible to implement in a mind (e.g., because it is computationally intractable) is like a ship without a sail. It is unlikely to be useful and to help make the world a better place.

My starting point is the Enlightenment theory of rational expectation. It was developed by the great 17th-century French mathematician Blaise Pascal, who together with Pierre Fermat laid down the principles of mathematical probability.

2. Moral behavior as rational expectation under uncertainty

Should one believe in God? Pascal's (1669/1962) question was a striking heresy. Whereas scores of earlier scholars, from Thomas Aquinas to René Descartes, purported to give a priori demonstrations of the divine existence and the immortality of the soul, Pascal abandoned the necessity of God's existence in order to establish moral order. Instead, he proposed a calculus to decide whether or not it is rational to believe that God exists (he meant God as described by Roman Catholicism of the time). The calculus was for people who were convinced neither by the proofs of religion nor by the arguments of the atheists and who found themselves suspended between faith and disbelief. Since we cannot be sure, the result is a bet, which can be phrased in this way:

Pascal's Wager: If I believe in God and He exists, I will enjoy eternal bliss; if He does not exist, I will miss out on some moments of worldly lust and vice. On the other hand, if I do not believe in God, and He exists, then I will face eternal damnation and hell.

However small the odds against God's existence might be, Pascal concluded that the penalty for wrongly not believing in Him is so large and the value of eternal bliss for correctly believing is so high that it is prudent to wager on God's existence and act as if one believed in God—which in his view would eventually lead to actual belief. Pascal's argu-

ment rested on an alleged isomorphism between decision problems where objective chances are known and those where the objective chances are unknown (see Hacking, 1975). In other words, he made a leap from what Jimmy Savage (1954), the father of modern Bayesian decision theory, called ''small worlds'' (in which all alternatives, consequences, and probability distributions are known, and no surprises can happen) to what I will call ''large worlds,'' in which uncertainty reigns. For Pascal, the calculus of expectation served as a general-purpose tool for decisions under uncertainty, from games of chance to moral dilemmas: Evaluate every alternative (e.g., to believe in God or not) by its n consequences, that is, by first multiplying the probabilities p_i with the values x_i of each consequence i ($i = 1$, …, n), and then summing up:

$$EV = \sum p_i x_i \tag{1}$$

The alternative with the highest expected value (EV) is the rational choice.

This new vision of rationality emphasized risk instead of certainty and subsequently spread through the Enlightenment in many incarnations: as Daniel Bernoulli's expected utility, Benjamin Franklin's moral algebra, Jeremy Bentham's hedonistic calculus, and John Stuart Mill's utilitarianism, among others. From its inception, the calculus of expectation was closely associated with moral and legal reasoning (Daston, 1988). Today, it serves as the foundation on which rational choice theory is built.

For instance, Gary Becker tells the story that he began to think about crime in the 1960s after he was late for an oral examination and had to decide whether to put his car in a parking lot or risk getting a ticket for parking illegally on the street. ''I calculated the likelihood of getting a ticket, the size of the penalty, and the cost of putting the car in a lot. I decided it paid to take the risk and park on the street'' (Becker, 1995, p. 637). In his view, violations of the law, be they petty or grave, are not due to an irrational motive, a bad character, or mental illness, but can be explained as rational choice based on the calculus of expectation. This economic theory has policy implications: Punishment works and criminals are not ''helpless'' victims of society. Moreover, city authorities should apply the same calculus to determine the optimal frequency of inspecting vehicles, the size of the fine, and other variables that influence citizens' calculations whether it pays to violate the law.

In economics and the cognitive sciences, full (unbounded) rationality is typically used as a methodological tool rather than as an assumption about how people actually make decisions. The claim is that people behave *as if* they maximized some kind of welfare, by calculating Bayesian probabilities of each consequence and multiplying these by their utilities. As a model of the mind, full rationality requires reliable knowledge of all alternative actions, their consequences, and the utilities and probabilities of these consequences. Furthermore, it entails determining the best of all existing alternatives, that is, being able to compute the maximum expectation.

The calculus of expectation provided the basis for various (act-)consequentialist theories of moral behavior, according to which actions are to be judged solely by their consequences, and therefore are not right or wrong per se, even if they use torture or betrayal.

The best moral action is the one that maximizes some currency—the expected value, welfare, or the greatest happiness of the greatest number. Depending on what is being maximized, many versions of consequentialism exist; for instance, maximizing happiness may refer to the total amount of happiness, not the total number of people, or vice versa (Braybrooke, 2004). Whereas consequentialism sets up a normative ideal of what one should do (as opposed to observing what people actually do), the calculus of expectation has also influenced descriptive theories of behavior. Versions of Eq. 1 have been proposed in theories of health behavior, consumer behavior, intuition, motivation, attitude formation, and decision making. Here, what *ought to be* provided the template for theories of *what is*. This ought-to-is transfer is a widespread principle for developing new descriptive theories of mind, as illustrated by Bayesian and other statistical optimization theories (Gigerenzer, 1991). Even descriptive theories critical of expected utility maximization, such as prospect theory, are based on the same principles: that people make decisions by looking at all consequences and then weighting and summing some function of their probabilities and values—differing only on specifics such as the form of the probability function (e.g., linear or S-shaped). Hence, the calculus of expectation has become one of the most successful templates for human nature.

2.1. Morality in small worlds

The beauty and elegance of the calculus of expectation comes at a price, however. To build a theory on maximization limits its domain to situations where one can find and prove the optimal solution, that is, well-defined situations in which all relevant alternatives, consequences, and probabilities are known. As mentioned earlier, this limits the experimental studies to ''small worlds'' (Binmore, 2009; Savage, 1954). Much of decision theory, utilitarian moral philosophy, and game theory focuses on maximizing, and their experimental branches thus create small worlds in which behavior can be studied. These range from experimental games (e.g., the ultimatum game) to moral dilemmas (e.g., trolley problems) to choices between monetary gambles. Yet this one-sided emphasis on small-world behavior is somewhat surprising given that Savage spent the second half of his seminal book on the question of decision making in ''large worlds,'' where not all alternatives, consequences, and probability distributions are known, and thus maximization is no longer possible. He proposed instead the use of heuristics such as minimax—to choose the action that minimizes the worst possible outcome, that is, the maximum loss. This part of Savage's work anticipated Herbert Simon's concept of bounded rationality, but few of Savage's followers have paid attention to his warning that his maximization theory should not be routinely applied outside small worlds (Binmore, 2009, is an exception).

Maximization, however, has been applied to almost everything, whether probabilities are known or not, and this overexpansion of the theory has created endless problems (for a critique, see Bennis, Medin, & Bartels, in press). Even Pascal could not spell out the numbers needed for his wager: the prior probabilities that God exists, or the probabilities and values for each of the consequences. These gaps led scores of atheists, including Richard Dawkins (2006), to criticize Pascal's conclusion and propose instead numerical values and

probabilities to justify that it is rational *not* to believe in God. None of these conclusions, however, follow from the maximization calculus per se, because both sides can always pick particular probabilities and utilities in order to justify their a priori convictions. The fact that maximization limits the domain of rationality and morality to small worlds is one of the motivations for searching for other theories.

2.2. When conditions for maximization cannot be fulfilled, should one try to approximate?

Many moral philosophers who propose maximization of some kind of utility as normative concede that in the real world—because of lack of information or cognitive limitations—computing the best moral action turns out to be impossible in every single case. A standard argument is that maximization should be the ideal to aspire for, that is, to reach better decisions by more approximation. This argument, however, appears inconsistent with the *general theory of the second-best* (Lipsey, 1956). The theory consists of a general theorem and one relevant negative corollary. Consider that attaining an optimal solution requires simultaneously fulfilling a number of preconditions. The general theorem states that if one of these conditions cannot be fulfilled, then the other conditions, although still attainable, are in general no longer desirable. In other words, if one condition cannot be fulfilled (because of lack of information or cognitive limitations), the second-best optimum can be achieved by departing from all the other conditions. The corollary is as follows:

> Specifically, it is *not* true that a situation in which more, but not all, of the optimum conditions are fulfilled is necessarily, or is even likely to be, superior to a situation in which fewer are fulfilled. It follows, therefore, that in a situation in which there exist many constraints which prevent the fulfillment of the Paretian optimum conditions, the removal of any one constraint may affect welfare of efficiency either by raising it, by lowering it, or by leaving it unchanged. (Lipsey, 1956, p. 12)

Thus, the theory of the second-best does not support the argument that when maximization is unfeasible because some preconditions are not fulfilled, it should nevertheless be treated as an ideal to be approximated by fulfilling other conditions in order to arrive at better moral outcomes. The theory indicates that maximization cannot be a sound gold standard for large worlds in which its conditions are not perfectly fulfilled.

I now consider an alternative analogy for morality: bounded rationality.

3. Moral behavior as bounded rationality

How should one make decisions in a large world, that is, without knowing all alternatives, consequences, and probabilities? The "heresy" of the 20th century in the study of rationality was—and still is considered so in many fields—to dispense with the ideal of maximization in favor of bounded rationality. The term *bounded rationality* is attributed to Herbert A. Simon, with the qualifier *bounded* setting his vision apart from that of

''unbounded'' or ''full'' rationality, as represented by the calculus of expectation and its modern variants. Bounded rationality dispenses with the idea that optimization is the sine qua non of a theory of rationality, making it possible to deal with problems for which optimization is unfeasible, without being forced to reduce these to small worlds that accommodate optimization. As a consequence, the bounds to information and computation can be explicitly included as characteristics of a problem. There are two kinds of bounds: those in our minds, such as limits of memory, and those in the world, such as noisy, unreliable samples of information (Todd & Gigerenzer, 2001).

When I use the term *bounded rationality*, I refer to the framework proposed by Herbert A. Simon (1955, 1990) and further developed by others, including Reinhard Selten and myself (Gigerenzer & Selten, 2001a,b; Gigerenzer, Todd, & the ABC Research Group, 1999). In short, bounded rationality is the study of the cognitive processes (including emotions) that people actually rely on to make decisions in the large world. Before I explain key principles in the next sections, I would like to draw your attention to the fact that there are two other, very different interpretations of the concept of bounded rationality.

First, Ken Arrow (2004) argued that bounded rationality is ultimately optimization under constraints, and thus nothing but unbounded rationality in disguise—a common view among economists as well as some moral philosophers. Herbert Simon once told me that he wanted to sue people who misuse his concept for another form of optimization. Simon (1955, p. 102) elsewhere argued ''that there is a complete lack of evidence that, in actual human choice situations of any complexity, these computations can be, or are in fact, performed.'' Second, Daniel Kahneman (2003) proposed that bounded rationality is the study of deviations between human judgment and full rationality, calling these cognitive fallacies. In Kahneman's view, although optimization is possible, people rely on heuristics, which he considers second-best strategies that often lead to errors. As a model of morality, Arrow's view is consistent with those consequentialist theories that assume the maximization of some utility while adding some constraints into the equation, whereas Kahneman's view emphasizes the study of discrepancies between behavior and the utilitarian calculus, to be interpreted as moral pitfalls (Sunstein, 2005). Although these two interpretations appear to be diametrically opposed in their interpretation of actual behavior as rational versus irrational, both accept some form of full rationality as the norm. However, as noted, optimization is rarely feasible in large worlds, and—as will be seen in the next section— even when it is feasible, heuristic methods can in fact be superior.

I now introduce two principles of bounded rationality and consider what view of morality emerges from them.

4. Principle one: Less can be more

Optimizing means to compute the maximum (or minimum) of a function and thus determine the best action. The concept of *satisficing*, introduced by Simon, is a Northumbrian term for *to satisfy*, and is a generic term for strategies that ignore information and involve

little computation. These strategies are called *heuristics*. Note that Simon also used the term *satisficing* for a specific heuristic: Choosing the first alternative that satisfies an aspiration level. I will use the term here in its generic sense. The classical account of why people would rely on heuristics is the accuracy–effort trade-off: Compared to relying on a complex calculus, relying on a heuristic can save effort at the cost of some accuracy. Accordingly, in this view, a heuristic is second-best in terms of accuracy, because less effort can never lead to more accuracy. This viewpoint is still prevalent in nearly all textbooks today. Yet research on bounded rationality has shown that this trade-off account is not generally true; instead, the heuristic can be both more accurate and less effortful (see Gigerenzer & Brighton, 2009):

> Less-can-be-more: If a complex calculus leads to the best outcome in a small world, the same calculus may lead to an outcome inferior to that of a simple heuristic when applied in a large world.

For instance, Harry Markowitz received his Nobel Prize for an optimal asset allocation method known as mean-variance portfolio (currently advertised by banks worldwide), yet when he made his own investments for retirement, he did not use his optimization method. Instead, he relied on an intuitive heuristic known as *1/N: Allocate your money equally to each of N alternatives* (Gigerenzer, 2007). Studies showed that 1/*N* in fact outperformed the mean-variance portfolio in terms of various financial criteria, even though the optimization method had 10 years of stock data for estimating its parameters (more than many investment firms use). One reason for this striking result is that estimates generally suffer from sampling error, unless one has sufficiently large samples, whereas 1/*N* is immune to this kind of error because it ignores past data and has no free parameters to estimate. For $N = 50$, one would need a sample of some 500 years of stock data in order for the optimization model to eventually lead to a better outcome than the simple heuristic (DeMiguel, Garlappi, & Uppal, 2009). This investment problem illustrates a case where optimization can be performed (the problem is computationally tractable), yet the error in the parameter estimates of the optimization model is larger than the error due to the ''bias'' of the heuristic. In statistical terminology, the optimization method suffers mainly from *variance* and the heuristic from *bias*; the question of how well a more flexible, complex method (such as a utility calculus) performs relative to a simple heuristic can be answered through the *bias–variance dilemma* (Geman, Bienenstock, & Doursat, 1992). In other words, the optimization method would result in the best outcome if the parameter values were known without error, as in a small world, but it can be inferior in a large world, where parameter values need to be estimated from limited samples of information. By analogy, if investment were a moral action, maximization would not necessarily lead to the best outcome. This is because of the error in the estimates of the probabilities and utilities that this method generates in an uncertain world.

The investment example also illustrates that the important question is an ecological one. In which environments does optimization lead to better outcomes than satisficing (answer

for the investment problem: sample size is ≥500 years), and in which does it not (answer: sample size is <500 years)? Contrary to the claim that heuristics are always second-best, there is now broad evidence that simple heuristics that ignore information can outperform strategies that use more information and computation (e.g., Brighton, 2006; Gigerenzer, 2008a; Makridakis & Hibon, 2000).

Heuristics perform well precisely because they ignore information. The take-the-best heuristic, which relies on one good reason alone and ignores the rest, has been shown in many situations to predict more accurately than a multiple regression that relies on all available reasons (Czerlinski, Gigerenzer, & Goldstein, 1999). The tit-for-tat heuristic memorizes only the last of the partner's actions and forgets the rest (a form of forgiving) but can lead to better cooperation and higher monetary gain than more complex strategies do, including the rational strategy of always defecting (e.g., in the prisoner's dilemma with a fixed number of trials). Similarly, $1/N$ ignores all previous information about the performance of investment funds. In each case the question is: In what environment does simplicity pay, and where would more information help?

As the Markowitz example illustrates, experts tend to rely on fast and frugal heuristics in order to make better decisions (Shanteau, 1992). Here is more evidence. To predict which customers are active and which are nonactive in a large database, experienced managers of airlines and apparel businesses rely on a simple *hiatus heuristic*: *Customers who have not made a purchase for 9 months are considered inactive*. This heuristic has been shown to be more accurate than sophisticated methods such as the Pareto/NBD (negative binomial distribution) model, which uses more information and relies on complex computations (Wübben & Wangenheim, 2008). British magistrates appear to base bail decisions on a fast-and-frugal decision tree (Dhami, 2003), most professional burglars' choice of target objects follow the take-the-best heuristic rather than the weighting and adding of all cues (Garcia-Retamero & Dhami, 2009), and baseball outfielders' intuitions about where to run to catch a fly ball are based on the gaze heuristic and its variants (Gigerenzer, 2007; Shaffer, Krauchunas, Eddy, & McBeath, 2004). Some of the heuristics used by people have also been reported in studies on birds, bats, rats, and other animals (Hutchinson & Gigerenzer, 2005).

In sum, there is evidence that people often rely on heuristics and, most important, that by ignoring part of the information, heuristics can lead to better decisions than more complex strategies do, including optimization methods. An important point is that optimizing and satisficing are defined by their process (optimizing strategies compute the maximum for a function using all information while heuristics employ limited search for a few important pieces of information and ignore the rest). The process should not be confused with the outcome. Whether optimizing or satisficing leads to better outcomes in the real, uncertain world is an empirical question. This result contradicts the widespread belief that complex calculation always leads to better decisions than are made using some simple heuristic. And, I believe, it provides a challenge to the related normative ideal that maximizing can define how people ought to behave in an uncertain world.

5. Principle two: Simon's scissors

Theories of unbounded rationality are typically based on logical principles, such as axioms of consistency and transitivity. Logic was also Piaget's metaphor for understanding thinking; when turning to moral judgment, he proposed that it follows the same development, with abstract logical thought as its final stage (Gruber & Vonèche, 1977). Similarly, Kohlberg (1968) entitled one of his essays ''The child as a moral philosopher,'' asserting that moral functioning involves systematic thinking, along with emotions. Bounded rationality, in contrast, is based on an ecological rather than a logical view of behavior, as articulated in Simon's scissors analogy:

> Human rational behavior (and the rational behavior of all physical symbol systems) is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor. (Simon, 1990, p. 7)

Behavior is a function of both mind and environment. By looking at only one blade, one will not understand how scissors cut. Likewise, by studying only the mind, be it by interview or brain imaging, one will only partially understand the causes of behavior. Similarly, the norm for evaluating a decision is not simply consistency, transitivity, or other logical principles, but success in the world, which results from the match between mind and environments. Consistency can still be relevant in a theory of bounded rationality if it fulfills a functional goal, such as in checking the validity of a mathematical proof. The study of *ecological rationality* asks in which world a given heuristic is better than another strategy, as measured by some currency.

By analogy, in this view of rationality, moral behavior is a function of mind and environments rather than the consequence of moral reasoning or character alone.

6. Moral satisficing

Based on these two principles, I will attempt to sketch out the basis for a theory of moral satisficing. The theory has two goals:

1. *Explanation of moral behavior.* This goal is descriptive, that is, to explain how moral behavior results from the mind on the one hand and environmental structure on the other.
2. *Modification of moral behavior.* This goal is prescriptive, that is, to use the results under (1) to derive how a given moral goal can be realized. The solution can involve changing heuristics, designing environments, or both.

As mentioned before, moral satisficing is not a normative theory that tells us what one's moral goals should be—whether or not you should sign up as an organ donor, divorce your spouse, or stop wasting energy to protect the environment. But the theory can tell us how to

help people reach a given goal more efficiently. Furthermore, a descriptive analysis can shed light on the plausibility of normative theories (Feltz & Cokely, 2009; Knobe & Nichols, 2008).

For illustration, let me begin with the organ donor problem (Johnson & Goldstein, 2003). It arises from a shortage of donors, contributes to the rise of black markets for selling organs, and sparks ongoing debates about government intervention and the rights of individuals.

The organ donor problem. Every year, an estimated 5,000 Americans and 1,000 Germans die waiting in vain for a suitable organ donor. Although most citizens say that they approve of postmortem organ donation, relatively few sign a donor card: only about 28% and 12% in the United States and Germany, respectively. Why do so few sign up as potential donors? Various explanations have been proposed, such as that many people are selfish and have little empathy for the suffering of others, that people are hypersensitive to a postmortem opening of their bodies, and that people fear that doctors will not work as hard to save them in an emergency room situation. Yet why are 99.9% of the French and Austrians potential donors?

Normally, the striking difference in the rates of potential donors between countries (Fig. 1) is likely to be explained by personality traits—selfishness, or lack of empathy. But character is unlikely to explain the big picture. Consider next reasoning as an explanation. If
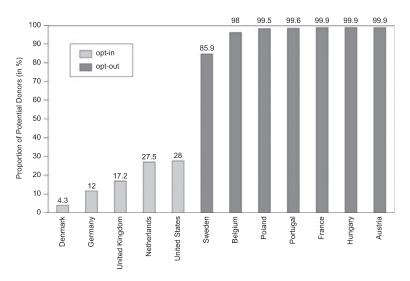

Fig. 1. Why are so few citizens in Denmark, Germany, United Kingdom, Netherlands, and the United States potential organ donors, compared to in the other countries? The answer is not a difference in character or knowledge. Rather, most citizens in all 12 countries appear to rely on the same heuristic: *If there is a default, do nothing about it.* The outcome, however, depends on the default setting, and therefore varies strikingly between countries with opt-in policies and opt-out policies. There are nuances with which these systems function in practice that are not shown here. In the United States, some states have an opt-in policy, whereas others force citizens to make a choice (based on Johnson & Goldstein, 2003).

moral behavior were the result of deliberative reasoning rather than personality traits such as selfishness, however, the problem might be that most Americans, Germans, or Dutch are not aware of the need for organs. This explanation would call for an information campaign. In the Netherlands, an exhaustive campaign was in fact conducted, with 12 million letters sent to a population of 16 million. As in similar attempts, the effect was practically nil. Nevertheless, in a survey, 70% of the Dutch said they would like to receive an organ from someone who has died, should it be necessary, and only 16% said they were not willing to donate (Persijn, 1997). This natural experiment suggests that missing information is not the issue either. What then lies at the root of the organ donor problem?

Here is a possible answer. Despite the striking differences in Fig. 1, most people seem to rely on the same heuristic:

If there is a default, do nothing about it.

The *default heuristic* leads to different outcomes because environments differ. In explicit-consent countries such as United States, Germany, and the Netherlands, the law is that nobody is a donor unless you *opt in*. In presumed-consent countries such as France and Austria, the default is the opposite: Everyone is a donor, unless you *opt out*. From a rational choice perspective, however, this should have little effect because people are assumed to ignore a default if it is inconsistent with their preference. As Fig. 1 also shows, a minority of citizens do not follow the default heuristic, and more of these opt in than out, consistent with the preference of the majority. These citizens might rely on some version of Pascal's moral calculus or on the golden rule. An online experiment found similar results (Johnson & Goldstein, 2003). Americans were asked to assume that they had just moved into a new state and were given the choice to confirm or change their donor status. In one group, being a donor was the default; in a second group, not being a donor was the default; and in a third group, there was no default. Even in this hypothetical situation where zero effort was necessary for overriding the default, more than 80% ended up as potential donors when the default was being a donor, compared to only half as many when the default was not being one. In the absence of a default, the far majority signed up as donors.

Why would so many people follow this heuristic? Johnson and Goldstein's experimental results indicate that it is not simply irresponsible laziness, because when participants could not avoid the effort of making a decision, the heuristic was still followed by many. Rather, the heuristic appears to serve a function that has been proposed as the original function of morality: the coordination of groups of individuals (Darwin, 1981; pp. 161–167; Wilson, 2002). Relying on defaults, specifically legal defaults, creates homogeneity within a society and thus helps cement it together. In general, I think that heuristics that guide moral behavior might be those that help to coordinate behavior, feeling, motion, and emotion.

I will now use the case of organ donation to illustrate two hypotheses that follow from the analogy with bounded rationality.

Hypothesis 1: Moral behavior = f(mind, environment).

Several accounts, normative and descriptive, from virtue theories to Kohlberg's six stages of moral reasoning, assume that forces inside the mind—moral intuition or reasoning—are

or should be the cause of moral behavior (unless someone actively prevents a person from executing it by threat or force). These theories provide a conceptual language for only one blade of Simon's scissors: the mind. As the organ donation problem illustrates, however, the second blade, the environment, contributes profoundly to the resulting outcome. Many who feel that donation is a good thing nevertheless do not opt in, and some of those who believe that donation is not a good thing do not opt out. Here, I would like to propose the notion of *ecological morality*, that is, that moral behavior results from an interaction between mind and environment. I think of ecological morality as a descriptive (how to explain moral behavior?) and prescriptive (how to improve moral behavior given a goal?) research program, but I can imagine that it also has the potential for the basis of a normative theory. Here, one would expand questions such as ''What is our duty?'' and ''What is a good character?'' into interactive questions such as ''What is a virtuous environment for humans?''

The main contribution of social psychology to the study of morality has been in demonstrating the power of social environments, including the power to make ordinary people exhibit morally questionable behavior. Milgram's (1974) obedience experiments and Zimbardo's (2007) prisoners' study are two classics. For instance, in one of Milgram's studies (Experiment 5), the experimenter instructed the participant to administer electric shocks of increasing intensity to a learner in a paired-associate learning task every time the learner gave an incorrect answer. The unsettling finding was that 83% of the participants continued to administer shocks beyond the 150-V level, and 65% continued to administer shocks all the way to the end of the range—450 V in 15-V increments. Would people still obey today? Yes, the power of the environment is still remarkable. In a partial replication (with the 150-V level), the obedience rate in 2006 was 70%, only slightly lower than what Milgram found 45 years earlier (Burger, 2009). This result was obtained despite participants having been told thrice that they could withdraw from the study at any time and still retain their $50 participation fee.

These studies document the strong influence of the situation on moral behavior, as opposed to internal accounts in terms of ''personality types'' and the like, as described in the literature on the authoritarian personality (Adorno, Frenkel-Brunswik, Levinson, & Sanford, 1950). When Milgram asked psychiatrists and Yale seniors to predict the outcome of the experiment, the predicted obedience rates were only 0.1% and 1.2%, respectively (Blass, 1991). Milgram (1974) himself said that his major motivation for conducting the studies was his own refusal to believe in the situation-based excuses made by concentration camp guards at the Nuremberg War Trials (e.g., ''I was just following orders''). His experimental results changed his mind: ''Often, it is not so much the kind of person a man is as the kind of situation in which he finds himself that determines how he will act'' (p. 205). A situational understanding of behavior does not justify behavior, but it can help us to understand it and avoid what is known as the *fundamental attribution error*: to explain behavior by internal causes alone. To better understand why more than 100 million people died violently at the hands of others during the entire 20th century (Doris, 2002), we need to analyze the environmental conditions that foster people's willingness to kill and inflict physical pain on others. This is not to say that character variables do not matter (Cokely & Feltz, 2009; Funder, 2001), but these express themselves in the interaction with specific environments.

Here I want to address two direct consequences of the ecological approach to moral behavior: inconsistency between moral intuition∕reasoning and behavior, and moral luck.

6.1. Systematic inconsistencies

Inconsistencies between moral intuition and behavior are to be expected from an ecological view of morality. Moreover, one can predict in what situation inconsistencies are more likely to arise, such as when there is no match between intuition and default. For instance, a survey asked citizens whether they would be willing to donate an organ after they have died; 69% and 81% of Danish and Swedish citizens, respectively, answered ''yes,'' compared to about 4% and 86% (Fig. 1) who are actually potential donors (Commission of the European Communities, 2007). The Danish appear to behave inconsistently; the Swedish do not. Inconsistencies or only moderate correlations between moral intuition and behavior have been reported in studies that both elicited people's moral intuitions and observed their behavior in the same situation (e.g., Gerson & Damon, 1978; Narvaez & Lapsley, 2005). Consider premarital sexual relations and American teenagers who publicly take a vow of abstinence. These teenagers typically come from religious backgrounds and have revived virginity as a moral value, as it had been up to the first half of the 20th century. One would expect that their moral intentions, particularly after having been declared in public, would guide their behavior. Yet teens who made a virginity pledge were just as likely to have premarital sex as their peers who did not (Rosenbaum, 2009). The difference was that when those who made the pledge had sex, they were less likely to use condoms or other forms of contraception. We know that teenagers' behavior is often guided by a coordination heuristic, called *imitate-your-peers*: *Do what the majority of your peers do*. If my friends make a virginity pledge, I will too; if my friends get drunk, I will too; if my friends already have sex at age 16, I will too; and so on. If behavior is guided by peer imitation, a pledge in itself makes little difference. Moreover, if the heuristic works unconsciously but teenagers nonetheless believe that their behavior is totally under their control, this would explain why they are not prepared for the event of acting against their stated moral values. The U.S. Government spends about $200 million a year on abstinence-promotion programs, which seem to be as ineffective in preventing unwanted pregnancies as the Dutch mass mail campaign was in boosting organ donations.

6.2. Moral luck

Matheson (2006) was troubled that bounded rationality implies a post-Enlightenment picture of ''cognitive luck'' because ''in that case, there is little we can do to improve our cognitive abilities, for—so the worry continues—such improvement requires manipulation of what it is within our power to change, and these external, cognitively fortuitous features are outside that domain'' (p. 143). I conjecture that cognitive luck is an inevitable consequence that one should not worry about or try to eliminate but instead use constructively for better theories. Similarly, moral philosophers have discussed the question of ''moral luck.'' It arises from the fact that moral behavior is in part determined by our environment and thus not entirely controlled by the individual, and it concerns the question whether behavior

should be evaluated as right or wrong depending on its result shaped by situational circumstances (Statman, 1993; Williams, 1981). Nagel (1993, p. 59) defines moral luck as follows: ''Where a significant aspect of what someone does depends on factors beyond his control, yet we continue to treat him in that respect as an object of moral judgment, it can be called moral luck.'' Nagel claims that despite our intuition that people cannot be morally assessed for what is not their fault, we nevertheless frequently make moral judgments about people based on factors out of their control.

The worry about moral luck is based on the assumption that internal ways to improve cognition and morality are under our control, or should be, whereas the external ways are not. Changing environments, however, can sometimes be more efficient than changing minds, and creating environments that facilitate moral virtue is as important as improving inner values (Gigerenzer, 2006; Thaler & Sunstein, 2008). The donor problem illustrates this conjecture, where thousands of lives could be saved every year if governments introduced proper defaults rather than continuing to bet on the wrong ''internal'' psychology and send letters to their citizens. For instance, a 2008 European Union Bulletin (Bundesärztekammer, 2008) maintains the importance of increasing public awareness of the organ donation problem and urges the member states of the European Union to disseminate more information on it. This governmental policy is unwilling to grant citizens the benefit of a proper environment that respects human psychology and helps people to reach their own goals.

Hypothesis 2: The same social heuristics guide moral and nonmoral behavior.

You may have noticed that I avoided using the term *moral heuristics*. There is a reason for this. The term would imply that there are two different kinds of heuristics, those for moral decisions and those for self-regarding ones, that is, matters of personal taste. On the contrary, I believe that, as a rule, one and the same heuristic can solve both problems that we call moral and those we do not (Gigerenzer, 2008b). Let me explain why. The boundaries between what is deemed a moral issue shift over historical time and between cultures. Although contemporary Western moral psychology and philosophy often center on the issues of harm and individual rights, such a constrained view of the domain of morality is unusual in history. There existed more important moral values than avoiding harm to individuals. Abraham was asked by the Lord to kill his son, and his unquestioning readiness to heed God's command signaled a higher moral value, faith. For the ancient world, where human sacrifice was prevalent, the surprising part of the story was that God stopped the sacrifice (Neiman, 2008). The story of the Sodomites who wanted to gang rape two strangers to whom Lot had offered shelter is another case in point. From a contemporary Western view, we might misleadingly believe that the major moral issue at stake here is rape or homosexuality, but hospitality was an essential moral duty at that time and remains so in many cultures. For Lot, this duty was so serious that he offered the raging mob his virgin daughters if they left his guests alone (Neiman, 2008). Similarly, in modern Europe, wasting energy, eating meat, or smoking in the presence of others were long seen as purely self-regarding decisions. However, environmental protection groups, vegetarians, and anti-smoking groups have reinterpreted these as moral infractions that cause environmental pollution, killing of animals, and lung

cancer through second-hand smoking. I refer to the line that divides personal taste and moral concerns as the *moral rim*. The location of the moral rim describes whether a behavior is included in the moral domain. My hypothesis is that wherever the rim is drawn, the underlying heuristic is likely to remain the same.

As an example, consider renewable energy, environmental protection, and ''green'' electricity. For some, these are deeply moral issues that will determine the living conditions of our great-grand children; for others, these are merely matters of personal preference. Yet the default heuristic appears to guide behavior on both sides of the moral rim. A natural experiment in the German town Schönau showed that when green electricity was introduced as a default, almost all citizens went with the default, even though nearly half had strongly opposed its introduction. In contrast, in towns where ''gray'' energy was the default, only about 1% opted for green energy, a pattern replicated in laboratory experiments (Pichert & Katsikopoulos, 2008). As in the case of organ donation (Johnson & Goldstein, 2003), the willingness to go with the default was stronger in the natural world than in the hypothetical laboratory situation. The same heuristic also seems to be used when drivers decide on which insurance policy to buy, which is rarely considered a moral issue. The states of Pennsylvania and New Jersey offer drivers the choice between an insurance policy with unrestricted right to sue and a cheaper one with suit restrictions (Johnson, Hershey, Meszaros, & Kunreuther, 1993). The unrestricted policy is the default in Pennsylvania, whereas the restricted one is the default in New Jersey. If drivers based their decision on preferences concerning the right to sue, one would expect them to ignore the default settings. If many instead followed the default rule, one would expect more drivers to buy the expensive policy in Pennsylvania. Indeed, only 30% of the New Jersey drivers bought the expensive policy, whereas 79% of the Pennsylvania drivers did so. Many people avoid making a deviating decision, be it on money, life, or death. What we do not know is whether those who rely on the default heuristic for moral issues are the same persons who rely on it for other decisions.

The hypothesis that humans do not have a special moral grammar but that the same social strategies guide moral and nonmoral behavior appears to be consistent with neuroscientific studies that failed to find a specific moral area in the brain or a specific moral activation pattern in several areas. Rather, the same network of activation that is discussed for moral decisions (e.g., Greene & Haidt, 2002) is also typical for social decisions without moral content (Amodio & Frith, 2006; Saxe, 2006).

In sum, I argue that the heuristics in the adaptive toolbox can be used for both moral and nonmoral behavior. This is why I do not attach the qualifier ''moral'' to heuristics (but see Sunstein, 2005). A behavior can be judged as moral or personal, depending on where a culture draws the moral rim. The underlying heuristics are likely the same.

7. The study of moral satisficing

The interpretation of moral behavior as a form of bounded rationality leads to three research questions:

1. Which heuristics underlie moral behavior?
2. What are the social (including legal) environments that, together with the heuristics, produce moral behavior?
3. How can we design environments so that people can reach moral goals more quickly and easily?

I can only provide a sketch of an answer to each of these questions.

7.1. Which heuristics underlie moral behavior?

One obvious answer would be ''don't kill,'' ''don't lie,'' and so on. In my view, this is the wrong direction, as argued in Hypothesis 2. We should not confuse present-day Christian humanist values with heuristics that guide moral behavior. Certain forms of killing, for instance, are legal in countries with capital punishment, and morally acceptable in religious communities, for instance, when a father is expected to kill his daughter if her conduct is considered morally repulsive, such as having sex before marriage (Ali, 2002). We might get a pointer to a better answer when we first ask about the original function of morality (not necessarily the only one in modern societies). Darwin (1871/1981), who thought that a combination of social instincts plus sufficient intellectual powers leads to the evolution of moral sense, proposed the coherence or coordination of human groups:

> There can be no doubt that a tribe including many members who, from possessing in a high degree the spirit of patriotism, fidelity, obedience, courage, and sympathy, were always ready to give aid to each other and to sacrifice themselves for the common good, would be victorious over most other tribes; and this would be natural selection. At all times throughout the world tribes have supplanted other tribes; and as morality is one element in their success, the standard of morality and the number of well-endowed men will thus everywhere tend to rise and increase. (p. 166)

> Selfish and contentious people will not cohere, and without coherence nothing can be effected. (p. 162)

If Darwin's assumption that one original function of morality was the coherence of groups is correct, then the heuristics underlying moral behavior should include those that can provide this function. The default heuristic and imitate-your-peers are apt examples: They can foster social coherence, whatever the default is or whatever the majority does. Note that this opens up a different understanding of what might be the nature of potential universals underlying moral behavior. For instance, Hauser (2006) argued that there is a universal moral grammar with ''hardwired'' principles such as: do as you would be done by; don't kill; don't cheat, steal, or lie; avoid adultery and incest; and care for children and the weak. In his critique, Pippin (2009) responded that these values may be ours but not those of other cultures and times: children sold into slavery by parents who feel entitled to do so; guilt-free spousal abuse by men who see it as their right; moral sanctioning of

pregnant unmarried women by humiliation or driving them into suicide; and so forth. A theory of moral behavior should avoid a Christian humanist bias. Darwin (1981, p. 73) captured this point long ago:

> If for instance, to take an extreme case, men were reared under precisely the same conditions as hive-bees, there can hardly be a doubt that our unmarried females would, like the worker-bees, think it a sacred duty to kill their brothers, and mothers would strive to kill their fertile daughters; and no one would think of interfering.

Terrorists, the Mafia, and crack-dealing gangs run on moral principles (e.g., Gambetta, 1996). For his film *Suicide Killers*, filmmaker Pierre Rehow interviewed would-be terrorists who survived because their bombs failed to explode. He relates: ''Every single one of them tried to convince me that it was the right thing to do for moralistic reasons'' (cited in Neiman, 2008, p. 87). Social psychologists have documented in our own cultures how a situation can stimulate evil behavior in ordinary people, and how easily physical abuse of others can be elicited (e.g., Burger, 2009; Zimbardo, 2007). I suggest that the heuristics underlying moral behavior are not the mirror images of the Ten Commandments and their modern equivalents, but embody more general principles that coordinate human groups.

Consider the following four heuristics as a starting point. Each guides both actions evaluated as moral or immoral and has the potential to coordinate groups; their success (the degree to which they reach a moral goal) depends on the structure of the environment.

1. *Imitate-your-peers:* Do what the majority of your peers do.

Unlike the default heuristic, which needs a default to be elicited, imitation can coordinate behavior in a wide range of situations. No species is known in which children and adults imitate the behavior of others as generally and precisely as *Homo sapiens*. Tomasello (2000) argued that the slavishness of imitation led to our remarkable culture. Imitation enables us to accumulate what our ancestors have learned, thus replacing slow Darwinian evolutionary learning by a Lamarckian form of cultural inheritance. Imitating the majority virtually guarantees social acceptance in one's peer group and fosters shared community values. For instance, the philosopher Otto Weininger (1903) argued that many men desire a woman not because of her features, but because their peers also desire her. Imitation can steer both good and bad moral action, from donating to charity to discriminating against minorities. Those who refuse to imitate the behavior and the values of their culture are likely to be called a coward or oddball if male, or a shame or dishonor to the family if female. A variant of this heuristic is *imitate the successful*, where the object is no longer the majority but an outstanding individual.

2. *Equality heuristic (1/N):* To distribute a resource, divide it equally.

The principle of allocating resources equally is used in self-regarding decisions, such as financial investment (as mentioned before), as well as in morally relevant decisions. For instance, many parents try to divide their love, time, and attention among their children equally to generate a sense of fairness and justice. As a just, transparent distribution principle, it can foster the coherence of a family or a larger group. Similar to the case of organ donation, the equality heuristic does not directly translate into a corresponding

behavior; rather, the result depends on the environment and can even generate systematic inequality. For instance, parents that try to divide their time every day between their N children equally will attain the long-term goal of providing each child with as much time as the other if they have only two children. But if there are three or more children (excepting multiple births), the goal will be missed, because the first-born and the last-born will end up receiving more time than the middle-borns (Hertwig, Davis, & Sulloway, 2002). This result illustrates again that a heuristic (divide equally) and its goal (all children should be given the same amount of time during childhood) is not the same—the environment has the last word.

3. *Tit-for-tat:* If you interact with another person and have the choice between being kind (cooperate) or nasty (defect), then: (a) be kind in the first encounter, there-after (b) keep a memory of size one, and (c) imitate your partner's last behavior (kind or nasty).

''Keep a memory of size one'' means that only the last behavior (kind or nasty) is imi-tated; all previous ones are ignored or forgotten, which can help to stabilize a relationship. Tit-for-tat can coordinate the behavior in a group in the sense that all actors will end up cooperating but are simultaneously protected against potential defectors. As with imitate-your-peers and the default heuristic, tit-for-tat illustrates that the same heuristic can lead to opposite behaviors, here kind or nasty, depending on the social environment. If a husband and wife both cooperate when engaging in their first interaction and subsequently always imitate the other's behavior, the result can be a long harmonious relationship. If, however, she relies on tit-for-tat but he on the maxim ''always be nasty to your wife, so that she knows who is the boss,'' her initially kind behavior will turn to being nasty to him as well. Behavior is not a mirror of a trait of being kind or nasty, but results from an interaction between mind and environment. An explanation of the tit-for-tat players' behavior in terms of traits or attitudes would miss this crucial difference between process (tit-for-tat) and resulting behavior (cooperate or not).

4. *Default heuristic:* If there is a default, do nothing about it (see above).

Whereas equality is a simple answer to the problem of allocating resources fairly among N alternatives, the default heuristic addresses the problem of which of N alternatives to pur-sue when one of these is marked as the default. It has the potential to create coherence in behavior even when no social obligation or prohibition exists.

These four heuristics are selected examples. Several others have been studied (see Cosmides & Tooby, 2008; Gigerenzer, 2007, 2008b; Haidt, 2001), but we do not have anything near a complete list. The general point is that moral satisficing assumes not one general calculus but several heuristics. This fact is reflected in the term ''adaptive toolbox,'' where the qualifier *adaptive* refers to the evidence that heuristic tools tend to be selected, consciously or unconsciously, to match the task at hand (e.g., Bröder, 2003; Dieckmann & Rieskamp, 2007; Mata, Schooler, & Rieskamp, 2007). Acting in a morally responsible way can thus be reinterpreted as the ability to choose a proper heuristic for a given situation. The question whether there are one or several processes underlying morality is an old one. Adam Smith (1761), for instance, criticized Hutchesonn's and Hume's theory of moral sense as a single feeling of moral approval. For him, the sense of virtue was distinct from the sense of

propriety, of merit, or of duty, which is why Smith spoke of ''moral sentiments'' in the plural.

7.1.1. Building blocks and core capacities

Heuristics can be composed from several building blocks, which enable new heuristics to be generated by recombination and modification. For instance, one weak point of tit-for-tat is that a single negative behavior (being nasty) can bring two people who play tit-for-tat into an endless cycle of violence or two social groups into a vendetta, each act of violence being justified as a fair response to the other's most recent attack. A modification of the second building block to a memory size of two can resolve this problem. This heuristic is called *tit-for-two-tats*, where a person turns nasty only if the partner behaved nasty twice in a row. Heuristics and their building blocks are based on evolved and learned core capacities, such as recognition of individuals, inhibitory control, and the ability to imitate (Stevens & Hauser, 2004). Like the default heuristic and imitate-your-peers, tit-for-tat appears to be rare among other animals, unless they are genetically related (Hammerstein, 2003).

7.2. What social environments, together with heuristics, guide moral behavior?

Structural features of the environment can interact with the mind in two ways. First, the presence or absence of a feature increases or limits the choice set of applicable heuristics. For instance, if no alternative is marked as the default, the default heuristic cannot be triggered; if there are no peers present whose behavior in a new task can be observed, the imitate-your-peers heuristic cannot be activated. Second, if more than one heuristic remains in the choice set, features of the environment can determine which heuristic is more likely to be relied on. Research on decision making has shown that people tend to select heuristics in an adaptive way (Payne, Bettman, & Johnson, 1993), and this selection process has been formalized as a reinforcement learning process in strategy selection theory (Rieskamp & Otto, 2006). Environmental features investigated in this research include time for decision, payoffs, and redundancy of cues.

One aspect of the environment is the structure of social relations that humans are born into or put into in an experimental situation. Fiske (1992) distinguished four kinds of relationships among which people move back and forth during their daily activities. On the basis of these, we might ask what social environments are likely to trigger the equality heuristic. According to Fiske's classification, these environments consist of *equality matching relations*, where people keep track of the balance of favors and know what would be required to restore the balance. Examples are turn-taking in babysitting co-ops and voting systems in democracies that allocate each adult one vote. Equality is a simpler rule for fair division than equity; the latter divides a cake among *N* persons according to some measure of effort, time, or input of each individual (Deutsch, 1975; Messick, 1993). In Fiske's taxonomy, the environments in which distribution according to equity can be expected comprise *market pricing relations*, that is, social relations that are structured by some kind of cost–benefit analysis, as in business relations. Milgram's

experiments implemented the third kind of relation, an *authority ranking relation*, where people have asymmetric relations in a hierarchy, subordinates react with respect and deference, and superiors take pastoral responsibility for them. Note that in this experimentally induced authority relation, the absence of monetary incentives—participants were paid independent of whether or not they applied shocks to a stranger—appeared to play little or no role (see above). Authority ranking relations tend to trigger heuristics of the kind: *If a person is an authority, follow requests.* It appears that not even a true authority relation is needed, but that mere signs of such a relation—for instance, a white coat—can trigger the heuristic and the resulting moral behavior (Brase & Richmond, 2004). The fourth relation in Fiske's taxonomy is *community sharing*, where people treat some group as equivalent or undifferentiated, as when sharing a commons.

That is not to say that triggering is a one-to-one process; conflicts can arise. For instance, in one condition of the obedience experiment, a confederate participant was introduced who sat next to the real participant and refused to continue the experiment after pressing the 90-V switch and hearing the learner's groan (Burger, 2009). This situation might trigger both *imitate-your-peers* and *if a person is an authority, follow requests*, which are conflicting behaviors. The majority of participants (63%) followed the authority and went on to give shocks of higher intensity, compared to 70% of those who did so without seeing someone decline.

These are single examples, but not a systematic theory of the structure of social environments relevant to moral behavior. Such a theory could be constructed by combining the research on heuristics with that on social structures (e.g., Boyd & Richerson, 2005; Fiske, 1992; Haidt & Bjorklund, 2008; Shweder, Much, Mahaptra, & Park, 1997).

The ecological view of morality has methodological consequences for the experimental study of morality.

1. *Study social groups in addition to isolated individuals.* If moral behavior is guided by social heuristics, these can hardly be detected in typical psychological experiments where individuals are studied in isolation. Heuristics such as imitate-your-peers and tit-for-tat can only unfold in the presence of peers.

2. *Study moral behavior in natural environments in addition to in hypothetical problems.* Hypothetical situations such as the trolley problems eliminate characteristic features of natural environments, such as the uncertainty about the full set of possible actions and their consequences. It needs to be addressed whether the results obtained from hypothetical small worlds and isolated individuals generalize to moral behavior outside the laboratory.

3. *Analyze moral behavior in addition to verbal reports.* Given that people are often unaware of the heuristics and environmental structures that guide their moral behavior, paper-and-pencil tasks and self-reports alone are insufficient research methods (Baumeister, Vohs, & Funder, 2007). This methodological point is consistent with the observation that people typically cannot explain why they feel that something is morally right or wrong, or why they did what they did (Haidt & Bjorklund, 2008).

These methodological consequences are exactly the same ones I recommend for studying decision making outside the moral domain, consistent with Hypothesis 2 above.

7.3. How can environments be designed so that people can better reach moral goals?

To begin with, we need public awareness that the causes of moral behavior are not simply inside the mind and that understanding the interplay between minds and environments is a useful starting point for moral policy. Note that this is not paternalistic as long as one helps, not forces, people to reach their own goals. Changing the law from opt-in to opt-out is an example of how to better reach the goal of increasing the number of potential organ donors and thus reducing the number of people who die for want of a donor. Abadie and Gay (2006) estimated that, once other factors influencing organ donation are taken into account (such as a country's rate of motor vehicle accidents involving fatalities, a main source of donors), actual donation rates are 25–30% higher on average in presumed-consent countries. However, as mentioned before, to this day various governmental agencies tend to bet on the internal causes and continue to proclaim the importance of raising public awareness and of disseminating more information. This program is based on an inadequate psychological theory and will likely fail in the future as it has in the past. Let me end with an illustration for the design of environments: crime in communities.

How should one deal with community crime, from littering to vandalism? Moral satisficing suggests considering potential candidate heuristics such as imitate-your-peers and then changing the environment to diminish the likelihood that it triggers the heuristic. In other words, as long as deviant behavior is publicly observed, this will trigger further deviant behavior. The program of ''fixing broken windows'' (Kelling & Coles, 1996) follows this line of thought, by repairing windows and fixing streetlights, cleaning sidewalks and subway stations, and so on. Supplemented by a zero-tolerance program, this change in the environment substantially decreased petty crime and low-level antisocial behavior in New York City and other places, and it may have even reduced major crime.

8. Satisficing in moral philosophy

In this article, I invited you to consider the question: How can we understand moral behavior if we look at it from the perspective of bounded rationality? My answer is summarized in Propositions 1–5 in the introduction. I will end by briefly comparing this perspective with the views of several moral philosophers in *Satisficing and Maximizing* (Byron, 2004). The key difference is that the two principles on which my essay is based, less-can-be-more and Simon's scissors, are absent in this interesting collection of essays. The common assumption is that an optimizing process leads to the optimal outcome and satisficing to a second-best outcome, because process and outcome are not distinguished (but see Proposition 2). Thus, less-can-be-more is not part of the picture. As a result, some philosophers make (misleading) normative statements such as: ''satisficing may be a conceptual tool that fits the facts, but it is not good enough. We can do better'' (Rich-

ardson, 2004, p. 127). The absence of the ecological dimension makes Propositions 2, 4, and 5 nonissues.

Although Herbert Simon is repeatedly invoked as the source of inspiration, some essays still define satisficing as a form of optimizing. As mentioned earlier, one common (mis)interpretation of Simon's concept of bounded rationality is that it is nothing other than optimization under constraints (e.g., Narveson, 2004, p. 62), overlooking the fact that optimization is typically impossible (e.g., computationally intractable) in large worlds (Proposition 1). A second more interesting interpretation starts from the observation that people have multiple goals. Here, satisficing means that people choose local optima for some goals (e.g., to spend less time than desired with family for one's career) but still seek the global optimum for their life as a whole (Schmidtz, 2004). This interpretation, in contradiction to Proposition 1, also assumes that optimization is always possible. An original third interpretation (which does not reduce satisficing to optimizing) is that satisficing means pursuing moderate goals, and that moderation can be a virtue and maximization a vice, the latter leading to greed, perfectionism, and the decline of spontaneity (Slote, 2004; Swanton, 2004). There is a related claim in the psychological literature that satisficers are more optimistic and satisfied with life, while maximizers excel in depression, perfectionism, and regret (Schwartz et al., 2002).

A further proposal (not covered in this volume) is various forms of rule-consequentialism (Braybrooke, 2004). Because it is typically impossible to anticipate all consequences of each possible action and their probabilities in every single case, rule consequentialists emphasize the importance of rules, which do not maximize utility in every single case but—if they are followed—do so in the long run. Once again, the idea of maximization is retained, although the object of maximization is not directly the action but instead a rule that generally leads to the best actions. The similarity between moral satisficing and rule-consequentialism is in the focus on rules, whereas moral satisficing assumes that rules are typically unconscious social heuristics that are elicited by the structure of the environment (see also Haidt, 2001).

This broad range of interpretations can be taken to signal the multiple ways in which bounded rationality can inspire us to rethink morality. But it also indicates how deeply entrenched the notion of maximization is, and how alien Simon's scissors remain to many of us.

9. Ecological morality

Herbert Simon (1996, p. 110) once said, ''Human beings, viewed as behaving systems, are quite simple. The apparent complexity of our behavior over time is largely a reflection of the complexity of the environment in which we find ourselves.'' An ecological view of moral behavior might be easier to accept for those who have experienced that their good and bad deeds are not always the product of deliberation but also of the social environment in which they live, for better or worse. Yet others will see too much ''moral luck'' in this vision, insisting that the individual alone is or should be responsible for what he or she does. Luck, I believe, is as real as virtue. It is the milieu in which we happen to grow up but also the environment we actively create for our children and students.

Acknowledgments

References

Abadie, A., & Gay, S. (2006). The impact of presumed consent legislation on cadaveric organ donation: A cross-country study. *Journal of Health Economics*, *25*, 599–620.

Adorno, T. W., Frenkel-Brunswik, E., Levinson, D., & Sanford, N. (1950). *The authoritarian personality*. New York: Harper & Row.

Ali, A. H. (2002). *The caged virgin: A Muslim woman's cry for reason*. London: Simon & Schuster.

Amodio, D. M., & Frith, C. D. (2006). Meeting of minds: The medial prefrontal cortex and social cognition. *Nature Review Neuroscience*, *7*, 268–277.

Arrow, K. J. (2004). Is bounded rationality unboundedly rational? Some ruminations. In M. Augier & J. G. March (Eds.), *Models of a man: Essays in memory of Herbert A. Simon* (pp. 47–55). Cambridge, MA: MIT Press.

Baumeister, R. F., Vohs, K. D., & Funder, D. C. (2007). Psychology as the science of self-reports and finger movements. Whatever happened to actual behavior? *Perspectives on Psychological Science*, *2*, 396–403.

Becker, G. S. (1995). *The essence of Becker*. Stanford, CA: Hoover Institution Press.

Bennis, W. M., Medin, D. L., & Bartels, D. M. (2010). The costs and benefits of calculation and moral rules. *Perspectives on Psychological Science*, *5*, 187–202.

Binmore, K. (2009). *Rational decisions*. Princeton, NJ: Princeton University Press.

Blass, T. (1991). Understanding behavior in the Milgram obedience experiment: The role of personality, situations, and their interaction. *Journal of Personality and Social Psychology*, *60*, 398–413.

Boyd, R., & Richerson, P. J. (2005). *The origin and evolution of cultures*. New York: Oxford University Press.

Brase, G. L., & Richmond, J. (2004). The white-coat effect: Physician attire and perceived authority, friendliness, and attractiveness. *Journal of Applied Social Psychology*, *34*, 2469–2481.

Braybrooke, D. (2004). *Utilitarianism: Restorations; repairs; renovations*. Toronto, Canada: University of Toronto Press.

Brighton, H. (2006). Robust inference with simple cognitive models. In C. Lebiere & R. Wray (Eds.), *AAAI spring symposium: Cognitive science principles meet AI-hard problems* (pp. 17–22). Menlo Park, CA: American Association for Artificial Intelligence.

Bröder, A. (2003). Decision making with the ''adaptive toolbox'': Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology*, *29*, 611–625.

Bundesärztekammer. (2008). *EU Bulletin, Nr. 9*, April 21, 2008.

Burger, J. M. (2009). Replicating Milgram: Would people still obey today? *American Psychologist*, *64*, 1–11.

Byron, M. (Ed.) (2004). *Satisficing and maximizing: Moral theorists on practical reason*. Cambridge, England: Cambridge University Press.

Cokely, E. T., & Feltz, A. (2009). Adaptive variation in judgment and philosophical intuition. *Consciousness and Cognition*, *18*, 355–357.

Commission of the European Communities. (2007, May 30). *Organ donation and transplantation: Policy actions at EU level. Impact assessment*. (Press Release No. IP/07/718). Available at: http://europa.eu/rapid/pressReleasesAction.do?reference=IP/07/718&format=HTML&aged=1&language=EN&guiLanguage=en

Cosmides, L., & Tooby, J. (2008). Can a general deontic logic capture the facts of humans' moral reasoning? How the mind interprets social exchange rules and detects cheaters. In W. Sinnott-Armstrong (Ed.), *Moral*

psychology: Vol 1. The evolution of morality: Adaptations and innateness (pp. 53–119). Cambridge, MA: MIT Press.

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 97–118). New York: Oxford University Press.

Darwin, C. (1981). *The descent of man*. Princeton, NJ: Princeton University Press. (Original work published 1871).

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

Dawkins, R. (2006). *The God delusion*. Boston, MA: Houghton Mifflin Harcourt.

DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, *22*, 1915–1953.

Deutsch, M. (1975). Equity, equality, and need: What determines which value will be used as the basis of distributive justice? *Journal of Social Issues*, *31*, 137–149.

Dhami, M. K. (2003). Psychological models of professional decision-making. *Psychological Science*, *14*, 175–180.

Dieckmann, A., & Rieskamp, J. (2007). The influence of information redundancy on probabilistic inference. *Memory & Cognition*, *35*, 1801–1813.

Doris, J. M. (2002). *Lack of character*. New York: Cambridge University Press.

Feltz, A., & Cokely, E. T. (2009). Do judgments about freedom and responsibility depend on who you are? Personality differences in intuitions about compatibilism and incompatibilism. *Consciousness and Cognition*, *18*, 342–350.

Fiske, A. P. (1992). Four elementary forms of sociality: Framework for a unified theory of sociality. *Psychological Review*, *99*, 689–723.

Funder, D. (2001). Personality. *Annual Review of Psychology*, *52*, 197–221.

Gambetta, D. (1996). *The Sicilian Mafia. The business of private protection*. Cambridge, MA: Harvard University Press.

Garcia-Retamero, R., & Dhami, M. K. (2009). Take-the-best in expert–novice decision strategies for residential burglary. *Psychonomic Bulletin & Review*, *16*, 163–169.

Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, *4*, 1–58.

Gerson, R., & Damon, W. (1978). Moral understanding and children's conduct. In W. Damon (Ed.), *New directions for child development* (pp. 41–59). San Francisco, CA: Jossey-Blass.

Gigerenzer, G. (1991). From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological Review*, *98*, 254–267.

Gigerenzer, G. (2006). Heuristics. In G. Gigerenzer & C. Engel (Eds.), *Heuristics and the law* (pp. 17–44). Cambridge, MA: MIT Press.

Gigerenzer, G. (2007). *Gut feelings: The intelligence of the unconscious*. New York: Viking. (UK version: London: Allen Lane/Penguin).

Gigerenzer, G. (2008a). *Rationality for mortals*. New York: Oxford University Press.

Gigerenzer, G. (2008b). Moral intuition = Fast and frugal heuristics? In W. Sinnott-Armstrong (Ed.), *Moral psychology: Vol 2. The cognitive science of morality: Intuition and diversity* (pp. 1–26). Cambridge, MA: MIT Press.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, *1*, 107–143.

Gigerenzer, G., & Selten, R. (Eds.) (2001a). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.

Gigerenzer, G., & Selten, R. (2001b). Rethinking rationality. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality. The adaptive toolbox* (pp. 1–12). Cambridge, MA: MIT Press.

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Greene, J., & Haidt, J. (2002). How (and where) does moral judgment work? *TRENDS in Cognitive Sciences*, *6*, 517–523.

Gruber, H. E., & Vonèche, J. J. (1977). *The essential Piaget*. New York: Basic Books.

Hacking, I. (1975). *The emergence of probability*. Cambridge, England: Cambridge University Press.

Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, *108*, 814–834.

Haidt, J., & Bjorklund, F. (2008). Social intuitionists answer six questions about moral psychology. In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Vol. 2: The cognitive science of morality: Intuition and diversity* (pp. 181–217). Cambridge, MA: MIT Press.

Hammerstein, P. (2003). Why is reciprocity so rare in social animals? A protestant appeal. In P. Hammerstein (Ed.), *Genetic and cultural evolution of cooperation* (pp. 83–93). Cambridge, MA: MIT Press.

Hauser, M. (2006). *Moral minds: How nature designed our universal sense of right and wrong*. New York: Ecco.

Hertwig, R., Davis, J. N., & Sulloway, F. (2002). Parental investment: How an equity motive can produce inequality. *Psychological Bulletin*, *128*, 728–745.

Hutchinson, J. M. C., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural Processes*, *69*, 97–124.

Johnson, E. J., & Goldstein, D. G. (2003). Do defaults save lives? *Science*, *302*, 1338–1339.

Johnson, E. J., Hershey, J., Meszaros, J., & Kunreuther, H. (1993). Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty*, *7*, 35–51.

Kahneman, D. (2003). A perspective on judgement and choice: Mapping bounded rationality. *American Psychologist*, *58*, 697–720.

Kelling, G. L., & Coles, C. M. (1996). *Fixing broken windows: Restoring order and reducing crime in our communities*. New York: The Free Press.

Knobe, J., & Nichols, S. (2008). An experimental philosophy manifesto. In J. Knobe & S. Nichols (Eds.), *Experimental philosophy* (pp. 3–14). New York: Oxford University Press.

Kohlberg, L. (1968). The child as a moral philosopher. *Psychology Today*, *2*, 25–30.

Lipsey, R. G. (1956). The general theory of the second best. *Review of Economic Studies*, *24*, 11–32.

Makridakis, S., & Hibon, M. (2000). The M3-competition: Results, conclusions, and implications. *International Journal of Forecasting*, *16*, 451–476.

Mata, R., Schooler, L. J., & Rieskamp, J. (2007). The aging decision maker: Cognitive aging and the adaptive selection of decision strategies. *Psychology and Aging*, *22*, 796–810.

Matheson, D. (2006). Bounded rationality, epistemic externalism and the Enlightenment picture of cognitive virtue. In R. Stainton (Ed.), *Contemporary debates in cognitive science* (pp. 134–144). Oxford, England: Blackwell.

Messick, D. M. (1993). Equality as a decision heuristic. In B. A. Mellers & J. Baron (Eds.), *Psychological perspectives on justice* (pp. 11–31). New York: Cambridge University Press.

Milgram, S. (1974). *Obedience to authority: An experimental view*. New York: Harper & Row.

Mischel, W. (1968). *Personality and assessment*. New York: Wiley.

Nagel, T. (1993). Moral luck. In D. Statman (Ed.), *Moral luck* (pp. 57–71). Albany, NY: State University of New York Press.

Narvaez, D., & Lapsley, D. (2005). The psychological foundations of everyday morality and moral expertise. In D. Lapsley & C. Power (Eds.), *Character psychology and character education* (pp. 140–165). Notre Dame, IN: University of Notre Dame Press.

Narveson, J. (2004). Maximizing life on a budget; or, if you would maximize, then satisfice! In M. Byron (Ed.), *Satisficing and maximizing: Moral theorists on practical reason* (pp. 59–70). Cambridge, England: Cambridge University Press.

Neiman, S. (2008). *Moral clarity: A guide for grownup idealists*. New York: Harcourt.

Pascal, B. (1962). *Pensées*. Paris: Editions du Seuil. (Original work published 1669).

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge, England: Cambridge University Press.

Persijn, G. (1997). Public education and organ donation. *Transplantation Proceedings*, *29*, 1614–1617.

Pichert, D., & Katsikopoulos, K. V. (2008). Green defaults: Information presentation and pro-environmental behavior. *Journal of Environmental Psychology*, *28*, 63–73.

Pippin, R. B. (2009). Natural & normative. *Daedalus, Summer 2009*, 35–43.

Richardson, H. S. (2004). Satisficing: Not good enough. In M. Byron (Ed.), *Satisficing and maximizing: Moral theorists on practical reason* (pp. 106–130). Cambridge, England: Cambridge University Press.

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*, 207–236.

Rosenbaum, J. (2009). Patient teenagers? A comparison of the sexual behavior of virginity pledgers and matched nonpledgers. *Pediatrics*, *123*, 110–120. doi: 10.1542/peds.2008-0407.

Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.

Saxe, R. (2006). Uniquely human social cognition. *Current Opinion in Neurobiology*, *16*, 235–239.

Schmidtz, D. (2004). Satisficing as a humanly rational strategy. In M. Byron (Ed.), *Satisficing and maximizing: Moral theorists on practical reason* (pp. 30–58). Cambridge, England: Cambridge University Press.

Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., & Lehman, D. R. (2002). Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology*, *83*, 1178–1197.

Shaffer, D. M., Krauchunas, S. M., Eddy, M., & McBeath, M. K. (2004). How dogs navigate to catch Frisbees. *Psychological Science*, *15*, 437–441.

Shanteau, J. (1992). How much information does an expert use? Is it relevant? *Acta Psychologica*, *81*, 75–86.

Shweder, R. A., Much, N. C., Mahaptra, M., & Park, L. (1997). The ''big three'' of morality (autonomy, community, and divinity), and the ''big three'' explanations of suffering, as well. In A. Brandt & P. Rozin (Eds.), *Morality and health* (pp. 119–169). New York: Routledge.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, *41*, 1–19.

Simon, H. A. (1996). *The sciences of the artificial* (3rd ed.). Cambridge, MA: MIT Press. (Originally published 1969).

Slote, M. (2004). Two views of satisficing. In M. Byron (Ed.), *Satisficing and maximizing: Moral theorists on practical reason* (pp. 14–29). Cambridge, England: Cambridge University Press.

Smith, A. (1761). *The theory of moral sentiments*. London: Millar.

Statman, D. (1993). *Moral luck*. Albany, NY: State University of New York Press.

Stevens, J. R., & Hauser, M. D. (2004). Why be nice? Psychological constraints on the evolution of cooperation *TRENDS in Cognitive Sciences*, *8*, 60–65.

Sunstein, C. R. (2005). Moral heuristics. *Behavioral and Brain Sciences*, *28*, 531–542.

Swanton, C. (2004). Satisficing and perfectionism in virtue ethics. In M. Byron (Ed.), *Satisficing and maximizing: Moral theorists on practical reason* (pp. 176–189). Cambridge, England: Cambridge University Press.

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven, CT: Yale University Press.

Todd, P. M., & Gigerenzer, G. (2001). Shepard's mirrors or Simon's scissors? *Behavioral and Brain Sciences*, *24*, 704–705.

Tomasello, M. (2000). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.

Weininger, O. (1903). *Geschlecht und Charakter [Sex and character]*. Vienna: Wilhelm Braumüller.

Williams, B. (1981). *Moral luck*. Cambridge, England: Cambridge University Press.

Wilson, D. S. (2002). *Darwin's cathedral: Evolution, religion, and the nature of society*. Chicago, IL: University of Chicago Press.

Wübben, M., & Wangenheim, F. v. (2008). Instant customer base analysis: Managerial heuristics often ''get it right.'' *Journal of Marketing*, *72*, 82–93.

Zimbardo, P. (2007). *The Lucifer effect: Understanding how good people turn to evil*. New York: Random House.