Ecological Rationality: Intelligence in the World
Peter M. Todd and Gerd Gigerenzer

Print publication date: 2012
Print ISBN-13: 9780195315448
Published to Oxford Scholarship Online: May-12
DOI: 10.1093/acprof:oso/9780195315448.001.0001

Designed to Fit Minds

Will M. Bennis, Konstantinos V. Katsikopoulos, Daniel G. Goldstein, Anja Dieckmann, Nathan Berg

DOI: 10.1093/acprof:oso/9780195315448.003.0126

Abstract and Keywords

Economic models of benefit–cost analysis assume decision makers choose so as to maximize net benefits given stable internal preferences. But institutional structure can play a central role in determining whether or not an agent's decisions promote their individual preferences. This chapter explores four cases where the interaction between institutional structure and non-optimizing human decision processes does a better job than optimizing models at explaining choice. These cases suggest that institutions are often designed to fit minds. In some cases these designs rely on existing heuristics and their building blocks (organ donation rules, slot-machine design), while in other cases institutions design new lexicographic heuristics to help make decisions fast and unambiguous (driving right-of-way rules, World Cup soccer team ranking rules).

Keywords: benefit−cost analysis, default heuristic, institutional design, lexicographic strategy, organ donation, preference stability, traffic right-of-way, slot machines, ranking rules

> Reform the environment, stop trying to reform the people. They will reform themselves if the environment is right.
>
> *Buckminster Fuller*

Only about 12% of Germans have given legal consent to donate their organs when they die. In contrast, in the neighboring country of Austria more than 99% are potential donors. To explain this large difference in consent rates for organ donation, social scientists using the standard decision-making model in economics have looked to differences in expected benefits and costs while controlling for income, education, and religion (Gimbel, Strosberg, Lehrman, Gefenas, & Taft, 2003). Regression models based on the benefit–cost theory, however, show little evidence that large differences in actual organ-donor consent rates are statistically or causally linked to perceived benefits and costs. Critics of the economic model have attempted to explain cross-country behavioral differences in terms of culture, social norms, and history. But the mostly small differences between Austria and Germany on these dimensions seem unlikely candidates for explaining the large gap in their donor consent rates.

Johnson and Goldstein (2003) did, however, identify an important *institutional* difference between Austria and Germany that seems to explain differential consent rates much better than economic, sociological, and historical approaches: different defaults written into law regarding organ donation consent status. In *presumed* ***(p. 410)*** *consent* countries such as Austria, individuals are from birth considered to be potential organ donors, which means there is effective legal consent for their organs to be harvested upon death for transplant to the living. *Explicit consent* countries such as Germany, on the other hand, use the opposite default: No organs can be legally harvested from the dead unless individuals opt in to organ-donor status by giving their explicit consent.

Switching away from either default is not especially costly in terms of time or effort. In Germany, according to current law, one can switch from the nondonor default to donor status by submitting this wish in writing.[1] In Austria, opting out of consent status requires a bit more effort and physical resources, but not much more: submitting an official form to the Austrian Federal Health Institute via post or fax, requiring approximately 5 minutes and perhaps a stamp. The main implication of these small switching costs is that, according to the stable preference assumption of standard economic theory, defaults should not influence behavior. For example, someone who has stable preferences that rank donor over nondonor status—and whose difference in payoffs across these two states more than offsets the cost of switching away from the default—should choose to be an organ donor regardless of how defaults are set. Yet, contrary to economic theory, defaults are strongly correlated with actual consent rates. Figure 16-1 shows consent rates for a range of countries, making clear the large difference in potential organ donation rates between presumed consent countries and explicit consent ones.

Johnson and Goldstein (2003) suggested a simple heuristic model of individual behavior that fits the data in Figure 16-1 much better than rival explanations investigated elsewhere in the literature. Their *default heuristic* consists of the following procedure: *When faced with a choice between options where one of them is a default, follow the default*. This heuristic —in contrast to other explanations—does not rely on inherent differences inside the minds of decision makers in different countries: It predicts distinct behavior on the part of Austrians and Germans because it depends on an institutional variable set to different values in those countries, namely, defaults regarding consent. The heuristic model does not rely on a theory of inherent preferences, and it attributes none of **(p. 411)**

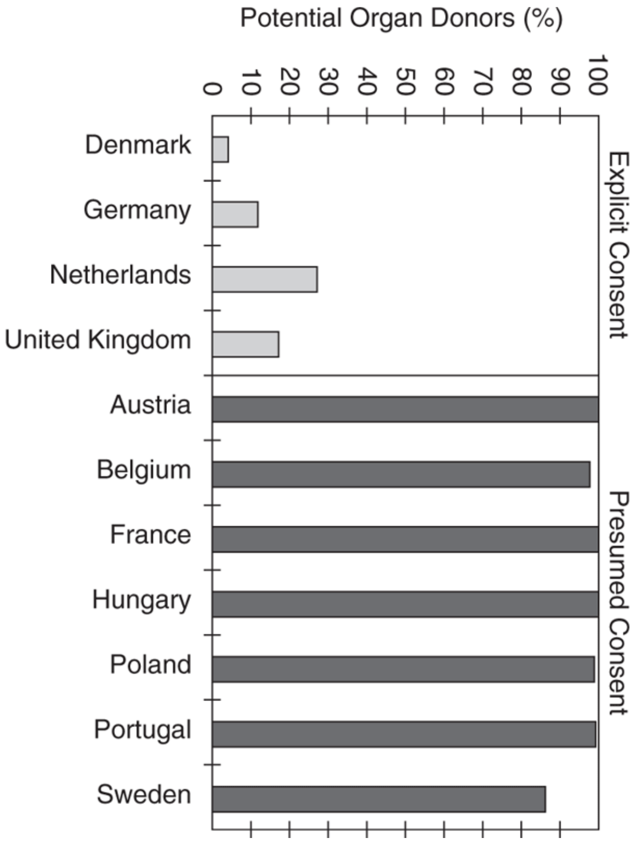

Figure 16-1: Population rates of potential organ donors by country. The first four bars indicate *explicit consent* countries, where individuals are assumed *not* to be organ donors but can take action to opt in to organ donor status. The remaining bars indicate *presumed consent* countries, where the default presumes that individuals are organ donors while allowing them to opt out if they choose. (Adapted from Johnson & Goldstein, 2003.)

the observed differences in behavior to essentialist concepts residing solely within individuals or exclusively outside. In this chapter, we explore cases such as this where ecological rationality can emerge—or be obscured—through interactions between the decision heuristics of individuals and the choice environments they face, which in turn have been structured by institutions with incentives that may or may not match those of the individual. (See chapter 17 for further examples of this interaction in health care.)

The institutional environment structures that shape people's behavior can be surprisingly subtle. To show this in the case of organ donation decisions, Johnson and Goldstein (2003) ran the following experiment. Participants were randomly assigned to two groups. One group saw the following *opt-in* cover story:

> Imagine that you just moved to a new state and must get a new driver's license. As you complete the application, you come across the following. *Please read and respond as you would if you were actually presented this choice today. We are interested in your honest response*: In this state every person is **(p. 412)** considered not to be an organ donor unless they choose to be. You are therefore currently not a potential donor. If this is acceptable, click here. If you wish to change your status, click here.

The second group saw the same message changed to an *opt-out* scenario with the script modified to read: "In this state every person is considered to be an organ donor unless they choose not to be. You are therefore currently a potential donor… ." The default has simply been changed. How much difference will this make for choices between the same two important outcomes?

In this environment constructed in the laboratory, 82% of participants in the opt-out scenario chose to be potential donors, while only 42% in the opt-in scenario did. This large gap between experimental consent rates mirrors the differences between European countries seen in Figure 16-1.

This experiment shows that the small change of adding or removing the word "not" on the organ donation form, thereby changing the default, has a large impact on the aggregate outcome as measured by consent rates. Similarly drawing on heuristic models of behavior, researchers have achieved large changes in aggregate behavior by modifying default settings of institutional parameters in other domains, such as personal savings (Thaler & Benartzi, 2004; Thaler & Sunstein, 2008). Additionally, using two natural experiments and two laboratory studies, Pichert and Katsikopoulos (2008) showed that defaults have a dramatic influence on whether people in Germany subscribe to a "green" electricity provider. On the other hand, large campaigns hoping to increase donation rates by providing information about costs and benefits, but *without* changing defaults, do not seem to work.[2] Such failed attempts to influence the public's behavior implicitly draw on the standard economic model of individual decision making as the rationale for intervention, which assumes that individual decisions result from systematic weighing of costs and benefits and so are best influenced by changing individuals' benefit and cost parameters. Following this economic model, for example, the Netherlands **(p. 413)** undertook a broad educational campaign that included sending out a mass mailing to more than 12 million people asking them to register their organ donation preference. The result: Donation consent rates did not improve (Oz et al., 2003). Consequently, calls are increasing to adopt the simpler and more effective path of following psychology and changing defaults as one way to overhaul ailing health care systems (e.g., in the U.S., as heralded in the *New York Times*—see Rose, 2009) and address other policy issues (Goldstein, Johnson, Herrmann, & Heitmann, 2008).

Heuristics Versus Standard Economic Approaches to Decision Making

In evolutionary game theory, strategies or behavioral rules that yield suboptimal payoffs are usually assumed to die out under competitive pressure from agents using strategies with higher average payoffs. Thus, decision processes such as the default heuristic, which are not derived as solutions to optimization problems, are often considered uninteresting. The logic behind this dismissive attitude is that heuristic behavior is unstable because it is likely to be supplanted by superior decision strategies, and therefore it need not be studied, since one would not expect to observe what is unstable for long. This exclusive focus on stable outcomes in standard economic theory has attracted its share of critics (e.g., Hayek, 1945; Schumpeter, 1942) yet remains a core tenet of economics as it is taught and practiced throughout most of the world.

Those who study heuristics as an alternative to the standard economic model must acknowledge that the viewpoint of economic theory poses a fair question: Why would someone use heuristics? In the case of the default heuristic, it is easy to see that it is well adapted to environments where institutional designers (i.e., those in charge of choosing defaults) have the interests of default users in mind and communicate their recommendations through their choice of available defaults. Of course, this confluence of interests will not always be the case, as in countries such as Germany and the United States, where 70–80% of those surveyed say they want to be an organ donor and yet consent defaults are not set to match this majority preference (Gallup Organization, 1993). Social preferences may also play a role in explaining why people follow defaults, for example, if people perceive social value in matching the action taken by the majority, or if they fear negative social consequences from behaving out of line with the majority (Ariely & Levav, 2000). Defaults may codify social norms or provide a coordination mechanism by which users of the default heuristic successfully wind up in the majority. The default heuristic also **(p. 414)** greatly reduces decision costs of time and deliberation, which are common benefits of fast and frugal decision making (Gigerenzer & Todd, 1999). Finally, the case of organ donation also raises the possibility that deliberating over some choice sets is inherently distasteful, forcing individuals to consider unpleasant contingencies such as one's own death, which may be substantially avoided by ignoring the full choice set and accepting defaults.

In this chapter we take up the theme of institutional design through the lens of ecological rationality instead of standard economic theory. Heuristics are models of individual behavior based on psychological plausibility and ecological effectiveness rather than axioms of logical consistency from economic theory. As the examples in this chapter are intended to show, the study of heuristics allows us to analyze institutions that economic theory would never predict and provides new explanations for the functioning of existing institutions according to institutional objectives, such as simplicity and transparency, that are difficult to motivate using standard informational assumptions of economic theory.

As critics (e.g., Hayek, 1945; Simon, 1955a) and defenders (e.g., Becker, 1978) have both pointed out, neoclassical economics and game theory are based on a well-defined, singular model of human behavior. This benefit–cost model assumes that choice sets are searched exhaustively, alternative choices are scored in terms of benefits and costs, and finally these scores are integrated to determine an optimal action or decision (for foundational examples, see Savage, 1954; von Neumann & Morgenstern, 1947). One key implication of the economic model is that behavior, which is taken to result from the process of optimization just described, should depend systematically on perceived benefits and costs. A second important implication that follows from this is that institutional modifications that leave choice sets and their net benefits unaltered, as do default rules for organ donation consent (apart from the costs of switching away from the default), should have no effect on observed behavior. Similarly, logically equivalent representations of a given set of information should not, according to the economic model, influence behavior (see chapter 17).

But once one considers the possibilities for designing institutions to fit *actual* human minds and the processes they follow rather than fictitious agents pursuing the economic model of optimization, new challenges and new possibilities arise. Some institutions that would not work in a world populated by economic agents work surprisingly well in the real world. For example, economists consider it something of a puzzle why voluntary compliance with income tax laws is so high, and why littering in some very clean public parks is not more of a problem, given that governments invest so little in enforcement. In other cases, institutions that **(p. 415)** assume forward-looking behavior, full information, and costless information processing encounter obvious problems when confronted with the human realities of limited information and cognition, as demonstrated by the case of organ donations and by numerous instances of well-intentioned institutions incorrectly assuming that complete information and unhindered choice is the best way to help people make good decisions (Thaler & Sunstein, 2008). The examples that follow illustrate a range of real-world institutions that one would never expect to be designed in the way that they are if the hypotheses built into the economic model of human behavior were universally valid. Our analysis provides initial steps toward an ecological rationality perspective on institutional design, exploring how the structure of institutions can fit or exploit the structure of tools in the mind's adaptive toolbox.

Transparency Without Trade-offs in Traffic and Soccer

When making a decision based on a list of factors, perhaps the most common recommendation in the decision sciences is to weigh many factors. The decision maker is supposed to apply implicit weights to various factors and trade off the relative value of one factor against another. Weighing many factors embodies the essence of oft-repeated adages about good decision making that insist on considering all the evidence, carefully analyzing trade-offs, not rushing to make snap decisions, and so on.

In this section, we examine two institutions that help agents to make transparent decisions *without* weighing many factors. Decision rules that require no trade-offs are referred to as noncompensatory, because decision factors have a fixed ranking of importance, and factors that are less important cannot overrule, or compensate for, higher ranking factors. The way we alphabetize words in the dictionary provides a good example of a particular type of noncompensatory decision strategy called a lexicographic rule, with the letters in each word representing the potential factors that contribute to the decision of which word is ordered first. In ordering the words *azimuth* and *babble*, for example, the first letter, or factor, by itself leads to an unequivocal decision: *azimuth* comes before *babble* because the first letter of the former comes before the first letter of the latter—the subsequent letters do not matter, even if they point in the “opposite” direction (e.g., “z” comes after “a”). This is precisely what allows us to alphabetize words quickly, without comparing all their letters.

Lexicographic rules have proven successful in the design of institutions in environments where decisions must be fast and at the same time transparent, that is, readily predictable by others so **(p. 416)** as to minimize uncertainty and misunderstanding in interactions. Speed and transparency are especially valuable when smooth temporal coordination between individual actors is required, as in the following brief analysis of traffic rules.

Determining Right-of-Way

Ancient Rome was a city of perhaps a million people, but it lacked traffic signs (let alone stoplights) to guide the many pedestrians, horse riders, and chariots on its roads. Right-of-way was determined by wealth, political status, and reputation. In case of ambiguity about which of these cues was more important, the issue was decided by how loudly accompanying slaves could yell, or by physical force. This led to much confusion and conflict on the roads of Rome. Historian Michael Grant even controversially hypothesized that traffic chaos pushed Nero over the edge, leading him to burn the city in the year 64 A.D. with hopes of subsequently building a more efficient road system (Gartner, 2004).

In contrast to the compensatory system of Nero’s time that required simultaneous consideration of multiple factors, right-of-way throughout most of the world is now governed by noncompensatory lexicographic rules that leave far less room for ambiguity, although the details differ between countries. In Germany, for example, the right-of-way rules for deciding which of two cars approaching an intersection gets to go through first include the following hierarchy:

> If you come to an intersection with a police officer regulating traffic, follow the officer's directions and ignore everything else.
>
> *Otherwise*, if there is a traffic light or stop sign, follow it and ignore everything else.
>
> *Otherwise*, if there is a yellow-diamond right-of-way sign, proceed.
>
> *Otherwise*, if there is a car approaching from the right, yield to it.
>
> *Otherwise*, proceed.

So, for example, the stopping gesture of a police officer cannot be overruled by any combination of lesser priority cues suggesting that one may drive through an intersection, including a green light, right-of-way sign, and being to the right of other approaching cars. This is the hallmark of a lexicographic system.

If drivers had to apply weights to various factors or cues and compute weighted sums to decide whether to drive through any given intersection, disastrous consequences would surely follow. Individual decision processes would slow down as more information **(p. 417)** would need to be looked up and processed. The possibility of overlooking information, computational errors, and individual variation in weights assigned to cues would make it almost impossible to anticipate how other drivers might act. Processing cues in a simple lexicographic fashion, and relying on other drivers to do so as well, frees cognitive resources for other important driving tasks and makes the roads safer. Noncompensatory rules also help settle arguments about fault quickly when accidents do occur. These benefits of the transparency of noncompensatory regulation can also be found in a variety of other institutions—for example, deciding outcomes in sports.

Making It to the Next Round

The International Federation of Football Associations (FIFA) is the governing body of the soccer world. It manages a number of major soccer competitions, including the World Cup, which attracts more than a billion television viewers around the world. Economists have studied the design of sports tournaments, focusing on designs that maximize profits (Knowles, Sherony,

& Haupert, 1992), or whether tournament rules satisfy certain axioms (Rubinstein, 1980). As it turns out, FIFA also employs lexicographic rules to increase transparency and minimize controversy.

World Cup tournaments involve a group and a knock-out stage. In the latter knock-out stage, teams are eliminated with a single loss. In the group stage, however, teams are usually arranged in groups of four, where each team plays all others in the group, and a single loss is not necessarily fatal. To determine which team advances to the next stage, FIFA uses a point system (with points being distinct from goals). The winner of each match is awarded three points, regardless of the final score, and the loser receives zero points. If a match's final score is a tie, then each team gets one point. After all group-stage matches are played, teams in each group are ranked according to points to determine who advances to the knock-out stage.

Because ties in these point totals can occur at the group stage, FIFA had to develop a system to produce an unambiguous ranking when a tie arose. FIFA considers multiple cues for ranking teams at the group stage. Following a lexicographic rule similar to take-the-best (Gigerenzer & Goldstein, 1996, 1999), a team is ranked above its competitor when it is favored by one of the following cues, considered in the listed order (starting with the point totals), taken from the FIFA 2010 tournament regulations (Regulations, 2010, pp. 47–48):

1. More points earned in all group matches;
2. Larger goal differential in all group matches;

- **(p. 418)**

3. More goals scored in all group matches;
4. More points earned in group matches against teams with the same values on cues 1, 2, and 3;
5. Larger goal differential in group matches against teams with the same values on cues 1, 2, and 3;
6. More goals scored in group matches against teams with the same values on cues 1, 2, and 3;
7. Random tie-breaker: If two or more teams tie according to the first six cues, then the ranking is made at random by drawing lots.

A similar set of cues was employed in the lexicographic rule used to decide the notorious "Shame of Gijón" group ranking in the 1982 World Cup in Spain, comprising teams from Algeria, Austria, Chile, and Germany. Only two teams were to advance to the next stage, but according to FIFA's group-stage point system, Germany, Algeria, and Austria all had four overall points, while Chile had zero.[3] Further cues were applied in order and determined that Austria and Germany would advance to the next round. But this result led to widespread suspicion and criticism, because the group-stage game between these two neighbors took place after the first five group-stage matches were finished. Germany and Austria knew, even before their match began, that a 1:0 result for Germany would allow both to advance. Many fans suspected that the teams somehow colluded to ensure their joint success over Algeria. After this incident, FIFA redesigned the timing of matches so that its ranking rule could not be exploited. The last two group-stage games now take place simultaneously.

Why does FIFA use a lexicographic rule to produce group-stage rankings rather than weighting and adding all the cues? Unlike the right-of-way example in fast-moving traffic, plenty of time and computing resources are available to process the final group-stage scores and arrive at rankings using more complex, compensatory ranking schemes. One reason a more complex method is not used, though, appears to be transparency. The hypothesis is that when stakeholders in any ranking scheme clearly understand the process by which results are obtained, they accept those rankings—or, as in the Shame of Gijón, are able to spot problems with them—more readily than they do when complex algorithms are employed. This is based on the idea that rankings, like tax schemes and constitutions in democracies (Berg, 2006), require a large degree of shared belief in their legitimacy in order to coordinate action effectively. **(p. 419)** The basic principles behind FIFA group-stage rankings are easy to understand: Points earned are more important than goal differentials, goal differentials are more important than goals scored, and all arguments about how much more important one cue is than the next are moot.

One way to measure the simplicity of a ranking device is via its informational requirements. A ranking device based on a regression model with the cues described above would rely on all available information to make any pair-wise comparison: Plug in cue values for two teams, apply beta weights from the regression model, and rank the team with the higher score ahead of the other. In contrast, the lexicographic ranking rule that FIFA uses operates much more frugally, in the sense that most pairs of teams can be ranked based on a single reason, without looking up each team's values for all cues. This reliance on typically little information also makes the application of the rule more transparent.

To determine how informationally frugal the FIFA strategy is, we calculated an empirical frequency distribution of how many cues in the list given above would have been needed historically to determine pair-wise team rankings.[4] In the 18 World Cups played before 2010, there were 88 groups and a total of 529 pair-wise rankings. For each of these 529 cases, we determined how many of the seven cues in that order would need to be looked up to specify the ranking. As can be seen in Figure 16-2, most of the time (471 out of 529 cases), the first cue alone (overall points earned) sufficed to specify the ranking. One ranking was decided by chance (i.e., cue 7, after no other cues were decisive). The average number of cues looked up was 1.2, indicating a high degree of informational frugality. This was due in large part to the high discrimination rate of the first cue (which was so high because the cue is nonbinary), allowing it to determine most of the ranking decisions.

Transparency is chief among the virtues of FIFA's lexicographic ranking rule. On the other hand, many organizations, such as casinos, are strategically designed for nontransparency—so that their customers, such as gamblers, cannot easily see how they operate. We next investigate the nontransparency of casinos and show how their strategies can be understood in terms of heuristic models of behavior that depart from the standard economic model. **(p. 420)**

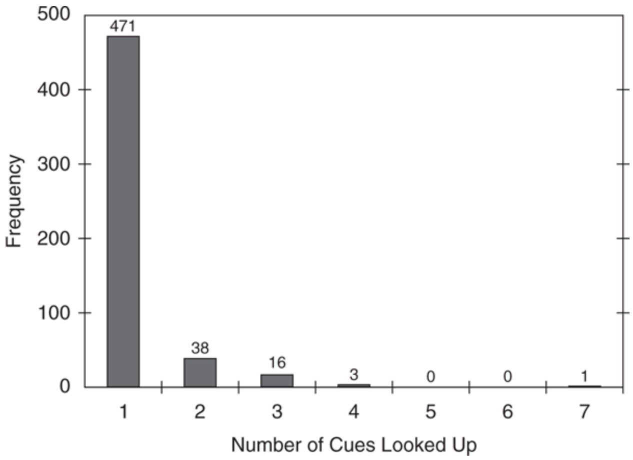

Figure 16-2: Frequency distribution of the number of cues looked up to determine pair-wise rankings in the group stage of World Cup tournaments 1930–2006. (Precise counts indicated above each bar.)

Beliefs About Winning on Slot Machines: It’s Not All in the Players’ Heads

In 2007, Americans spent $34 billion gambling in commercial casinos (American Gaming Association, 2008), perhaps half what they spent across all forms of institutionalized gambling (Christiansen, 2006). This figure is on the same scale as the entire fast food industry ($150 billion) and greatly exceeds the value of another entertainment industry, the $600 million worth of movie tickets purchased (American Gaming Association, 2008). To make a profit, gambling institutions are designed so that the average gambler loses money. Because gamblers can expect this loss, the fact that so many people who turn out to be risk averse in other decision domains still choose to gamble presents a perplexing challenge to the economic model of individual decision making (Eadington, 1988; Wagenaar, 1988; Walker, 1992b).

Nonetheless, many economists could see this paradoxical gambling behavior as readily explained by the standard economic model by pointing to nonmonetary utility as compensation for monetary losses (Becker, 1978). When people choose to gamble, this reasoning goes, they willingly forgo a sum of money (the expected monetary loss from gambling) as the purchase price for their entertaining or exciting experience (Eadington, 1988). Indeed, empirical research supports the view that the utility of gambling stems from many nonmonetary sources along with the obvious monetary one **(p. 421)** (Bennis, 2004; Smith & Preston, 1984; Wagenaar, Keren, & Pleit-Kuiper, 1984; Zola, 1963). Nonetheless, although other sources of utility besides expected winnings are undoubtedly part of what motivates gamblers, there is abundant evidence that many people gamble because they have false beliefs about their ability to win. Often this is a belief that they have an advantage over the casino, but casino gamblers also systematically overestimate their chances of winning, overestimate the role of skill in games that are largely determined by chance, and use gambling strategies that do not work (Ladouceur, 1993; Lambos & Delfabbro, 2007; Miller & Currie, 2008; Sundali & Croson, 2006; Wagenaar, 1988; Walker, 1992b). Thus, at least part of why people gamble seems to stem from a systematic failure to estimate their expected payoffs correctly.

Theories attempting to account for this faulty payoff estimation fall into two broad categories. The first, and far more common, type of theory identifies the source of the problem as originating inside gamblers’ minds. According to such theories, people gamble because of shortcomings in how they think and reason, including, among other things, a failure to understand the nature of probability and randomness (Gaboury & Ladouceur, 1988, 1989; Ladouceur & Dubé, 1997; Ladouceur, Dubé, Giroux, Legendre, & Gaudet, 1995; Lambos & Delfabbro, 2007; Metzger, 1985; Steenbergh, Meyers, May, & Whelan, 2002; Sundali & Croson, 2006; Wagenaar, 1988; Walker, 1990, 1992a).

The second type of explanation, to which we subscribe, focuses on factors in the external environment: While acknowledging that gamblers may sometimes have false beliefs about their chances of winning and use the wrong heuristics, we argue that the source of these shortcomings lies not so much in biased or irrational thinking, but rather in the gamblers' environment and their interactions with it (see, e.g., Bennis, 2004; Dickerson, 1977; Griffiths & Parke, 2003a; Harrigan, 2007, 2008; Parke & Griffiths, 2006). Specifically, there is a *mismatch* between the (otherwise usually adaptive) heuristics used by gamblers on the one hand, and the structure of the casino environment on the other—the opposite of the ecologically rational *match* between heuristics and environments explored extensively elsewhere in this book.

Why does this mismatch come about? Because it is in the casinos' interest for this mismatch to exist, and they construct the gamblers' environment so that it does. The degree to which casinos *intentionally* design games to exploit normally adaptive heuristics, or alternatively simply select the games that end up garnering the greatest profits and which turn out to be the ones that promote this mismatch, is an open question. But the result is a wide range of casino games exquisitely designed to exploit otherwise adaptive heuristics to the casinos' advantage. They produce representations **(p. 422)** in the environment that provide the cues that the gamblers' heuristics rely on; as we will see, these cues are about the success and failure of gambling heuristics and about the ways machines operate. (This is similar to how companies exploit the often-adaptive use of recognition to lead people to buy the products that they recognize through advertisement—see Goldstein & Gigerenzer, 1999, 2002.) Unlike the organ-donor example, in which some environments were *inadvertently* designed in a way that discouraged organ donation, the casino industry has a powerful incentive to design environments that contribute to false beliefs and a corresponding maladaptive application of heuristics, since their economic success stems from their ability to get and keep people gambling.

We focus here on slot machine environments constructed by Las Vegas resort casinos to encourage use of misleading cues (Bennis, 2004). In the standard economic model, logically equivalent representations of information are irrelevant, because deductive logic, which is equally capable of utilizing information in any format, is assumed to underlie behavior. But psychologically, different representations of the same information can have a large impact on how people use it to reach decisions (see, e.g., chapter 17 on the impact of different representations of medical information). Thus, the casinos' ability to influence gambling through the strategic representation of information becomes understandable only when the economic model is revised to incorporate psychologically realistic theories of cognition.

Representing the Experience of Winning

Major hotel-casino resorts in Las Vegas have one or more casino floors where hundreds, sometimes thousands, of slot machines are arranged in aisles with lines of machines on both sides, back to back against other lines of machines. During play, contemporary slot machines generate an abundance of audio and visual cues that are difficult to miss or ignore.

When slot machine players cash out winnings, metal tokens typically drop several inches onto a metal tray, generating loud clanking sounds that can be heard almost constantly and from virtually every direction in busy casinos.[5] Many machines amplify the clanking of coins, which makes winning a very public and familiar (if vicarious) event to those who spend time in a casino. If slot **(p. 423)** players do not immediately collect their tokens, wins are announced with escalating beeping music, marking the increasing credits that players can cash out in the future. In this case, the amplified sound of growing credits often accrues at a faster pace than the credits themselves, contributing to a subjective perception that players have won more than they actually have.

In addition to audio, slot machines can generate visual cues that can be seen by others from a distance. For example, most slot machines in Las Vegas are equipped with a spinning siren light on top, which flashes whenever a major jackpot has been hit. Larger jackpots need to be paid by hand, and during the time it takes for slot machine attendants to walk to the winner and deliver their money, the winning machine continues to flash and blare, sometimes for more than half an hour. Slot machine players regularly complain about how slow attendants are to pay off major jackpots. These long waits serve to advertise large jackpots in a manner that makes their occurrence appear more frequent than it is. On busy nights, many large-jackpot winners can be observed, often at the same time, due in part to extended payoff wait times. Some casinos prominently display posters of past winners of major jackpots, photographed while being paid with over-sized publicity checks.

While winnings are emphasized and communicated through a wide variety of cues in the casino environment, losses are hardly signaled at all. This raises questions about gamblers' perceptions of win and loss probabilities: Where environments have been constructed to highlight winnings and hide losses, can we expect individuals to see through the selectively represented cues and formulate hard-nosed expectations based on the logic that casinos must profit to stay in business, that gambling is a zero-sum game, and therefore that they should expect to suffer losses? Or might gamblers too often expect to win because instances of winning are almost always visible in the casino?

Heuristics designed to adaptively guide foraging behavior by following the observed successes of others, such as an "imitate the successful" rule (Boyd & Richerson, 1985), run into problems in the casino environment. To the extent that frequencies of success are processed unconsciously by observing other gamblers in a casino, the casinos' nonrepresentative construction of cues, which include uninformative or misleading signals from sirens and flashing lights, may significantly promote gambling behavior, to the detriment of most gamblers.

Representing How Slot Machines Work

Another way that nonrepresentative cues distort gamblers' perceptions of the constructed casino environment revolves around **(p. 424)** the inner workings of slot machines. Until the 1960s, slot machines worked much as their exterior design suggests. A machine had three reels covered with symbols, each with around 20 possible stop positions where the reel could come to rest showing one of the symbols, and each stop had an equal probability of occurring (Cardoza, 1998; Kiso, 2004; Nestor, 1999). Given this design, there would be 20^3 (i.e., 8,000) possible outcomes, and a jackpot requiring a unique combination of three symbols would occur with probability 1 in 8,000, or .000125. After observing the pay line (i.e., the payoff-determining three symbols shown when the reels stop spinning) on several spins on an old machine, along with a view of the symbols above and below the pay line, savvy players could estimate the actual number of stops and the frequency of each symbol on each reel. They could then compare this assessment with the payout chart for winning combinations to determine the expected value of playing a particular machine.

Figure 16-3 shows an old and a new slot machine side by side. On the surface, new slot machines look very much like older machines, but their internal mechanics are entirely different. New slot machines use digital random number generators rather than physically spinning reels to determine wins and misses. Nevertheless, contemporary machines continue to display spinning reels, providing nonrepresentative cues meant to distort the true payoff-generating process. If, for example, the largest jackpot requires

Figure 16-3: Left: The “Liberty Bell,” the father of the contemporary slot machine (image courtesy of Marshall Fey), released to the public in 1899 (Legato, 2004). Right: A contemporary 25¢ banking slot machine with a siren light on top (image courtesy of Paul and Sarah Gorman).

(p. 425) three red sevens, it would be possible for the microchip designers to assign a 1 in 1 billion chance of this outcome, even while the machine’s external design falsely suggests a 1 in 8,000 chance of winning, as would have been the case on older machines. Similarly, inflated frequencies of hope-inspiring near-jackpot misses can also be created. Such strategically nonrepresentative design is standard practice in the casino environment (Griffiths & Parke, 2003b; Harrigan, 2008; Turner & Horbay, 2004). Institutional designers go to great lengths to represent information in ways that should not matter in the standard economic model (e.g., rational Bayesian updaters making inferences about winning probabilities should not be influenced by sirens, flashing lights, and uninformative spinning wheels). But this strategy works for the interests of the casinos because gamblers use decision processes built on psychological mechanisms that are sensitive to the structure of their environment and which can thus be subverted by situations constructed to provide misleading and irrelevant cues.

Ecological Rationality in Institutional Design

Unlike the axiomatic definitions of rationality that economic models draw upon, ecological rationality implies that evaluations of decision processes cannot be undertaken in isolation, strictly at the level of one individual's internal logical consistency. Rather, decision processes should be evaluated contextually according to how well they match the environments in which they are used. These distinct notions of rationality have important implications for the analysis of institutions.

According to the standard economic model, there is no need to study or analyze strategic interactions between institutional designers and nonoptimizing heuristic users, because people would eventually abandon such heuristics in favor of optimal behavior. The space of problems to which the economic model is applicable is therefore rather narrowly circumscribed because of its stringent behavioral assumptions, such as exhaustive search for information and options, optimal weighing of costs and benefits, and adherence to logical norms of probabilistic reasoning. These assumptions rule out consideration of institutions that are built to work with populations of real humans using heuristics.

The organ-donor example shows how the standard economic model misses an important institutional determinant of real-world behavior: the setting of defaults that do not change feasible choice sets yet influence heuristic-based decision making nevertheless. In the same way, psychological theories that try to understand behavior solely in terms of knowledge and beliefs also miss the **(p. 426)** importance of heuristics interacting with institutions. The examples of noncompensatory rules regulating traffic and professional soccer rankings highlight psychologically important objectives that are difficult to motivate using the standard economic model: decision simplicity and transparency. These factors are critical for many institutional designs, and designers can achieve them not by trying to manipulate economic models of behavior but by creating systems that fit human lexicographic decision strategies.

The last example of the casino environment shows how institutions can be designed to exploit vulnerable heuristics that rely on transparent information structure to produce adaptive choices in other domains. People typically expect transparency and use simple rules exploiting straightforward relationships between cues and outcomes, such as "where I've seen success (or near success) up to now, I will expect success in the future."

Casinos can exploit this by subverting the cue–outcome relationship and leading gamblers to think mistakenly that they are on the path to likely success. Such conflict of interest between institutional designers and agents who interact with those institutions is also commonly analyzed within the standard economic model framework. However, the ongoing systematic exploitation of gamblers by casinos is understood much more easily using the concept of designed mismatch between heuristics and decision environments than through complicated rationalizations of gambling as a positive-surplus-yielding activity where intrinsic, nonpecuniary gains outweigh monetary losses.

In the book *Simple Rules for a Complex World*, Richard Epstein (1995) similarly builds a case for the benefits of designing institutions with simple transparent rules and the dangers of going in the opposite direction. He argues that in the United States, the law has become excessively complex and nontransparent, resulting in an overly litigious environment where complexity is exploited by lawyers. According to his view, complexity in the legal code makes outcomes more malleable to intervention by skilled legal craftsmanship and, thus, more volatile and less robust. The result has been a kind of arms race where more and more lawyers are necessary to protect individual and corporate interests against the claims of others, with the outcome depending on who has the money to hire the best team of lawyers rather than on more ideal standards of justice. Epstein advocates that we reduce our complex avalanche of laws to just six simple mandates, such as property protection. This will save on legal costs and, more importantly, reduce uncertainty through greater transparency, thereby increasing public trust in government institutions, and as a consequence, compliance with the law. (For an extensive investigation of the **(p. 427)** general question of how legal institutions shape heuristics and vice versa, see Gigerenzer & Engel, 2006.)

The central point is that environmental structure is not simply an independent variable on which decision processes and their performance depend. Environments themselves can be, and often are, actively structured, selected, and intentionally designed (both by humans and by other animals —see Hansell, 2005; Odling-Smee, Laland, & Feldman, 2003). A crucial ingredient for successfully analyzing the institutional dynamics in which environments and behavior co-evolve is understanding the decision heuristics that are actually used by the population under consideration (see, e.g., Todd & Heuvelink, 2007, and chapter 18), not unrealistic optimizing strategies derived from the standard economic model. The descriptive question of how well, or poorly, people make decisions in particular environments is thus also, fundamentally, a question about how well environments are tuned to particular decision tasks. From the standpoint of ecological rationality, the normative question is not simply how our reasoning processes can be improved, but also how to design environments to better match the ingenious human cognitive hardware that comes for free.

Notes:

(1.) The law is the Gesetz über die Spende, Entnahme und Übertragung von Organen, BGBI 1997, Article 2631. A German government website (www.organspende-kampagne.de/) provides an official form that one can use for the purpose of changing donor status. The official form is not required, however, nor any formal registration. In some cases where relatives have been clearly informed of the individual's wish to become an organ donor should the occasion arise, verbal consent may even substitute for written consent.

(2.) This is not to say that educational campaigns and increased knowledge about the issues cannot make a difference or that a default heuristic explains everything. Many people do not know they face an organ donation decision at all (including one author of this paper who thought he was a donor but discovered he needed to send in a letter in addition to marking his preference on his driver's license application). But for those (many) who do know they have a choice, most go with the default. If people assume that defaults were designed to represent the average person's preference or the greater good, and if this assumption is generally correct, then following the default heuristic would be appropriate.

(3.) In 1982 the winner of a game was allocated two points (not three as is the case at the time of this writing).

(4.) Again, note that this set of cues is not exactly the same as that used in some of the World Cups we analyzed.

(5.) Coin and token payouts are rapidly being replaced with paper vouchers such that this method of manipulating subjective experience may soon be a thing of the past.