Ecologically Rational Choice and the Structure of the Environment

Timothy J. Pleskac
Michigan State University

Ralph Hertwig
Max Planck Institute for Human Development, Berlin, Germany

In life, risk is reward and vice versa. Unfortunately, the big rewards people desire are relatively unlikely to occur. This relationship between risk and reward or probabilities and payoffs seems obvious to the financial community and to laypeople alike. Yet theories of decision making have largely ignored it. We conducted an ecological analysis of life's gambles, ranging from the domains of roulette and life insurance to scientific publications and artificial insemination. Across all domains, payoffs and probabilities proved intimately tied, with payoff magnitudes signaling their probabilities. In some cases, the constraints of the market result in these two core elements of choice being related via a power function; in other cases, other factors such as social norms appear to produce the inverse relationship between risks and rewards. We offer evidence that decision makers exploit this relationship in the form of a heuristic—the *risk–reward heuristic*—to infer the probability of a payoff during decisions under uncertainty. We demonstrate how the heuristic can help explain observed ambiguity aversion. We further show how this ecological relationship can inform other aspects of decision making, particularly the approach of using monetary lotteries to study choice under risk and uncertainty. Taken together, these findings suggest that theories of decision making need to model not only the decision process but also the environment to which the process is adapted.

Keywords: decisions under uncertainty, heuristic, bounded rationality, ecological rationality, risk and reward

While God may not gamble, animals and humans do, and . . . they cannot help but to gamble in an ecology that is of essence only partly accessible to their foresight.

—Egon Brunswik (1955a, p. 236)

In *An Essay Concerning Human Understanding,* the empiricist philosopher-cum-psychologist John Locke (1690/1959) emphasized the inescapable bounds of human knowledge. According to Locke, humans conduct their lives—apart from a few essential things they can be certain of and see in broad daylight, such as their own existence and that of God—in the "twilight, as I may so say, of probability, suitable, I presume, to that state of mediocrity and probationership, he [God] has been pleased to place us in here" (p. 498). How should we reason and act in this twilight? Part of an answer had already been suggested in 1654, some decades before Locke's depiction of the human condition. In that year, the mathematicians Blaise Pascal and Pierre Fermat began an epistolary exchange in which they discussed various gambling problems—in

other words, lotteries. Their correspondence formed the foundation of the calculus of probability and mathematical expectation. Expectation, in turn, was soon to become the central doctrine of a new brand of rationality, grounding choice in the maximization of expected outcomes. Thus, a pragmatic rationality emerged that abandoned ideals of certainty and was thought to describe how the educated *homme éclairé* reasoned in what Locke portrayed as the twilight of probability (Hacking, 1975).

Since Pascal and Fermat, lotteries have time and again played a prominent role in the development of normative and descriptive accounts of choice. Sets of equiprobable outcomes as produced by gambling devices such as a fair coin or a die make it possible to define probabilities in a nonarbitrary fashion (Gigerenzer et al., 1989). Moreover, lotteries offer a simplified model of risky choices in the world, displaying the key characteristics of risk and rewards. Consequently, as Lopes (1983) put it, "the simple, static lottery or gamble is as indispensable to research on risk as is the fruitfly to genetics" (p. 137).

The focus on lotteries has impacted decision theory in many different ways, including how researchers classify decisions. A common distinction is between decisions under risk and decisions under uncertainty (Knight, 1921). Decisions under risk refer to situations where the probability distribution over the possible outcomes is known, as is commonly the case in monetary lotteries. Most experimental investigations of how people choose have focused on decisions made under risk (see Weber, Shafir, & Blais, 2004). However, countless everyday decisions—which candidate to support, which witness to believe, which venture to invest in, which scientific theory to endorse, when to cross a busy street—are made with limited or no knowledge of probabilities. In reality,

humans often navigate not in the twilight of probability but in the fog of uncertainty.

How people should and do reckon with uncertainty is one of the most vexing problems in theorizing about choice. A number of solutions have been proposed. One is to turn uncertainty into risk. Savage (1954) and Ramsey (1926) proposed that "for a 'rational' man . . . *all* uncertainties can be reduced to *risks*" (Ellsberg, 1961, p. 645) by replacing the objective probabilities in expected utility theory with subjective ones. Moving from rational man to actual people, psychologists have investigated two ways in which people respond to uncertainty. One response is to intuit subjective probabilities from knowledge and memory (Tversky & Fox, 1995). The second is to sample the external world for information, thus replacing uncertainty with what Knight (1921) called *statistical probabilities* or probability values known only via samples of information (Hertwig & Erev, 2009). We consider still another process people may enlist to navigate the fog of uncertainty. The process is inferential in nature and rests on the notion that probabilities can be approximated from statistical regularities that govern real-world gambles. Our account also aims to bring an adaptive view to the inquiry of rational choice.

An Adaptive View of Rational Choice

According to Simon (1956), "a great deal can be learned about [boundedly] rational decision making . . . by taking account of the fact that the environments to which it [the organism] must adapt possess properties that permit further simplification of its choice mechanisms" (p. 129). On this view, human cognition reflects and exploits statistical regularities in the world. This has been found to occur with perceptual processes (Brunswik, 1940; Gibson, 1979; Proffitt, 2006; Shepard, 1990, 1994), memory and categorization (Anderson, 1991; Anderson & Schooler, 1991), and inferential processes (Fiedler, 2000; Fiedler & Juslin, 2006; Juslin, 1994). Moreover, environmental regularities have been suggested to be the foundation of simplifying mechanisms that exploit such regularities in the service of good performance (Gigerenzer, Todd, & the ABC Research Group, 1999; Hertwig, Hoffrage, & the ABC Research Group, 2013; Simon, 1955; Todd, Gigerenzer, & the ABC Research Group, 2012).

One way to conceptualize how people exploit ecological regularities is with Brunswik's (1940, 1952) *probabilistic functionalism* (see also Dhami, Hertwig, & Hoffrage, 2004; Hammond & Stewart, 2001). Probabilistic functionalism rests on two principles. The first principle is that people do not act or perceive for the sake of acting or perceiving. They do so to achieve specific ends (functionalism). These ends can be basic, such as obtaining food, avoiding danger, or predicting a competitor's behavior, or they can be more intricate, such as investors identifying the most profitable stocks to buy. Irrespective of the ends, people often do not have direct access to or direct knowledge of the criterion of interest (e.g., the distance of an object, the likely strength of a competitor, the likely value of a good or an alternative). Instead, they need to infer the distal criterion from proximal cues that are probabilistic indicators of the criterion. This probabilistic inference process is the second principle, thus the name probabilistic functionalism. In terms of decision making under risk and uncertainty, payoffs, probabilities, ambiguity or lack thereof (Ellsberg, 1961), and anticipated or experienced affect (Damasio, 1994; Mellers, Schwartz,

Ho, & Ritov, 1997) are just some of the many possible cues that a decision maker uses to infer the value of an alternative.

The Lens Model in Decisions Under Uncertainty

Brunswik's (1952; see also Cooksey, 1996) lens model draws together the key concepts within probabilistic functionalism (see Figure 1). It shows the distal criterion on the left and the observed response on the right. The functional validity, represented by the arc between these two variables, measures the degree to which the organism's response (e.g., decision) attains the distal criterion. The proximal cues mediate the relationship between the distal criterion in the environment and the organism's response. Ecological validities measure the predictive strength of a cue with regard to the distal criterion. Utilization validities measure the extent to which the cues are predictive of the organism's response. Finally, the cues themselves are typically interrelated, thereby allowing redundancy in predicting the distal criterion. When some cues become unreliable or unavailable, the organism can exploit this redundancy by substituting or alternating between different cues. This is possible because of what Brunswik (1955b) called the *mutual substitutability* or *vicarious functioning* of cues.

It is these properties of intercue relationships and substitutability that we suggest offer a new perspective on how people make decisions under uncertainty. Under uncertainty, cues such as the payoffs associated with different courses of actions may be accessible, whereas other cues—in this case, the probability with which those payoffs occur—are not. This missing probability information has been problematic for choice theories as typically both payoffs and probabilities are used in determining the value of options and in choosing (e.g., Bernoulli, 1954; Kahneman & Tversky, 1979; Savage, 1954; Tversky & Fox, 1995; Tversky & Wakker, 1995; von Neumann & Morgenstern, 1947). However, if payoffs and probabilities are interrelated, then this ecological property can permit the decision maker to infer hidden or unknown probability distributions from the payoffs themselves, thus easing the problem of making decisions under uncertainty.
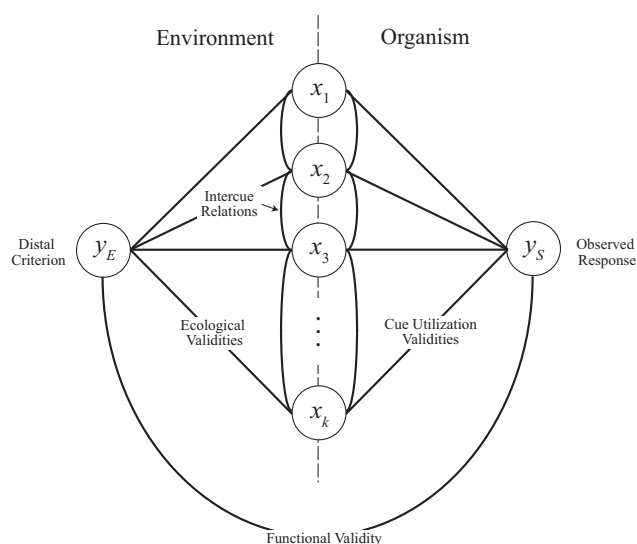

Figure 1. Brunswik's lens model (adapted from Brunswik, 1952).

The Road Map

Environmental analyses of the relationship between payoffs and probabilities or risks and rewards are lacking however. Therefore, we first analyze the ecological relationships between payoffs and probabilities across a variety of real-world environments. We also discuss the possible driving forces behind the observed relationships. Then, we investigate the extent to which people take advantage of this relationship to make choices. In particular, we provide evidence that people appear to use a simple heuristic, the *risk–reward heuristic,* to infer unknown probabilities from observable payoffs during decisions under uncertainty. We conclude by examining the implications of the heuristic and, more broadly, the structural relationship between payoffs and probabilities for decision making under uncertainty.

An Ecological Analysis: Are Risk and Reward Correlated?

Is there an environmental relationship such that reward could be a proxy for risk? According to literary intuition, for instance, the relationship is self-evident: "One doesn't discover new lands without consenting to lose sight of the shore for a very long time" (Gide, 1973, p. 309). The same idea is echoed in the proverb "Nothing ventured, nothing gained." Likewise, the relationship is taken for granted in the field of finance, where any casual search of the Internet or of self-help books reveals statements such as "understanding the relationship between risk and reward is a key piece in building your personal investment philosophy" (Little, 2013, para. 2).

In order to go beyond anecdotal evidence, we conducted several ecological analyses involving environments in which probabilities are known or obtainable. This focus follows necessarily from our goal to *quantify* the ecological relationship between payoffs and probabilities. For two reasons, we start our ecological analyses by considering the environments of monetary lotteries. First, monetary gambles and chance devices are a ubiquitous property of human ecology. In 2009, for instance, $34.6 billion was spent at commercial casinos in the United States (American Gaming Association, 2011). In that same year, state lotteries sold $52 billion dollars' worth of lottery tickets (U.S. Census Bureau, 2009). On average, this means households spent $484 on lotteries, which is more than the $435 that the average U.S. household spent on alcoholic beverages that year (Bureau of Labor Statistics, 2009).[1] At the other end of the lottery spectrum, the average U.S. household spent approximately $309 on life and other personal insurance (Bureau of Labor Statistics, 2009). Because of their omnipresence, monetary gambles may afford people one of the clearest views onto the statistical structure of an environment. Our second reason for initially focusing on lotteries was, admittedly, convenience: Lotteries provide reasonably accessible ecologies for a quantitative examination of the risk–reward relationship. We first consider one of the most popular and common casino games, roulette.

Roulette

In roulette, a ball is spun around a wheel with 38 pockets labeled 1 to 36, plus 0 and (in American, but not European, roulette) 00.

The odd numbers are red, and the even numbers are black. The player places a bet (e.g., $1) on one or more outcomes occurring (e.g., the ball stopping on pocket 12). If the ball lands on one of the player's selected outcomes, the payout is $(36 - n)/n$, where n is the number of outcomes the player is betting on plus the original wager. Each of the 155 possible bets (22 unique bets) has an expected loss to the player (in American roulette) of 5.3 cents for every $1 spent (2.7% of the wager in European roulette). The simplicity of roulette renders obvious the inverse relationship that exists between payoffs and probabilities across all 155 gambles, where the higher the payoff, the less likely the event is to occur (see Figure 2). Less obvious is that the functional relationship between probabilities and payoffs takes the form of a power function, as the log-log plot in Figure 3A illustrates. We return to why it takes the form of a power function shortly.

Horse Track

Does this relationship generalize to gambles that involve what Knight (1921) called statistical probabilities whose values are known only via samples of information? In order to find out, we analyzed 34,443 bets offered at flat horse races in Britain between 1985 and 1995. Obtained by Jullien and Salanie (2000) from RACETRACK Limited, the data set contains the number of horses in each race, the odds or starting price of each horse at the time the race began, and the order of arrival of the horses. With the large number of races, we could estimate the probability that a horse with particular odds won across all the races and years. There was an average of 10.7 horses in each race. The odds ranged from clear favorites (1 to 50) to incredible long shots (1,000 to 1), with the average odds being 16.6 ($Mdn = 12$, $SD = 18.1$). The average expected loss per £1 wagered was £0.16, with the loss falling off as the odds increased. The power function relationship observed in roulette also exists in horse betting. Figure 3B shows this relationship between payoffs relative to a £1 bet and probabilities. Low payouts are more likely to win the race, and the probability decreases as the payout increases in accordance with a power function relationship. Departure from linearity in the log-log space occurs when the odds offered are less than 1, which is also consistent with a power function.

Life Insurance

Does the negative relationship between payoffs and probabilities also extend to life insurance? We examined the risk–reward relationship in the pricing of voluntary group term life insurance offered in 2010 at a large Midwestern university. This insurance is an optional benefit offered in addition to the life insurance plan in which employees are automatically enrolled. The term life premiums consist of biweekly payments that differ according to age group (under 35, 35–39, 40–44, . . . , and 80–84).[2] The payouts are expressed relative to $1 spent on premiums. In order to determine the probability of a payout, we used a 2007 period life table from the U.S. Social Security Administration (n.d.) and averaged across the given ages and across gender (for the "under

---

[1] Forty-two states had lotteries in 2009. The number of households in each state was obtained from the U.S. Census Bureau (2011).
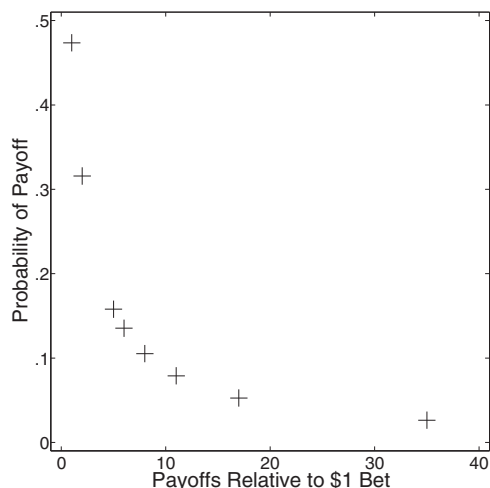[2] In the analysis, we used the premiums for nonsmokers.

Figure 2. Relationship between payoffs and probabilities in the game of roulette (American).

35" group, ages 22–34 were averaged). As Figure 3C shows, a negative power function relationship also emerged between payoffs and probabilities in this domain.

Farmers' Gamble: Artificial Insemination of Dairy Cows

What about real-world gambles beyond monetary lotteries? Farmers, on a daily basis, take highly consequential bets, from when to plant and harvest their crops to how to manage their livestock. One bet dairy farmers regularly make is on the specimen of bull semen used to artificially inseminate their cows. To do so, they are given a catalogue listing the semen currently available from different sires. For each specimen, approximately 32 different properties associated with the sire and his offspring are reported (e.g., fat and protein content of the milk from daughters of that bull, daughters' calving ease). In order to render selection easier, farmers are typically advised to focus on two numbers: an index called lifetime net merit and the sire conception rate (Kringle, 2008; Schnell, 2008).[3] The lifetime net merit is an estimate of the average improvement in profit that a dairy farm can expect from the daughters of a specific bull; the sire conception rate is an estimate of the conception rate of the sire's sperm.[4] In other words, the lifetime net merit is the payoff, and the sire conception rate is the probability.

Do farmers make their selections in an environment in which risk and reward are negatively correlated? We obtained data for all Holstein sires from a Genex Cooperative catalogue in May 2012. The Holstein breed had the largest list of available samples for sale ($N = 165$), and the sire conception rate (in terms of predicted deviations from the average rate for the farm) was reported for 136 of them. Figure 4A plots these deviations. Although the relationship is too noisy to evaluate whether it takes the form of a power function, there is nevertheless a negative relationship between the lifetime net merit and the sire conception rate, with a Pearson correlation of $r = -.39$ (95% highest density interval [HDI]: $-.53$ to $-.25$).[5]

Academics' Gamble: Publishing in Good Journals

Where dairy farmers worry about the quality of their livestock, academics worry about where to publish their manuscripts. Publishing fast and in high-impact journals is the supposed prerequisite for a successful professional career and tenure—thus the phrase *publish or perish* (Garfield, 1996). As a result of this pressure, academics face the gamble of where to submit a manuscript for publication. Beyond the fit between the manuscript's content and the journal's scope and readership, authors tend to consider two key aspects: the relative importance of the journal and its acceptance rate. Importance is usually indexed by a journal's impact factor, defined as the ratio of the number of citations of articles published in the journal relative to the total number of articles published in it over a 2-year period (Garfield, 1955, 1972). The acceptance rate is the proportion of all submissions that are published in a given year.

Impact factors and acceptance rates map onto payoffs and probabilities, respectively. Does a negative correlation between payoffs and probabilities exist in this domain as well? Many authors suspect that journals with the largest impact factors tend to have the lowest acceptance rates, and at least in the area of ecology, this seems to be the case (Aarssen et al., 2008). In order to investigate this question further, we compiled a list of all journals associated with psychology and/or psychiatry in 2010 by either Thomson Reuters Journal Citation Reports ($N = 608$), Cabell's Directories of Publishing Opportunities ($N = 743$), or the American Psychological Association (2011; $N = 59$), resulting in in $N = 1,043$ journals. Thomson Reuters Journal Citation Reports was our primary source of impact factors and Cabell's Directories of Publishing Opportunities, our primary source of acceptance rates. Impact factor and acceptance rate were listed for 285 journals. An initial investigation of the data revealed that subdisciplines within psychology had different baselines for impact factors and acceptance rates.[6] To account for these differences, Figure 4B plots the relationship between acceptance rate and impact factor as z scores for each subdiscipline. To estimate the correlation, we modeled this relationship hierarchically, with each subdiscipline having its own mean and standard deviation of acceptance rate and impact factor, as well as correlation between the two (see Appendix A for a description of the model). The posterior estimate of the modal correlation at the level of psychology is $r = -.26$ (95% HDI: $-.38$ to $-.13$).

---

[3] Personal communication with Jeff Pleskac (August 2008).

[4] The sire conception rate is calculated by the U.S. Department of Agriculture. It measures the rate of pregnancies via artificial insemination for a bull. It is technically a deviation score controlling for various aspects of the herd.

[5] We estimated the posterior distribution over the correlation coefficient using Markov chain Monte Carlo methods (Lee & Wagenmakers, 2013).

[6] Subdisciplines were determined by the 11 subdisciplines of psychology/psychiatry journals used by Thomson Reuters Journal Citation Reports. Of the 608 journals, 84 had multiple subdisciplines. These were assigned to the subdiscipline listed first in the database, which in nearly all cases appeared to be the primary subdiscipline of the journal.
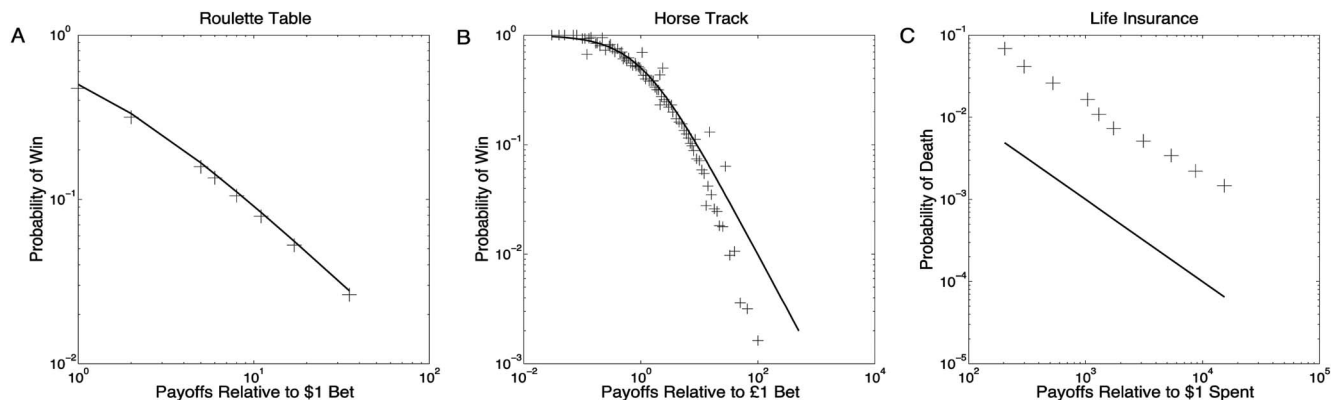
Figure 3. Relationship between payoffs and probabilities (plotted in log-log space) in three domains of life's gambles. All the payoffs have been transformed to be relative to $1 or £1 spent. The solid line is the estimated probability assuming a fair bet.

Proposers' Gamble: Suggesting a Division in the Ultimatum Game

Publishing in scientific journals is a social game in which authors compete for a scarce commodity, that is, a slot in the limited pages of a high-impact journal. A more generic social game is the *ultimatum game* that experimental economists and others have long employed to study bargaining behavior (Güth, Schmittberger, & Schwarze, 1982). In its simplest form, the ultimatum game involves two people who play a single round in which one person, the proposer, suggests how to split a fixed monetary pie. This split represents a take-it-or-leave-it offer (an ultimatum) that the other person, the responder, must accept or reject.

Normatively, according to standard economic theory, there ought to be no uncertainty in this game. That is, the self-interested payoff-maximizing responder should accept any positive payoff, no matter how small, and the self-interested payoff-maximizing proposer should allocate the smallest possible amount to the responder and pocket the lion's share. Thus, rational economic theory predicts that no relationship should emerge between the payoff and the probability of rejection. In practice, however, real responders frequently reject low offers (Camerer, 2003; Güth & Tietz, 1990). Consequently, the proposer faces a gamble in which offering too little increases the chances of the offer being rejected and offering too much means overpaying in the settlement reached. Given this tradeoff, it is possible that a negative relationship arises between risk and reward even in this social game. To find out whether this is indeed the case, we reanalyzed data from Fischbacher, Fong, and Fehr (2009), where the proposer was given 100 points (exchangeable for real money).[7] Figure 4C plots the probability of an offer being accepted as a function of the amount the proposer kept. Indeed, as the amount the proposer retains increases, the probability of the offer being accepted decreases. In other words, the risk–reward relationship exists not only in games against nature (e.g., monetary lotteries) but also in social games.

To summarize, we examined six diverse environments of risky choice—some involving probabilities known (in principle) a priori because of well-defined random mechanisms, others involving statistical probabilities and less obvious mechanisms generating

relationships between outcomes and payoffs. We consistently observed a negative relationship between probabilities and outcomes, with larger rewards (payoffs) being associated with more risk (smaller probabilities). This negative relationship appears to take the form of a power function, especially in those environments in which probabilities and outcomes can be measured without much noise and where there is a large range of payoffs.

How Does the Ecological Risk–Reward Relationship Arise?

What causes risk and rewards to be inversely tied in the environments considered? With no claim to completeness, we identify three relevant forces, of which only the first two are consistent with standard economic theory; the third illustrates that factors beyond a competitive marketplace can also independently produce the risk–reward relationship.

Fair Bets

The forces of the marketplace are sufficient to give rise to a negative relationship between probabilities and payoffs and sometimes even sufficient to produce a power function relationship between the two (as shown in Figure 3).[8] This is because decisions involving risk and uncertainty often occur in the course of an interaction between two parties: a player (e.g., a gambler) and a seller (e.g., the house). The seller has a gamble that has the prospect of gaining g if an event A occurs with a probability p_g and offers the player this gamble at a price l. Naturally, the player desires a high payoff, with the highest of probabilities and at the lowest cost. The seller has the opposite aspiration. These opposing

---

[7] We are grateful to Urs Fischbacher for this data (Fischbacher et al., 2009). In this condition of Fischbacher et al.'s (2009) experiment, a total of 46 participants made 10 offers anonymously to 10 other real participants.

[8] Many environmental magnitudes exhibit a power law relationship with their corresponding frequencies (Bak, 1999; Gutenberg & Richter, 1965; Ijiri & Simon, 1977; Mandelbrot, 1982; Zipf, 1949). However, the relationship between payoffs and probabilities in gambles is different in that it arises between, and not within, distributions.
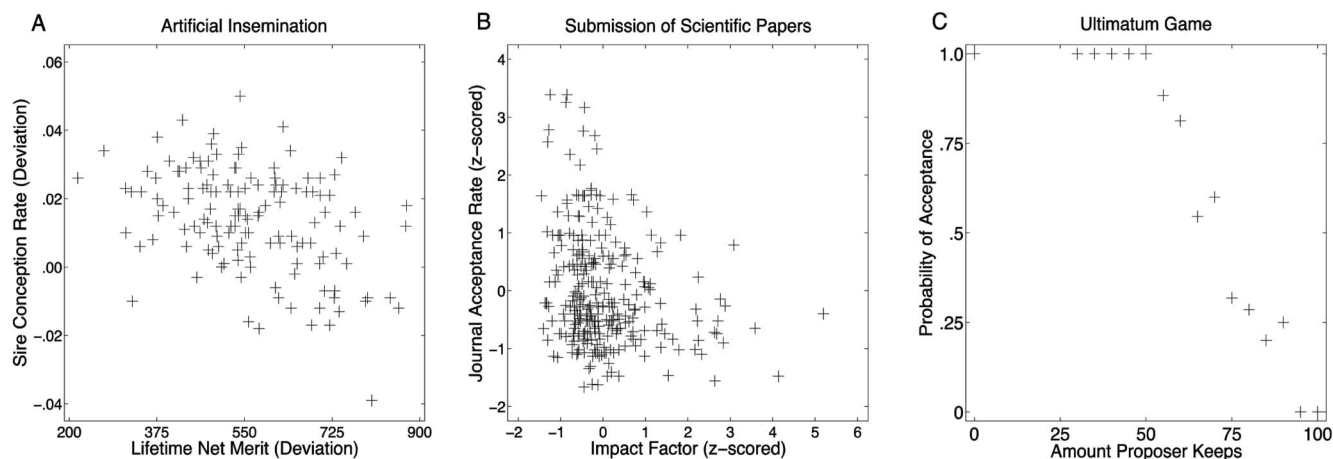
Figure 4. Relationship between payoffs and probabilities in three real-world domains beyond monetary lotteries. Panel A plots the relationship between probabilities and payoffs from gambles a dairy farmer takes when selecting semen to artificially inseminate his or her cows. Both lifetime net merit (payoff) and sire conception rate (probability) are deviation scores from the average value of the herd/farm. Panel B plots the relationship between payoffs and probabilities in a gamble a scientist in the psychological sciences faces when choosing where to submit a paper. The impact factor and acceptance rates are *z*-scored within a subdiscipline. Panel C plots the relationship between probabilities and payoffs in the ultimatum game from the perspective of the proposer (see text).

forces lead to a Pareto-efficient state such that lotteries are pushed toward a fair bet, in which expected gains correspond to expected losses such that

$$p_g \cdot g = [1 - p_g] \cdot l \qquad (1)$$

(Osherson, 1995; Skyrms, 1966). Reworking this equation, we can say that a gamble is fair if and only if

$$p_g = l/(l + g). \qquad (2)$$

The fair-bet definition is fundamental to many aspects of economic theory, including the efficient market hypothesis (Samuelson, 1965). One immediate psychological implication is that a decision maker can take advantage of the ratio $c = l/(l + g)$ to heuristically decide whether to accept a bet. Specifically, she can evaluate where her personal belief that an event will occur, sp_g, rests relative to c. If $sp_g > c$, then she should take the bet; otherwise, she should reject the bet. Indeed, poker players are instructed to use this heuristic when they estimate the chances of winning (card odds) and compare it with the ratio of the cost of a call to the total amount of possible winnings (pot odds).

A second implication of the fair-bet assumption concerns the structure of the ecology. In a competitive marketplace, gambles will be selected that offer approximately fair bets, thus resulting in a Pareto frontier between probabilities and payoffs. This frontier suggests that for gambles that offer a chance to win a single payoff for the cost of *l,* the probability of a gain p_g will be equal to the ratio of losses to the total amount being wagered:

$$p_g = [(g + l)/l]^{-1}. \qquad (3)$$

This functional relationship arises *regardless* of the true probability distribution of the event in question: Fair bets about normally distributed events, uniformly distributed events, exponentially dis-

tributed events, and so on will result in a power function relationship between payoffs and probabilities.

How closely do our choice environments correspond to this power function relationship? We plotted the predicted probability under the assumption of a fair bet for each of the binary monetary gambles in our ecological analysis (see Figure 3). Table 1 lists the mean absolute deviation between the estimated probabilities (assuming fair bets) and the implied or true probability of the event occurring. Across all three domains, the probabilities and payoffs fall surprisingly close to the predicted probabilities from a fair bet. In roulette and horse racing, the predicted probabilities of winning from a fair bet fall short of the true probabilities. This is, of course, by design. In these domains, the seller implements various adjustments that bias the gambles toward the house, and the predicted probability from a fair bet thus overestimates the chances of winning in these domains. In contrast, in the domain of life insurance, the predicted probability from a fair bet undershoots the probability of a death. Again, this is by design. The underestimation is due to the fact that employees pay only part of the premium and to the natural pooling of premiums across people helping to offset costs. Nevertheless, across all three domains, people would

Table 1
Mean Absolute Deviation (MAD) Between Observed Probability of a Win and the Estimate From a Fair Bet Across the Three Domains

Domain	MAD (*SD*)
Roulette table	.005 (.005)
Horse track	.04 (.04)
Life insurance	.017 (.020)

arrive at a relatively accurate estimate of the probability of a payoff by simply taking the ratio of losses to wins.

Filtering of Dominated Options

Although this quantitative result appears specific to monetary gambles, in qualitative terms, market forces can also be understood as responsible for the relationship observed in the domains of artificial insemination of dairy cows and publishing scientific papers. Dairy farmers, of course, do not want to purchase dominated options: semen samples that have a low sire conception rate and a low lifetime net merit. Thus, companies only sell semen samples that are high on lifetime net merit, sire conception rate, or both, removing semen samples with both low sire conception rate and low lifetime net merit from the reference class. This filtering of inferior options is sufficient to produce a negative association between the two quantities.[9] In a similar manner, the market can be understood to drive the relationship between impact factor and acceptance rate in academic journals. Journals with the largest impact factors are the most prestigious and receive most submissions. Space is limited in journals, leading to lower acceptance rates in journals with the highest impact factors.

Fairness During Bargaining

There are, however, cases in which forces other than market competition produce the risk–reward relationship. The ultimatum game is one example. According to the canonical model in economics, which considers people to be rational and self-regarding (see Camerer & Fehr, 2006), no risk–reward relationship should arise in this game. Yet it does. How, then, can it be explained? As is the case in many actual social exchanges, concerns of fairness enter the bargaining process during the ultimatum game (Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999; Rabin, 1993). Violations of fairness (here, exceedingly unequal distributions of the pie) risk punishment (i.e., rejection of the proposed allocation). Indeed, many people are willing to voluntarily incur costs to punish others' violations of fairness (de Quervain et al., 2004; Yamagishi, 1986). Because of this willingness and because proposers anticipate this threat, a potentially risk-free exchange turns into an uncertain gamble, exhibiting the same risk–reward relationship that market competition produces.

To conclude, competition in the marketplace can often give rise to a negative relationship between probabilities and payoffs and is sometimes even sufficient to produce a power function relationship between the two. However, this is not always the case. In some situations, like the ultimatum game, economic forces should give rise to no relationship at all between risk and reward, and yet a risk–reward relationship still emerges due to what appears to be social concerns of fairness.

The Risk–Reward Heuristic: A Way to Reckon With Uncertainty

Across a range of environments, we found two cues to a prospect's value—the payoffs and their probabilities—to be interrelated. This intercue relationship can be exploited when one cue becomes unavailable, as is the case when probabilities are unknown in decisions under uncertainty. This capacity is captured by

the concept of vicarious functioning: An organism can cope with cues that have become unreliable or unavailable by learning to alternate between different cues (cue substitution) through their mutual intersubstitutability (Brunswik, 1952). Thus, a person can gauge the probability of payoffs within a lottery by estimating the ratio of the cost of playing to the total amount of possible winnings. We call this shortcut the *risk–reward* heuristic, and its process can be stated as follows:

> For gambles that offer a single positive payoff and otherwise nothing, infer the probability of winning a payoff, p'_g, from the ratio of the cost of playing to the total amount of possible winnings as follows:

$$p'_g = l/(l + g).$$

The probability of not winning and being out the price l of the gamble can be estimated by taking the complement of the probability of a positive payoff, $p'_l = 1 - p'_g$. We have specified the heuristic to estimate the probability of a payoff in those decisions in which a player pays a fixed price to participate in a gamble offering the chance to win a desirable payoff and otherwise nothing. Many decisions like the ones in our ecological analyses take this structure. Of course, there are also other types of gambles, such as those dealing only with losses or with multiple payoffs. For now, we focus on gambles that primarily have a pay-to-play structure with the opportunity of a single possible payoff. Later in the discussion, we address how the heuristic may be applied to other types of gambles.

From Inference to Choice

The risk–reward heuristic is agnostic regarding the further processing of the inferred probability. It does not specify a choice mechanism that spells out how the inferred probability of winning the bet (p'_g) is used in making a choice. In principle, however, the same processes are feasible that have been proposed in decision making under risk, when probabilities and outcomes are known. For instance, the decision maker could calculate the expected payoff of an option offering the payoff g (and 0 otherwise) by multiplying g with p'_g and could then maximize (expected value theory). Similarly, the inferred probability could be entered in an expected utility model that multiplies the subjective value of the payoffs $u(g)$ with p'_g. Alternatively, it could be subjected to some form of distortion via a probability weighting function (as in the two-stage model of decisions under uncertainty; Fox & Tversky, 1998), or it could represent one input into a choice heuristic (Brandstätter, Gigerenzer, & Hertwig, 2006; Payne, Bettman, & Johnson, 1993; Tversky, 1969).

Alternative Conceptions of the Risk–Reward Relationship

The risk–reward heuristic envisions that when faced with choice under uncertainty people infer that the probability of an event is

---

[9] Consistent with this explanation, the correlation between sire conception rate and lifetime net merit in the sample ($N = 1,905$) of all possible Holstein bull semen specimens from which companies take a subsample is merely $r = -.13$ (P. Vanraden, personal communication, September 5, 2012).

negatively related with the magnitude of the payoffs. This view conflicts with other hypotheses about the relationship between these two variables during decision making. For instance, according to subjective expected utility theory (Savage, 1954)—the normative account of how people ought to make these decisions—payoffs and probabilities are two *independent* factors that determine the value of an alternative and, ultimately, choice. That is, the utility of an alternative that yields outcome *x* if the event *A* occurs otherwise 0, (*x*, *A*), is

$$u(x, A) = p(A) \cdot u(x) + p(\sim A) \cdot u(0) = p(A) \cdot u(x), \quad (4)$$

where event *A* is a subset of possible states of the world *S*, $A \subset S$. The *u* is the utility function describing the subjective value of those consequences. The *p* is a probability measure on the state space *S* and reflects the decision maker's subjective beliefs about the likelihood of different states of the world occurring. However, note that the probability is based on the event only and not on the consequence of the event. Consequently, in subjective utility theory, payoffs and probabilities are ultimately compensating but not interacting factors in determining the value of the alternative. Thus, if subjective expected utility theory is taken at first approximation as a descriptive theory of choice, then a consequence of this independence assumption is that the probabilities people use to make decisions under uncertainty must be estimated independently of the magnitudes of the payoffs.[10]

Still another view of the risk–reward relationship is implied by the desirability bias. According to this bias, when payoffs become more desirable, they are perceived as more likely than less desirable outcomes (Bar-Hillel & Budescu, 1995; Edwards, 1962; Irwin, 1953; Krizan & Windschitl, 2007; Windschitl, Smith, Rose, & Krizan, 2010). This bias, however, is typically thought of as a cognitive deficiency causing the desire for an outcome to inflate people's optimistic belief that good things coincide with attractive probabilities. The risk–reward heuristic posits the opposite dependency. Moreover, the root of the inverse relationship does not reside in people's pessimism but in the ecological structure of gambles in the world.

Next, we examine these conflicting conceptions about how the payoffs affect people's beliefs during choice under uncertainty. We first test people's estimates in lotteries involving actual monetary consequences before exploring in a second study whether they also apply their ecological intuitions to hypothetical lotteries. Finally, we explore the extent to which the heuristic can shed new light on choices under ambiguity.

Study 1: Do People Exploit the Ecological Correlation Between Payoffs and Probabilities?

To examine the extent to which people exploit the ecological correlation between payoffs and probabilities, we offered participants the opportunity to play a gamble for real stakes, with the chance of winning a prize. In order to play the gamble, they had to pay a fee. They were not told what the chances of winning the respective prizes were. Left ignorant of the chances of winning, they could potentially employ the risk–reward heuristic, according to which the estimated number of winning chances should decrease as the value of the payoff increases. Alternatively, their choice could be driven by the desirability bias, according to which

the estimated number of winning chances should increase as payoff magnitude increases, or they could strive to keep their estimated probabilities independent of the payoffs, as subjective expected utility theory would predict.

Method

Participants. One hundred and thirty-eight students from the undergraduate research participant pool of the psychology department of Michigan State University took part in the study for course credit. They could also earn an additional but unstated amount of money based on their performance. After giving their informed consent, participants first searched for answers to difficult trivia questions on the Internet. They were informed that if they answered 20 questions correctly, they would earn $2. The task was difficult, but everyone earned $2. After receiving the money, they were informed that they could now play a gamble for real money.

Design and procedure. The gamble offered participants the opportunity, at the expense of a $2 entry fee, to play a gamble with the chance of winning $2.50, $4, $10, or $20 (between participants), respectively. The setup involved an envelope and a bingo basket. Participants learned that if they chose to play, they would first open the envelope and find out how many of the 100 marbles (numbered 1 to 100) in the bingo basket would yield a winning outcome. The number in the envelope was between 1 and 100. After finding out this number, they would draw a marble, and if its number was less than or equal to the number in the envelope, the prize was theirs. The gambles were constructed—unbeknownst to participants—to be fair, that is, the expected value of the gamble matched the entry price of $2. Consequently, the number in the envelope was, on average, 80, 50, 20, and 10 (winning marbles) for the $2.50, $4, $10, and $20 lottery conditions, respectively. The actual criterion value was determined by drawing the number from a random uniform distribution with a width of 10 centered on each fair value using a 10-sided die. We went to great effort to ensure that participants were not aware of other prize conditions, including only running one condition at a time (among the four possible participants that could be run at a time) and asking participants not to disclose the nature of the gambles they completed.

Before choosing whether to play the gamble, participants estimated which number between 1 and 100 they expected to find in the envelope (responses of 0 were permitted) and entered that number on the computer. Next, they rated on a 100-point scale anchored with *not at all* and *absolutely* their belief that each five-unit interval (1 to 5, 6 to 10, . . . , 96 to 100) contained the number. Finally, participants had the opportunity to play the gamble. If they chose to do so, they paid $2 and completed the gamble; otherwise, they left with the $2.

---

[10] It is important to emphasize that Savage's (1954) subjective expected utility theory is a theory of choice. Utilities and probabilities are derived from preferences over acts. It does not explicitly state how probabilities are to be calculated. It does assume payoffs and probabilities are two independent constructs that determine the value of the construct. For this reason, we have stated the independence prediction—that probabilities be estimated independently from the magnitude of the payoffs—as a consequence that follows from the theory.

Results

Consistent with the predictions of the risk–reward heuristic, the estimated number of winning marbles in the basket decreased as the dollar value of the gamble increased (see Figure 5).[11] This finding was supported by a beta regression with the values of the prizes as predictors of the estimated number in the envelope (winning probability). In the regression, the mean of the beta distribution was a linear function of the dollar value condition via a logit link, making the regression coefficients interpretable in terms of changes in subjective odds (Ferrari & Cribari-Neto, 2004; Smithson & Verkuilen, 2006).[12] We used Bayesian analysis methods to estimate the posterior distribution over the model parameters. The distributions were estimated with Markov chain Monte Carlo methods (see Appendix A). With vague priors, the modal posterior estimate for the regression coefficient on the dollar value was $b = -0.025$ (95% HDI: -0.053 to -0.002). In terms of odds of winning, for every $10 decrease in the payoff of the gamble, participants were 1.28 (95% HDI: 1.70 to 1.02) times more likely to believe they could win.

As mentioned earlier, the risk–reward heuristic is not a choice heuristic. A choice rule is therefore needed to make predictions about people's choices. One straightforward rule is that participants use the estimate from the risk–reward heuristic to calculate the expected value of the lottery and then choose the option with the larger expected value. This choice rule combined with the risk–reward heuristic predicts indifference between paying to play the lottery or not across all prize conditions. However, the choice proportions (see Table 2) reveal an overall propensity to play the gamble that increases with dollar outcome value. This conclusion is supported by a Bayesian logistic regression analysis of the choice to pay to play using the respective dollar value of the condition as a predictor (Kruschke, 2010). With vague priors, the modal posterior estimate of the intercept was $b_0 = 0.21$ (95% HDI: $-.43$ to $.81$), indicating a bias to play the gamble. Choice proportions also suggest that the preference to play increased with payoff condition. Indeed, the modal posterior estimate of the regression coefficient for the dollar value was $b = 0.097$ (95% HDI: 0.03 to 0.17), indicating that with a $10 increase
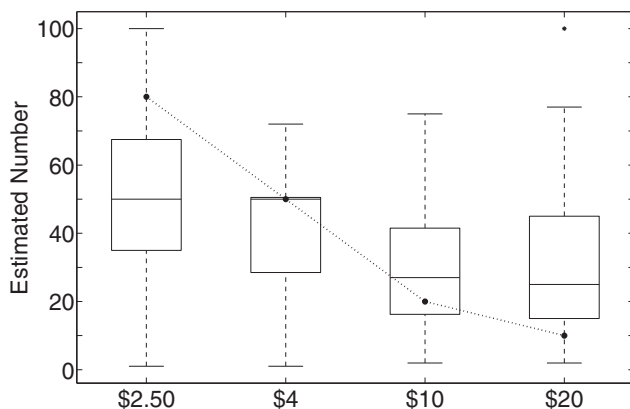
Table 2

Percentage of Participants Choosing to Play the Lottery in the Bingo Basket Study With Real Monetary Consequences

Condition	N	Percentage who gambled
$2.50	33	58%
$4	36	61%
$10	35	86%
$20	34	85%
Average		73%

Note. The study was implemented in the laboratory and involved real monetary consequences (participants paid $2 for the right to play and received the prize if they won).

in the payoff, participants were 2.65 times (95% HDI: 1.36 to 5.48) more likely to play the gamble.

Participants' estimates of the number of winning marbles impacted their choice to play the gamble. This can be seen at the aggregate and individual levels. On the aggregate level, the growing appetite to play across the conditions was consistent with the regressive pattern in the estimates of the number of winning marbles, relative to the corresponding probabilities for a fair bet. Specifically, participants' estimates of the chances of winning were lower than the fair-bet probabilities for the low-payoff conditions, and they were also less likely to play the gamble. At the same time, participants' estimates were higher than the fair-bet probabilities for the high-payoff conditions, and they were more likely to play the gamble. At the individual level, within each condition, people who estimated a lower number were less likely to accept the lottery. This was confirmed by a logistic regression using the numerical prediction and the payoff prediction as simultaneous predictors. The HDI for both coefficients did not contain 0. The modal posterior distribution estimate of the coefficient for the numerical prediction was $b_{pre} = 0.021$ (95% HDI: 0.001 to 0.045). The estimate of the coefficient for the payoff amount was $b_{pay} = 0.13$ (95% HDI: 0.05 to 0.21). Thus, even when the payoff condition was controlled, participants' judgments of the chances of winning were predictive of their choice to play.

Discussion

Consistent with the risk–reward heuristic, participants' estimates of the chances of winning decreased as the size of the payoff increased. Furthermore, their choices of whether to play were predicted by those estimates (and the prizes). This pattern of results implies that not only did the magnitude of the lottery payoff inform people's inferences about the probability of winning but people's inferred probabilities also impacted their preferences. These findings are inconsistent with the hypothesis derived from Savage's (1954) subjective expected utility theory that the probabilities are estimated independent of the payoffs. They are also inconsistent with the desirability bias's prediction that the per-

Figure 5. Box and whisker plots of the estimated number of winning marbles in the bingo basket (see text). The dotted line plots the estimates predicted by the risk–reward heuristic.

[11] The estimates from the rating method proved highly correlated with the single point estimates and thus are not reported here.

[12] Following Smithson and Verkuilen (2006), we rescaled the 0 to 100 estimates to be on the 0 to 1.0 interval, thus effectively shrinking the interval to a .0001 to .9999 interval and so avoiding zeros and ones.

ceived likelihood of an event occurring will increase as its outcome is made more desirable.

The observed estimates did deviate from those of fair bets, the latter being what an error-free variant of the risk–reward heuristic would suggest. In particular, as Figure 5 shows, the estimates were also regressive relative to the predictions of the risk–reward heuristic (dotted line): They overshot the probabilities corresponding to a fair bet for the high prizes and undershot the probabilities for the low prizes. One possible explanation for this deviation is that decision makers' use of the risk–reward heuristic—like other judgments—is not error free but perturbed by error from other sources of information (Erev, Wallsten, & Budescu, 1994; Hertwig, Pachur, & Kurzenhäuser, 2005; Juslin, Olsson, & Bjoerkman, 1997). This would be sufficient to produce the regressive effect observed in the estimated number of winning marbles.

Study 2: Do People Deem Payoffs and Probabilities to Be Dependent Even in Hypothetical Bets?

Hypothetical gambles are frequently used in research on behavioral decision making, and it is argued that they are equally as valid a measure of preference as lotteries involving real monetary consequences. Moreover, they afford the experimenter control over payoffs and probabilities (Kahneman & Tversky, 1979) and permit inexpensive, fast, and convenient studies (Thaler, 1987). Do people also apply their intuitions about the statistical structure of lotteries to bets that, unlike those in Study 1, have no actual monetary consequences? If this is the case, it will call into question the assumption that probabilities are processed independently from payoffs in lotteries that are constrained by market forces.

Method

Participants. One hundred and ninety-eight volunteers on Amazon Mechanical Turk completed a short survey on decision making. The sample included 81 males (41%), and the average age was 32 years ($SD = 11.5$). The modal level of education was a bachelor's degree (41%), and the modal description of employment status was a wage earner (45%). Participants were paid $0.15 for completing a 10-question survey that took just under 5 minutes. Two respondents failed to complete the survey and were excluded.

Design and procedure. Initial pilot studies suggested that the bingo basket paradigm used previously was too complicated to communicate in an online format. Thus, we presented the following simplified scenario:

> Imagine you have been asked to play the following lottery. The lottery offers the opportunity to win $X, but it costs you $2 to play. If you choose to play you would pay the $2 and, without looking, draw a ball from a basket. In the basket there are 100 balls. The balls are either black or red. If the ball is red you will win $X; otherwise, if the ball is black you will receive nothing. Thus, the number of red balls in the basket determines the probability that you will win. You are less likely to win the $X the lower the number of red balls in the basket.

Between participants, we set the dollar value to $X = $2.50, $4, $10, $50, or $100, respectively. Participants first estimated the number of red balls (out of 100) in the basket. We then asked whether they would pay $2 to play the gamble for $X if they had the opportunity.

Results

As in Study 1, the estimated number of winning red balls in the basket decreased as payoff increased (see Figure 6). A Bayesian beta regression model supported this conclusion. With vague priors, the modal posterior estimate for the regression coefficient on the dollar value was $b = -0.007$ (95% HDI: -0.012 to -0.004). In terms of their odds of winning, for every $10 decrease in the payoff of the gamble, participants were approximately 1.08 (95% HDI: 1.12 to 1.03) times more likely to believe they could win. As expected, these were smaller effects than those obtained in Study 1, conducted in the laboratory and involving actual monetary consequences. The distributions of the estimated number of winning red balls also revealed even greater skepticism relative to the previous study. For example, the median estimates for the $2.50, $4, and $10 conditions were 20%, 50%, and 39%, respectively, lower than in the previous study.[13]

The choice proportions, listed in Table 3, were also substantially lower than in Study 1. Participants in the online study showed no apparent preference to play the gamble. If anything, they showed the opposite inclination. This was confirmed by a Bayesian logistic regression of the choice to play with the payoff as predictor. The modal posterior estimate of the intercept term was negative, $b_0 = -1.23$ (95% HDI: -1.67 to -0.81), indicating no bias to play the gamble. As before and inconsistent with the assumption of risk neutrality, the likelihood of participants choosing to play the gamble increased with the size of the prize, $b_{pay} = 0.017$ (95% HDI: 0.009 to 0.026; see Table 3). In terms of odds, these estimates imply that with a $10 increase in the payoff, participants were 1.18 times (95% HDI: 1.10 to 1.29) more likely to play the gamble.

Similar to Study 1, participants' estimates of the number of winning marbles appeared to shape their choice to play the gamble. At the aggregate level, people's estimates overshot the fair probabilities for high payoffs and undershot those for low payoffs (see Figure 6). Finally, at the individual level, participants who estimated a lower number were again less likely to play the gamble. This was confirmed by a Bayesian logistic regression analysis of the choice to play with the payoff and the numerical estimates for the number of winning red balls as simultaneous predictors. The HDI for both coefficients did not contain 0. The modal posterior estimate of the coefficient for the payoff amount was $b_{pay} = 0.02$ (95% HDI: 0.01 to 0.03), and the coefficient for the numerical prediction was $b_{pre} = 0.02$ (95% HDI: 0.002 to 0.03).

Discussion

Even when facing simple lotteries with hypothetical consequences, participants use the size of the prize to infer the probability of winning. The evidence emerging from Studies 1 and 2 does not support a strong (error-free) form of the risk–reward heuristic. Nevertheless, the observed negative correlations (see Figures 5 and 6) between payoffs and inferred probabilities sug-

---

[13] Comparison of the intercepts of the beta regressions from Studies 1 and 2 reveals the same pattern. In Study 2, $b_0 = -0.0860$ (95% HDI: -0.395 to 0.173), whereas in Study 2 with hypothetical gambles, $b_0 = -0.704$ (95% HDI: -0.904 to -0.499).
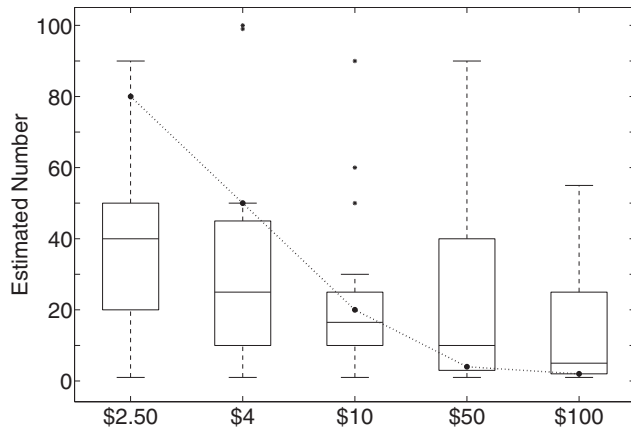
Figure 6. Box and whisker plots of the estimated number of winning red balls in the basket in the online study. The dotted line plots the estimates predicted by the risk–reward heuristic.

gest that the risk–reward heuristic strongly anchors people's regressed estimates.

Study 3: Ambiguity Aversion—A Possible Consequence of Using the Risk–Reward Heuristic?

One interesting consequence of exploiting ecological structures is that the risk–reward heuristic can make its users appear averse to ambiguity. Ambiguity aversion refers to the phenomenon that most people given a choice between two otherwise equivalent options—one in which the probability information is given and the other in which it is missing—avoid the option with missing probability information (Camerer & Weber, 1992). Consider, for example, the choice between these two lotteries:

Lottery A: Win $100 if a red marble is drawn from an urn with 50 red marbles and 50 black marbles.

Lottery B: Win $100 if a red marble is drawn from an urn with 100 marbles with an unknown proportion of red and black marbles.

Presented with such choices, people generally prefer A to B (Camerer & Weber, 1992; Ellsberg, 1961; Fox & Tversky, 1995; Rode, Cosmides, Hell, & Tooby, 1999; Yates & Zukowsky, 1976). This preference is commonly interpreted in terms of ambiguity aversion, based on the assumption that people are indifferent to betting on red or black in either lottery. Consequently, they should be indifferent between the two lotteries. The preference for *A* thus suggests that, as Ellsberg (1961) argued, uncertainty cannot simply be reduced to a probability, as subjective expected utility theory would assume (Savage, 1954).

Why do people behave as if they were averse to ambiguity? Numerous explanations have been proposed. These include the role of (comparative) ignorance (Fox & Tversky, 1995), an evolutionary predisposition to avoid variance (Rode et al., 1999), the probability estimation process (Curley & Yates, 1989; Einhorn & Hogarth, 1985; Trimmer et al., 2011), and suspicion that the ambiguous urn is rigged (Binmore, 2007; Frisch & Baron, 1988; Kühberger & Perner, 2003; Yates & Zukowsky, 1976). These last two explanations—probability estimation and suspicion—are of

particular interest in the context of our ecological analysis, which suggests that the belief in a skewed urn may reflect an ecologically rational intuition about the structure of lotteries in the world and not some internal bias. That is, can the risk–reward heuristic and its exploitation of an ecological regularity account for *some* of the ambiguity aversion effect?

Method

One hundred and twelve students from the University of Basel were randomly assigned to two groups and were presented with the aforementioned (hypothetical) choice between lotteries *A* and *B*. The first group ($n = 60$) was told to assume that the payoff for drawing a red marble from either urn was 10 Swiss francs. The second group ($n = 52$) was told to assume that the payoff for drawing a red marble was 1,000 Swiss francs.

Results and Discussion

The risk–reward heuristic predicts that the ambiguous or uncertain option will become increasingly undesirable as the magnitude of the payoff increases. The reason for this prediction is again the negative relationship between payoffs and probabilities: With increasing magnitudes of payoffs, their probability decreases (relative to indifference priors of .5). Consequently, the larger the magnitude of the payoff (and the smaller its probability), the more people can be expected to prefer the unambiguous urn. Because the relative variance between the two lotteries remains constant, the comparative ignorance that participants may experience (Fox & Tversky, 1995) or their inclination to avoid variance (Rode et al., 1999) should remain unchanged between the two groups. Consequently, the preference for the unambiguous urn should remain constant across the two groups.

The results support the risk–reward prediction. In the 10 Swiss francs condition, 63% (38/60) preferred the unambiguous urn with a known 50:50 composition. In the 1,000 Swiss francs condition, by contrast, 81% (42/52) preferred the unambiguous urn. With vague priors, the modal estimate of the posterior difference between the estimates was .17 (95% HDI: .01 to .33). This change in preference between 10 and 1,000 Swiss francs is consistent with the notion that choices are based on intuitions that exploit payoffs as proxies for probabilities. Indeed, Ward Edwards appears to have anticipated such a partial explanation of ambiguity aversion: "High positive or negative payoffs are typically associated with very low probabilities, while mediocre or zero payoffs are typically associated with high probabilities. This, I assert, is simply a fact about

Table 3

Percentage of Participants Choosing to Play the Lottery in the Online Bingo Basket Study With Hypothetical Monetary Consequences

Condition	N	Percentage who would gamble
$2.50	43	14%
$4	38	29%
$10	38	27%
$50	38	53%
$100	39	56%
Average		35%

the world in which we live" (personal communication quoted in Lopes, 1983, p. 141). Our ecological analyses indeed corroborate Edwards's assertion.[14]

Accounting for people's preference for the unambiguous urn in terms of ecologically rational intuitions is not meant to supplant the notion of ambiguity aversion. The point though is that ecological structures and cognitive processes can and should be viewed as working in tandem (see also Anderson, 1998; Brunswik, 1952, 1955b; Lopes, 1983; Simon, 1956). To this end, our ecological analyses and the risk–reward heuristic complement the estimation accounts of ambiguity aversion. They posit that, when faced with the choice between lotteries *A* and *B* (see above), people gather information to form an estimate of the probability of winning and arrive at an estimate of less than .5. This is done via either an asymmetrical adjustment downward from an anchor of .5 (Einhorn & Hogarth, 1985) or an application of Bayes's rule with prior beliefs that treat probabilities less than .5 as more credible than those above (Trimmer et al., 2011; Viscusi, 1989). Why these properties exist is less clear. The risk–reward heuristic offers a principled explanation of why, by suggesting that aversion to ambiguity is (at least partly) due to an ecologically grounded belief that high payoffs have a low probability of occurring.

General Discussion

How people reckon with uncertainty and make choices when the probabilities of outcomes are not known has been a vexing problem for choice theories. We suggest that people navigate the fog of uncertainty using, among other mental tools, a heuristic that infers the probabilities from the known payoffs. We conducted ecological analyses in diverse real-world environments and showed that a robust environmental relationship exists between risk and reward. Three empirical studies demonstrated different ways in which this relationship informs people's behavior when they choose among monetary gambles in the gain domain. We now discuss our adaptive view of choice under uncertainty and review how and where this relationship between payoffs and probabilities may arise in other domains.

How Generalizable Are the Risk–Reward Relationship and Its Consequences?

Like any heuristic, the risk–reward heuristic is not a domain-general tool. Its smart use is limited to domains in which there is an inverse relationship between payoffs and probabilities. Accordingly, it should not be used in environments with a positive relationship between valued outcomes and their relative frequency, where its application would lead decision makers astray. Such a relationship may be found in the social domain, especially in areas where what is the best return for the public is not identical with the best return for the individual. For example, the public is better off if individuals do not rob banks, run red lights, or evade taxes. The individual, in contrast, may be better off doing any of these things. As a result, formal (e.g., police, laws) and informal social institutions (e.g., norms, conventions) aim to enforce the collectively desirable behavior. Consequently, this collectively desirable behavior can be expected to be more common than the behavior that may maximize (some) individual self-interests.

Repeated risk taking and the risk–reward relationship. Our analysis of the ultimatum game (see Figure 4C) suggests that

market competition is not the only force underlying the risk–reward relationship. Figure 7 offers another illustration of this fact. It plots the relationship between female college students' number of sexual partners (1, 2–4, ≥5; self-reported) and percentage of those women without a sexually transmitted disease (STD; e.g., chlamydial infection, gonorrhea, syphilis; self-reported). Women with five or more sexual partners were 8 times more likely to have one or more STDs than were women with just one partner (Joffe et al., 1992). Equating a higher number of sexual partners with higher reward—obviously a controversial notion—and the risk of STD with the price for the chance of reaping this return, then another strong risk–reward relationship emerges.

This time, however, it is not market competition or social norms that produce it, but the forces of chance. Let us view the sexual encounter with a new partner as a Bernoulli trial with two possible outcomes, success (no STD infection) and failure (STD infection). Each new partner represents an independent trial. Probability of success k is then p^k, where p is the probability of a success on any given trial. An exponential relationship between risk and reward is thus produced: The higher the reward (number of partners), the lower the chance to escape an infection. This example illustrates that for risky behaviors that people can engage in repeatedly—from having unprotected sex to using illegal drugs—a negative relationship may be brought about between probabilities and payoffs as a result of the stochastics alone.

Losses. Does the inverse relationship between probabilities and the magnitude of outcomes extend to situations in which people incur costs to prevent or offset losses? We suggest that the answer is yes, especially in domains shaped by market competition. For instance, the failure rate within a product domain is likely to be inversely related to the magnitude of the price. Yet there is not always a straightforward rarity–severity relationship in the loss domain. In the 19th and early 20th centuries, tens of millions were affected by smallpox, and at least one fifth of those infected died (Fenner, Henderson, Arita, Jezek, & Ladnyi, 1988). Smallpox was a widespread and often lethal disease. Thus, as in the gain domain, the ability to infer the probability from the magnitude of the loss is domain specific. Moreover, other factors may shape the inference process in the area of losses, such as a motivation to overestimate severely negative events as a measure to prevent costly errors and experiences of severe loss (Harris, Corner, & Hahn, 2009; Weber, 1994).

Multiple-outcome gambles. Does a risk–reward relationship also arise for gambles with more than one possible payoff, such as lottery tickets? We analyzed all 32 scratch-off lottery tickets

---

[14] Edwards made this assertion in trying to explain a consequence of ambiguity aversion known as Ellsberg's paradox. One version of this paradox asks people who have chosen between lotteries *A* and *B* (see text) to choose again between *A* and *B* using the same urns, but now betting on a black marble (rather than a red one) being drawn. People tend to choose *A* again (the unambiguous urn), which is inconsistent with subjective expected utility theory. Edwards suggested that an equivalent of the risk–reward heuristic could explain this result if respondents ignored the instruction that only the payoffs changed and not the urns themselves. Our tests of this conjecture suggest this is not the case. People generally believe the purely verbal manipulation and still choose *A* both times. Moreover, this pattern of choices proves largely independent of the dollar magnitude. This suggests that the risk–reward heuristic can explain some aspects of ambiguity aversion but not all of people's behavior in the Ellsberg paradox.
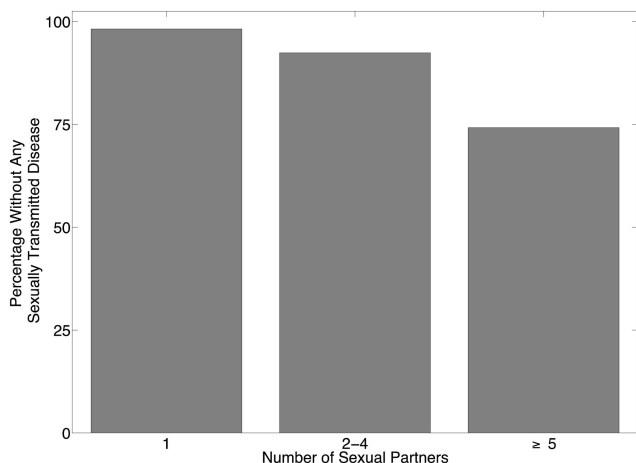
Figure 7. Percentage of women without any sexually transmitted disease during a 3.5-year period. Data come from Joffe et al. (1992).

available in the state of Michigan in March 2008. The average and median number of prizes per ticket was 21 ($SD = 7$), and the average expected loss per $1 spent was $0.36 ($SD = 0.08$).[15] Figure 8A plots the relationship between probabilities associated with each payoff (relative to $1 spent). The relationship (both between and within tickets) takes an approximately linear form in log-log space, again suggesting a power function relationship between probabilities and payoffs.

Across the lotteries, the probabilities, much like their counterparts in the binary gambles, fall close to a Pareto frontier defined by a fair bet. Calculating these fair-bet probabilities is by no means as straightforward as in the previous ecologies (see Figure 3): In multiple-outcome gambles, there are, by definition, degrees of freedom in the distribution of probabilities. Nevertheless, Appendix B describes how—using the stated outcomes and an estimate of the expected winnings—probabilities can be inferred for gambles like state lotteries that offer a chance to win a payoff from a range of possible gains at fixed cost. On average, the true probabilities deviated from these estimates of the fair-bet probabilities by .017 ($SD = .03$). Figure 8 shows the relationship between the true and fair-bet probabilities for the set of lottery tickets with the smallest (see Figure 8B), median (see Figure 8C), and largest (see Figure 8D) mean absolute deviations. In general, as should be expected given that the lotteries are designed to make a profit, the true probabilities fell short of the fair-bet probabilities. The deviations, as the figures show, tended to be largest for the payoffs with the highest magnitudes.

Nevertheless, the close correspondence between the true probabilities and the fair-bet probabilities raises the question whether people use a similar process to estimate the probabilities of multiple-outcome gambles and use them to make a choice among lottery tickets. It is not impossible that people use some expanded version of the risk–reward heuristic to estimate the probabilities of each of the possible payoffs. However, cognitively, the process seems quite taxing. Alternatively, people in multiple-outcome gambles may simplify the choice situation and focus only on the overall probability of winning and losing, respectively, using their personal estimates of the expected gains to estimate these probabilities and to make a decision.

Risk as variance. The risk-as-probabilities approach is a common definition in the area of discrete monetary gambles. Alternatively, risk can be operationalized as a function of variance—a definition that is often used in research on risk perception (e.g., Luce & Weber, 1986), finance (e.g., Sharpe, 1964), and in areas dealing with near-continuous distributions of possible outcomes. Is the risk–reward relationship limited to risk defined in terms of probability? For an answer, we can turn to theorizing in finance. Both the efficient frontier in Markowitz's (1952) portfolio theory and the capital market line (Sharpe, 1964) depict an inverse relationship between risk and expected return in a risk–return space. Furthermore, the assumption is that in selecting among different assets, investors trade off greater variability for greater expected returns. As a result, across different asset portfolios, a positive relationship can be predicted between the variability of a portfolio and the average return (or a negative risk–return relationship). Indeed, there is some support for this prediction in the performance of mutual funds (Sharpe, 1966; but see Fama & French, 1992). An adaptive view of cognition, in turn, raises the interesting question of how choice processes might have adapted to exploit this version of the risk–reward relationship. We leave this question to future research, but it opens the door for a better understanding of decision making outside the laboratory.

Ecological Rationality and Other Strategies to Reckon With Uncertainty

The risk–reward heuristic is an *ecologically rational* strategy insofar as it exploits the statistical structure of gambling environments to substitute for missing probability knowledge (Todd et al., 2012). To appreciate this aspect, it is useful to review how other strategies for making choices under uncertainty handle the missing probability information. Subjective expected utility theory implies that a decision maker's subjective beliefs concerning the event in question are formed independently of the payoffs (Savage, 1954). This independence assumption is reflected in more descriptively plausible accounts of decision making under uncertainty whether the probabilities are estimated—without reference to the associated payoffs—from a sample of possible outcomes (Hertwig & Erev, 2009) or from a search of declarative or episodic knowledge stored in memory (Fox & Tversky, 1998; Tversky & Fox, 1995). These solutions waste an opportunity to arrive at an informed estimate of the likelihood of obtaining a payoff by exploiting the information encapsulated in the payoffs themselves.

Other strategies for making a decision in the face of uncertainty forgo probabilities altogether and instead make a decision based solely on the possible outcomes. Yet these strategies can also be understood as making a bet about the structure of the environment. The maximin rule (Savage, 1954), for instance, identifies the lowest potential payoff across the alternatives available and chooses the alternative whose lowest payoff is highest (thus maximizing the minimum payoff). Relatedly, the minimax rule (Savage, 1951) also focuses on the worst possible case but minimizes the maximum amount of regret a person would feel (regret defined as the difference between the actual payoff and the maximum payoff possible if the true state of the world were known). These

---

[15] When the prize was something other than money (e.g., a car), we used its cash value.
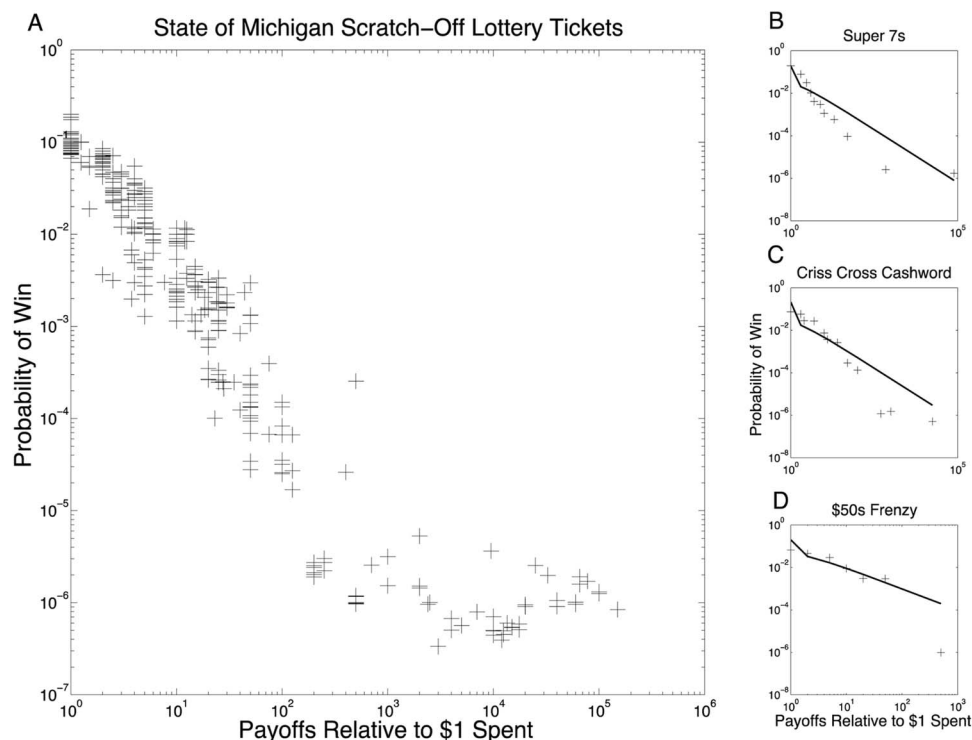
Figure 8. The relationship between payoffs and probabilities across all 32 scratch-off tickets for state lotteries (Panel A). Panels B–D illustrate the relationship between the estimates of the fair-bet probabilities and the true probabilities (see Appendix B). The Super 7s ticket has the smallest mean absolute deviation between the estimated probabilities and the true probabilities (.010; Panel B), Criss Cross Cashword has the median mean absolute deviation (.018; Panel C), and $50s Frenzy has the largest mean absolute deviation due to overestimation of the probability of the largest payoff (.023; Panel D).

pessimistic rules are optimal in hostile decision environments where the worst that could happen will always happen. Alternatively, decision makers might take an optimistic stance and choose as if the best that can happen will always happen, maximax. Or, as Hurwicz (1951) suggested, they could form a weighted average of the minimum and maximum payoff for each alternative, thus forming a pessimism–optimism index, and then choose the option with the largest index. This approach assumes that the unknown probabilities lie between the worst and best possible distributions (Luce & Raiffa, 1957).

These strategies are indeed possible ways to make decisions under uncertainty. However, as we have demonstrated, gambles in the world often tend not to be composed of the least favorable probability distribution (as assumed by maximin or minimax) or the most favorable probability distribution (as assumed by maximax). Instead, social and economic forces push the probability distributions toward a Pareto-efficient frontier, thus permitting ecologically smart choices based on inferences from the risk–reward heuristic.

How Representative Are Laboratory Gambles?

According to our analyses, payoffs and probabilities are not independent. This finding also pertains to the assumption that monetary lotteries can stand in for all of life's risky choices. Kahneman and Tversky (2000) summarized such a notion, stating,

"As with the fruit fly, we study gambles in the hope that the principles that govern the simple case will extend in recognizable form to complex situations" (p. xi). One way to evaluate this implicit or explicit hope has been to test predictions from the theories developed around monetary gambles in the field with real-world situations (for a review, see Camerer, 2000). Expected utility theory, for instance, has grown into such a powerful theory because it can account for a wide range of real-world phenomena. Prospect theory has some of the same abilities. It appears to predict actual horse-racing bets (using the very gambles we analyzed earlier; Jullien & Salanie, 2000). Prospect theory also offers an explanation for macrolevel phenomena such as the equity premium puzzle (Benartzi & Thaler, 1995) or the purchasing of lottery tickets (Clotfelter & Cook, 1989).

Our concern, though, is not one of fit. No doubt, with the right parameters and additional auxiliary assumptions, existing theories of risky choice can account for a wide range of real-world phenomena. Our concern is that the ecological structure of the monetary gambles may not be adequately represented in the laboratory. Brunswik (1952, 1955b) articulated the problem more generally. In his view, psychology's accepted experimental method of systematic design, in which variables are often artificially untied, creates stimuli that leave no room for organisms to exploit the ecological structure they have adapted to. As a result, the generalizability of the findings will be limited to the experimental

situation itself. Instead, Brunswik advocated the method of *representative design,* in which an experimenter samples stimuli—much like participants—representatively (e.g., randomly) on a number of dimensions (e.g., cue validities, intercorrelations) from a defined population of stimuli (Dhami et al., 2004).

It can, however, be difficult to evaluate the degree to which studies in the laboratory are representative of the ecologies outside it. Our analyses suggest a benchmark against which the representativeness of monetary gambles in the laboratory can be evaluated, namely, the degree to which they embody the negative correlation between payoffs and probabilities. We examined this benchmark across numerous laboratory gambling environments, including the extensive set of studies compiled in a meta-analysis of risky choice (Weber et al., 2004). Across gambles, the structures in the world are very different from the risk–reward structures realized in the gambles representing risky choices in laboratories. Figure 9 shows just one representative set of environments (many other published studies exhibit the same pattern) used to test prominent theories of risky decision making. The relationships between payoffs and probabilities between gambles depart greatly from the environments in Figures 3 and 4 and reflect experimenters' use of systematic and factorial manipulations.

One difference between the laboratory ecologies and our sample of real-world ecologies is the format of choice. All of the real-world gamble environments that we investigated involve an economic exchange in which a person pays a fixed amount of money for the right to play the gamble. None of the laboratory lotteries assume this pay-to-play structure. Instead, researchers commonly ask participants to choose between two or more lotteries or to judge the certainty equivalent of a particular gamble (though see Juni, Gureckis, & Maloney, 2012). Doing so can be understood as treating each decision problem as a separate ecology. So, perhaps collapsing across each decision problem is an unfair depiction of laboratory ecologies, and each problem should be treated as its separate ecology. In fact, commonly, each decision problem only contains nondominated options (i.e., no option is included that is inferior in all respects to the other option). As we have discussed, eliminating dominated options is sufficient to produce a negative relationship between payoffs and probabilities. Thus, payoffs and probabilities *within* each decision problem will be inversely related. However, a critical difference between the domains in our ecological analysis and the decision problems studied in the laboratory is that the laboratory problems typically contain only two options per problem. This structure offers enormous latitude to the negative relationship between payoffs and probabilities—so much latitude that, across decision problems, one and the same payoff magnitude can be associated with very small or with very large probabilities. To appreciate this, consider the hypothetical gambles from Kahneman and Tversky (1979). At the time of their data collection, so they mentioned (Kahneman & Tversky, 1979, p. 264), the median net monthly income of an Israeli family was 3,000 Israeli pounds. The payoffs in the neighborhood of this value were associated with probabilities ranging from as little as .002 to as much as .9 (see Figure 9, left-most plot in the top panel)—a vast range of probabilities for what was the equivalent of a family's monthly income. Such a nearly unconstrained alignment of payoffs and probabilities is not evident in any of the real-world ecologies we examined.
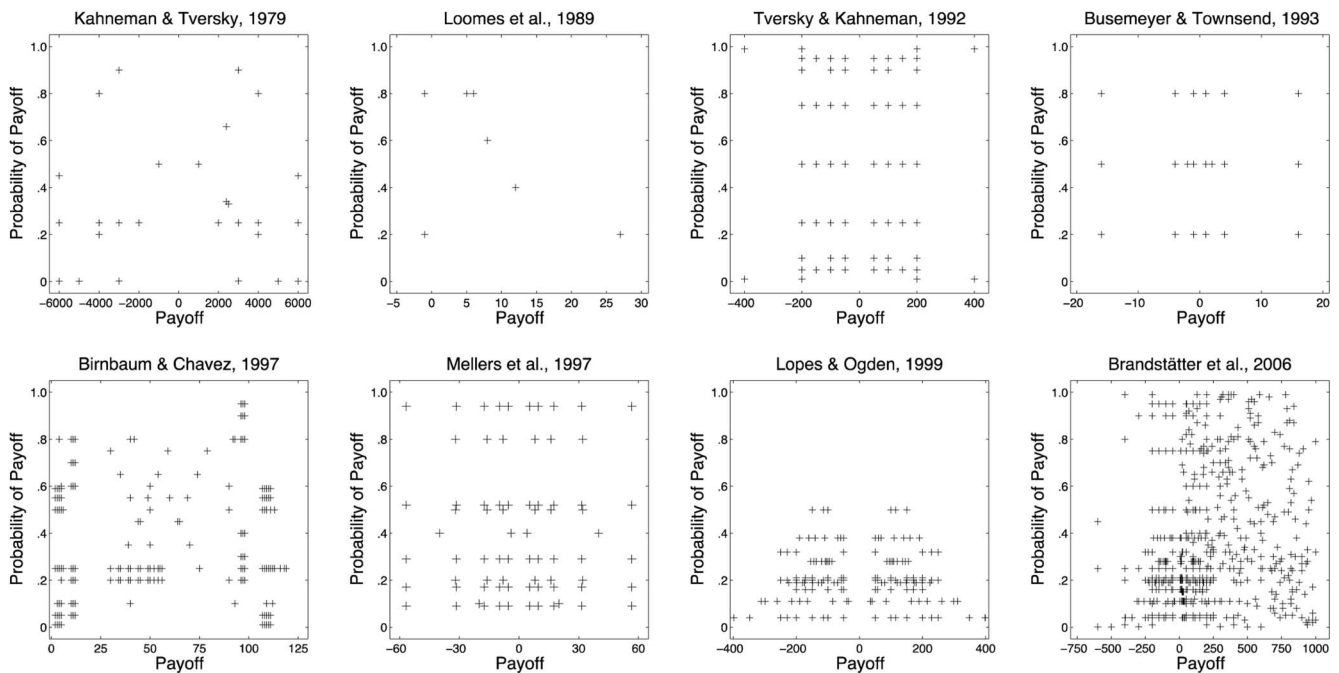

Figure 9. Relationship between payoffs and probabilities as realized in laboratory investigations of prominent theories of risky choice. Note that Brandstätter, Gigerenzer, and Hertwig (2006) used many of the gambles from the other investigations in the plot, including Kahneman and Tversky (1979), Tversky and Kahneman (1992), Mellers, Schwartz, Ho, and Ritov (1997), and Lopes and Oden (1999). Because of the variety of gambles, we normalized all payoffs to be on the same payoff scale of −1,000 to 1,000.

Unlike Brunswik, who often treated representative and systematic design as antagonistic tools, we believe that there is a division of labor between them (Dhami et al., 2004). For instance, if the goal is to evaluate competing choice models, then research is likely to rely on systematically designed stimuli that discriminate between them. If the goal is to understand how an organism functions and performs in an uncertain environment, then representative design is likely to be the method of choice. Employing representative and systematic designs as complimentary tools promises novel questions. For example, do violations of axioms of normative choice theories such as the sure-thing principle, substitution axiom, and transitivity generalize to investigations implementing representative designs, or are they restricted to the stimuli systematically designed to demonstrate them? Relatedly, does our limited ability to predict real-world risk taking improve when choice models are estimated in experimental choice environments that are representative of the gambles in the real world? Indeed, some evidence suggests that individual investors prefer stocks that share the same structural features as lottery tickets (Kumar, 2009).

Conclusion

Risk and reward: We have shown with ecological analyses that these two variables are often intimately tied in a wide range of environments that people experience and make decisions in every day. Moreover, we propose and provide some evidence that this ecological relationship enters the process by which people make decisions in the form of the risk–reward heuristic. This heuristic is one tool that empowers people to navigate uncertainty by inferring an unknown probability of a payoff from the magnitude of the payoff itself. These findings challenge the often reductionist approach of theories of decision making under risk and uncertainty, which seek to understand decision making by decomposing complex real-world risky decisions into choices between monetary lotteries that are defined by the independent constructs of payoffs and probabilities. An adaptive view of rational choice, however, suggests that this is a questionable approach. In real environments, these two constructs are often not independent, and people appear to exploit their dependency.

References

Aarssen, L. W., Tregenza, T., Budden, A. E., Lortie, C. J., Koricheva, J., & Leimu, R. (2008). Bang for your buck: Rejection rates and impact factors in ecological journals. *Open Ecology Journal, 1,* 14–19. doi: 10.2174/1874213000801010014

American Gaming Association. (2011). *State of the states: The AGA survey of casino entertainment.* Washington, DC: Author.

American Psychological Association. (2011). Summary report of journal operations, 2010. *American Psychologist, 66,* 405–406. doi:10.1037/a0023978

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review, 98,* 409–429. doi:10.1037/0033-295X.98.3.409

Anderson, J. R. (1998). Foreword. In M. Oaksford & N. Chater, *Rational models of cognition* (pp. v–vii). New York, NY: Oxford University Press.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2,* 396–408. doi:10.1111/j.1467-9280.1991.tb00174.x

Bak, P. (1999). *How nature works: The science of self-organized criticality.* New York, NY: Copernicus.

Bar-Hillel, M., & Budescu, D. (1995). The elusive wishful thinking effect. *Thinking and Reasoning, 1,* 71–103. doi:10.1080/13546789508256906

Benartzi, S., & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics, 110,* 73–92. doi:10.2307/2118511

Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrika, 22,* 23–36. doi:10.2307/1909829

Binmore, K. (2007). Rational decisions in large worlds. *Annales d'Economie et de Statistique, 86,* 25–41.

Birnbaum, M. H., & Chavez, A. (1997). Tests of theories of decision making: Violations of branch independence and distribution independence. *Organizational Behavior and Human Decision Processes, 71,* 161–194. doi:10.1006/obhd.1997.2721

Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review, 90,* 166–193. doi:10.1257/aer.90.1.166

Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review, 113,* 409–432. doi:10.1037/0033-295X.113.2.409

Brunswik, E. (1940). Organismic achievement and environmental probability. *Psychological Review, 47,* 69–78. doi:10.1037/h0057921

Brunswik, E. (1952). *The conceptual framework of psychology.* Oxford, England: University of Chicago Press.

Brunswik, E. (1955a). In defense of probabilistic functionalism: A reply. *Psychological Review, 62,* 236–242. doi:10.1037/h0040198

Brunswik, E. (1955b). Representative design and probabilistic theory in a functional psychology. *Psychological Review, 62,* 193–217. doi:10.1037/h0047470

Bureau of Labor Statistics. (2009). *Consumer expenditure survey.* Retrieved from http://www.bls.gov/cex/csxstnd.htm#2009

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100,* 432–459. doi:10.1037/0033-295X.100.3.432

Camerer, C. F. (2000). Prospect theory in the wild: Evidence from the field. In D. Kahneman & A. Tversky (Eds.), *Choice, values, and frames* (pp. 288–300). New York, NY: Cambridge University Press.

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction.* Princeton, NJ: Princeton University Press.

Camerer, C. F., & Fehr, E. (2006, January 6). When does "economic man" dominate social behavior? *Science, 311,* 47–52. doi:10.1126/science.1110600

Camerer, C. F., & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty, 5,* 325–370. doi:10.1007/BF00122575

Clotfelter, C. T., & Cook, P. J. (1989). *Selling hope: State lotteries in America.* Cambridge, MA: Harvard University Press.

Cooksey, R. W. (1996). *Judgment analysis: Theory, methods, and applications.* San Diego, CA: Academic Press.

Curley, S. P., & Yates, J. F. (1989). An empirical evaluation of descriptive models of ambiguity reactions in choice situations. *Journal of Mathematical Psychology, 33,* 397–427. doi:10.1016/0022-2496(89)90019-9

Damasio, A. R. (1994). *Descartes' error.* New York, NY: Harper Collins.

de Quervain, D. J.-F., Fischbacher, U., Treyer, V., Schellhammer, M., Schnyder, U., Buck, A., & Fehr, E. (2004, August 27). The neural basis of altruistic punishment. *Science, 305,* 1254–1258. doi:10.1126/science.1100735

Dhami, M. K., Hertwig, R., & Hoffrage, U. (2004). The role of representative design in an ecological approach to cognition. *Psychological Bulletin, 130,* 959–988. doi:10.1037/0033-2909.130.6.959

Edwards, W. (1962). Utility, subjective probability, their interaction, and variance preferences. *Journal of Conflict Resolution, 6,* 42–51. doi:10.1177/002200276200600106

Einhorn, H. J., & Hogarth, R. M. (1985). Ambiguity and uncertainty in probabilistic inference. *Psychological Review, 92,* 433–461. doi:10.1037/0033-295X.92.4.433

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics, 75,* 643–669. doi:10.2307/1884324

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review, 101,* 519–527. doi:10.1037/0033-295X.101.3.519

Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance, 47,* 427–465. doi:10.1111/j.1540-6261.1992.tb04398.x

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics, 114,* 817–868. doi:10.1162/003355399556151

Fenner, F., Henderson, D., Arita, I., Jezek, Z., & Ladnyi, I. (1988). *Smallpox and its eradication.* Geneva, Switzerland: World Health Organization.

Ferrari, S. L. P., & Cribari-Neto, F. (2004). Beta regression for modeling rates and proportions. *Journal of Applied Statistics, 10,* 1–18.

Fiedler, K. (2000). Beware of samples! A cognitive–ecological sampling approach to judgment biases. *Psychological Review, 107,* 659–676. doi:10.1037/0033-295X.107.4.659

Fiedler, K., & Juslin, P. (2006). *Information sampling and adaptive cognition.* New York, NY: Cambridge University Press.

Fischbacher, U., Fong, C. M., & Fehr, E. (2009). Fairness, errors and the power of competition. *Journal of Economic Behavior & Organization, 72,* 527–545. doi:10.1016/j.jebo.2009.05.021

Fox, C. R., & Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *Quarterly Journal of Economics, 110,* 585–603. doi:10.2307/2946693

Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science, 44,* 879–895. doi:10.1287/mnsc.44.7.879

Frisch, D., & Baron, J. (1988). Ambiguity and rationality. *Journal of Behavioral Decision Making, 1,* 149–157. doi:10.1002/bdm.3960010303

Garfield, E. (1955, July 15). Citation indexes for science: A new dimension in documentation through association of ideas. *Science, 122,* 108–111. doi:10.1126/science.122.3159.108

Garfield, E. (1972, November 3). Citation analysis as a tool in journal evaluation: Journals can be ranked by frequency and impact of citations for science policy studies. *Science, 178,* 471–479. doi:10.1126/science.178.4060.471

Garfield, E. (1996, June 10). What is the primordial reference for the phrase "publish or perish"? *Scientist, 10*(12), 11.

Genex Cooperative. (2012). *Dairy: Sire catalog.* Retrieved from http://genex.crinet.com/dairy/index.php?action=BYBREED&Breed=HO&country_code=UJV&lang=EN

Gibson, J. (1979). *The ecological approach to visual perception.* Boston, MA: Houghton Mifflin.

Gide, A. (1973). *The counterfeiters.* New York, NY: Vintage.

Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J., & Krueger, L. E. (1989). *The empire of chance: How probability changed science and everyday life.* doi:10.1017/CBO9780511720482

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

Gutenberg, B., & Richter, C. F. (1965). *Seismicity of the earth and associated phenomena.* New York, NY: Hafner.

Güth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization, 3,* 367–388. doi:10.1016/0167-2681(82)90011-7

Güth, W., & Tietz, R. (1990). Ultimatum bargaining behavior: A survey and comparison of experimental results. *Journal of Economic Psychology, 11,* 417–449. doi:10.1016/0167-4870(90)90021-Z

Hacking, I. (1975). *Emergence of probability.* Bristol, England: Cambridge University Press.

Hammond, K. R., & Stewart, T. R. (2001). *The essential Brunswik.* New York, NY: Oxford University Press.

Harris, A. J. L., Corner, A., & Hahn, U. (2009). Estimating the probability of negative events. *Cognition, 110,* 51–64. doi:10.1016/j.cognition.2008.10.006

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences, 13,* 517–523. doi:10.1016/j.tics.2009.09.004

Hertwig, R., Hoffrage, U., & the ABC Research Group. (2013). *Simple heuristics in a social world.* New York, NY: Oxford University Press.

Hertwig, R., Pachur, T., & Kurzenhäuser, S. (2005). Judgments of risk frequencies: Tests of possible cognitive mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 621–642. doi:10.1037/0278-7393.31.4.621

Hurwicz, L. (1951). *A class of criteria for decision-making under ignorance* (Cowles Commission Discussion Paper: Statistics No. 356). Retrieved from http://cowles.econ.yale.edu/P/ccdp/st/s-0356.pdf

Ijiri, Y., & Simon, H. A. (1977). *Skew distributions and the sizes of business firms.* New York, NY: North-Holland.

Irwin, F. W. (1953). Stated expectations as functions of probability and desirability of outcomes. *Journal of Personality, 21,* 329–335. doi:10.1111/j.1467-6494.1953.tb01775.x

Joffe, G. P., Foxman, B., Schmidt, A. J., Farris, K. B., Carter, R. J., Neumann, S., . . . Walters, A. M. (1992). Multiple partners and partner choice as risk factors for sexually transmitted disease among female college students. *Sexually Transmitted Diseases, 19,* 272–278. doi:10.1097/00007435-199209000-00006

Jullien, B., & Salanie, B. (2000). Estimating preferences under risk: The case of racetrack bettors. *Journal of Political Economy, 108,* 503–530. doi:10.1086/262127

Juni, M. Z., Gureckis, T. M., & Maloney, L. T. (2012). One-shot lotteries in the park. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th annual conference of the cognitive science society* (pp. 1745–1749). Austin, TX: Cognitive Science Society.

Juslin, P. (1994). The overconfidence phenomenon as a consequence of informal experimenter-guided selection of almanac items. *Organizational Behavior and Human Decision Processes, 57,* 226–246. doi:10.1006/obhd.1994.1013

Juslin, P., Olsson, H., & Bjoerkman, M. (1997). Brunswikian and Thurstonian origins of bias in probability assessment: On the interpretation of stochastic components of judgment. *Journal of Behavioral Decision Making, 10,* 189–209. doi:10.1002/(SICI)1099-0771(199709)10:3<189::AID-BDM258>3.0.CO;2-4

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–291. doi:10.2307/1914185

Kahneman, D., & Tversky, A. (Eds.). (2000). *Choices, values, and frames.* New York, NY: Cambridge University Press.

Knight, F. H. (1921). *Risk, uncertainty, and profit.* New York, NY: Sentry Press.

Kringle, A. (2008). *Understanding the Genex fertility advantage.* Retrieved from http://genex.crinet.com/page1868/UnderstandingTheGenexFertilityAdvantage

Krizan, Z., & Windschitl, P. D. (2007). The influence of outcome desirability on optimism. *Psychological Bulletin, 133,* 95–121. doi:10.1037/0033-2909.133.1.95

Kruschke, J. K. (2010). *Doing Bayesian data anlaysis: A tutorial with R and BUGS.* Burlington, MA: Academic Press.

Kühberger, A., & Perner, J. (2003). The role of competition and knowledge in the Ellsberg task. *Journal of Behavioral Decision Making, 16,* 181–191. doi:10.1002/bdm.441

Kumar, A. (2009). Who gambles in the stock market? *Journal of Finance, 64,* 1889–1933. doi:10.1111/j.1540-6261.2009.01483.x

Lee, M. D., & Wagenmakers, E. J. (2013). *Bayesian modeling for cognitive science: A practical course.* New York, NY: Cambridge University Press.

Little, K. (2013). *Understanding risk: Risk and reward are part of investing*. Retrieved from http://stocks.about.com/od/riskreward/a/Understandrisk.htm

Locke, J. (1959). *An essay concerning human understanding*. New York, NY: Dover. (Original work published 1690)

Loomes, G., Starmer, C., & Sugden, R. (1989). Preference reversal: Information-processing effect or rational non-transitive choice? *The Economic Journal, 99,* 140–151. doi:10.2307/2234076

Lopes, L. L. (1983). Some thoughts on the psychological concept of risk. *Journal of Experimental Psychology: Human Perception and Performance, 9,* 137–144. doi:10.1037/0096-1523.9.1.137

Lopes, L. L., & Oden, G. C. (1999). The role of aspiration level in risky choice: A comparison of cumulative prospect theory and SP/A theory. *Journal of Mathematical Psychology, 43,* 286–313. doi:10.1006/jmps.1999.1259

Luce, R. D., & Raiffa, H. (1957). *Games and decisions*. New York, NY: Dover.

Luce, R. D., & Weber, E. U. (1986). An axiomatic theory of conjoint, expected risk. *Journal of Mathematical Psychology, 30,* 188–205. doi:10.1016/0022-2496(86)90013-1

Mandelbrot, B. B. (1982). *The fractal geometry of nature*. New York, NY: Times Books.

Markowitz, H. (1952). The utility of wealth. *Journal of Political Economy, 60,* 151–158. doi:10.1086/257177

Mellers, B. A., Schwartz, A., Ho, K., & Ritov, I. (1997). Decision affect theory: Emotional reactions to the outcomes of risky options. *Psychological Science, 8,* 423–429. doi:10.1111/j.1467-9280.1997.tb00455.x

Osherson, D. N. (1995). Probability judgment. In E. E. Smith & D. N. Osherson (Eds.), *Thinking* (pp. 35–75). Cambridge, MA: MIT Press.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. doi:10.1017/CBO9781139173933

Proffitt, D. R. (2006). Distance perception. *Current Directions in Psychological Science, 15,* 131–135. doi:10.1111/j.0963-7214.2006.00422.x

Rabin, M. (1993). Incorporating fairness into game-theory and economics. *American Economic Review, 83,* 1281–1302.

Ramsey, F. P. (1926). Truth and probability. In R. B. Braithwaite (Ed.), *The foundations of mathematics and other logical essays* (pp. 156–198). New York, NY: Harcourt Brace.

Rode, C., Cosmides, L., Hell, W., & Tooby, J. (1999). When and why do people avoid unknown probabilities in decisions under uncertainty? Testing some predictions from optimal foraging theory. *Cognition, 72,* 269–304. doi:10.1016/S0010-0277(99)00041-4

Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review, 6,* 41–49.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association, 46,* 55–67. doi:10.1080/01621459.1951.10500768

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

Schnell, S. (2008). *The cow of the future*. Retrieved from http://genex.crinet.com/page1685/CowOfTheFuture

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance, 19,* 425–442. doi:10.2307/2977928

Sharpe, W. F. (1966). Mutual fund performance. *Journal of Business, 39,* 119–138. doi:10.1086/294846

Shepard, R. N. (1990). *Mind sights*. New York, NY: Freeman.

Shepard, R. N. (2001). Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin & Review, 1,* 2–28. doi:10.3758/BF03200759

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics, 69,* 99–118. doi:10.2307/1884852

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review, 63,* 129–138. doi:10.1037/h0042769

Skyrms, B. (1966). *Choice and chance*. Belmont, CA: Dickenson.

Smithson, M., & Verkuilen, J. (2006). A better lemon squeezer? Maximum-likelihood regression with beta-distributed dependent variables. *Psychological Methods, 11,* 54–71. doi:10.1037/1082-989X.11.1.54

Thaler, R. H. (1987). The psychology of choice and the assumptions of economics. In A. E. Roth (Ed.), *Laboratory experimentation in economics: Six points of view* (pp. 99–130). doi:10.1017/CBO9780511528316.004

Todd, P. M., Gigerenzer, G., & the ABC Research Group. (2012). *Ecological rationality: Intelligence in the world*. doi:10.1093/acprof:oso/9780195315448.001.0001

Trimmer, P. C., Houston, A. I., Marshall, J. A., Mendl, M. T., Paul, E. S., & McNamara, J. M. (2011). Decision-making under uncertainty: Biases and Bayesians. *Animal Cognition, 14,* 465–476. doi:10.1007/s10071-011-0387-4

Tversky, A. (1969). Intransitivity of preferences. *Psychological Review, 76,* 31–48. doi:10.1037/h0026750

Tversky, A., & Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review, 102,* 269–283. doi:10.1037/0033-295X.102.2.269

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5,* 297–323. doi:10.1007/BF00122574

Tversky, A., & Wakker, P. (1995). Risk attitudes and decision weights. *Econometrica, 63,* 1255–1280. doi:10.2307/2171769

U.S. Census Bureau. (2009). *State governance finances*. Retrieved from http://www.census.gov/govs/state/historical_data_2009.html

U.S. Census Bureau. (2011). *State & county quickfacts*. Retrieved from http://quickfacts.census.gov/qfd/download_data.html

U.S. Social Security Administration. (n.d.). *Period life table, 2007*. Retrieved from http://www.ssa.gov/oact/STATS/table4c6.html

Viscusi, W. K. (1989). Prospective reference theory: Toward an explanation of the paradoxes. *Journal of Risk and Uncertainty, 2,* 235–263. doi:10.1007/BF00209389

von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

Weber, E. U. (1994). From subjective probabilities to decision weights: The effect of asymmetric loss functions on the evaluation of uncertain outcomes and events. *Psychological Bulletin, 115,* 228–242. doi:10.1037/0033-2909.115.2.228

Weber, E. U., Shafir, S., & Blais, A. R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review, 111,* 430–445. doi:10.1037/0033-295X.111.2.430

Windschitl, P. D., Smith, A. R., Rose, J. P., & Krizan, Z. (2010). The desirability bias in predictions: Going optimistic without leaving realism. *Organizational Behavior and Human Decision Processes, 111,* 33–47. doi:10.1016/j.obhdp.2009.08.003

Yamagishi, T. (1986). The provision of a sanctioning system as a public good. *Journal of Personality and Social Psychology, 51,* 110–116. doi:10.1037/0022-3514.51.1.110

Yates, J. F., & Zukowsky, L. G. (1976). Characterization of ambiguity in decision making. *Behavioral Science, 21,* 19–25. doi:10.1002/bs.3830210104

Zipf, G. K. (1949). *Human behavior and the principle of least effort*. Cambridge, MA: Addison-Wesley.

(Appendices follow)

Appendix A

Description of Bayesian Analysis Methods

Bayesian Hierarchical Model of Pearson *r*

The following code describes the implementation of the Bayesian hierarchical model of Pearson *r*:

```
model {
#Likelihood
        for (i in 1:n ) {
                x[i,1:2] ~ dmnorm(mu[groupj[i],1:2],T1[,])
        }
        for (j in 1:nGroupj){
                mu[j,1] ~ dnorm(muG1,tauG1)
                mu[j,2] ~ dnorm(muG1,tauG2)
        }

        #Prior distributions
        muG1 ~ dnorm(0, 1.0E-12)
        muG2 ~ dnorm(0, 1.0E-12)
        tauG1 ~ dgamma(.001,.001)
        tauG2 ~ dgamma(.001,.001)
        lambda[1] ~ dgamma(.001,.001)
        lambda[2] ~ dgamma(.001,.001)
        r ~ dunif(-1,1)

        #reparameterization of the covariance matrix
        sigma[1] <- 1/sqrt(lambda[1])
        sigma[2] <- 1/sqrt(lambda[2])
        T[1,1] <- 1/lambda[1]
        T[1,2] <- r*sigma[1]*sigma[2]
        T[2,1] <- r*sigma[1]*sigma[2]
        T[2,2] <- 1/lambda[2]
        T1[1:2,1:2] <- inverse(T[1:2,1:2])
}
```

Data supplied from outside the JAGS software:

- x = matrix with two columns, one for each observed variable, with a length of the number of observations.
- groupj = column vector denoting group membership of each pair of observations.

Bayesian Beta Regression Model

The following code describes the implementation of the Bayesian beta regression model:

```
model {
for(i in 1: Ntotal) {
y[i] ~ dbeta(a[i], b[i])
a[i] <- mu[i]*tau
b[i] <- (1-mu[i])*tau
mu[i] <- 1/(1+exp(-(b0+b1*x[i]))) %using a logit link
}
```

(*Appendices continue*)

```
b0 ~ dnorm(0, 1.0E-12)
b1 ~ dnorm(0, 1.0E-12)
tauinv ~ dgamma(.1, 0.1)
tau <- 1/tauinv
}
```

Data supplied from outside JAGS:

- Ntotal = number of independent observations.
- y = observed estimates of the number in the envelope, scaled to be between .0001 and .9999.
- x = dollar value of the payoff.

Appendix B

Estimation Procedure for the Probabilities of Fair Bets for Multiple-Outcome Gambles

Our solution for estimating the probabilities of gambles that offer the chance to win more than a single possible payoff in exchange for a fixed loss (e.g., state lotteries) that make them fair requires an estimate of the expected gain given that one has won, $EV(gain|win)$. This process involves three steps.

Step 1

Estimate the probability of a loss assuming fair bets. In multiple-outcome gambles like scratch-off tickets, where there is a fixed cost l for a possibility of a range of payoffs, the gamble is fair when

$$(1 - p_l) \cdot EV(gain|win) = p_l \cdot l,$$

where $EV(gain|win)$ is the expected gain conditional on a win. The probability of a loss is then

$$p_l = \frac{EV(gain|win)}{EV(gain|win) + loss}.$$

In our calculations, we set the expected gain conditional on a win to $EV(gain|win) = \$2.75$. This value is roughly equal to the average gain in our sample of scratch-off lotteries.

Step 2

Estimate the conditional probabilities of each of the possible $i = 1, 2, \ldots, n$ gains given that one has won (p_i^*) using the value of $EV(gain|win)$ from Step 1. Setting

$$EV(gain|win) = p_1^* \cdot x_1 + p_2^* \cdot x_2 + \ldots + p_i^* \cdot x_i + \ldots + p_n^* \cdot x_n,$$

where $x_1 \le x_2 \le \ldots \le x_i \ldots \le x_n$ and $\sum_{i=1}^{n} p_i^* = 1$. Clearly, when there are more than two payoffs, there are not enough constraints to estimate each of the probabilities individually. To work around this problem, we estimated each of the probabilities for outcomes x_2 through x_n conditional on winning and the payoff not being x_1 using the expression for the probability of a gain for a fair bet with single possible gain (see Equation 3 in the main text), so that for $j = 2, \ldots, n$,

$$r_j = \frac{l}{x_j + l}.$$

To make them probabilities, the estimates are normalized to sum to 1. They are then used to estimate the probability of the smallest outcome x_1.

$$p_1^* = 1 - \frac{EV(gain|win) - x_1}{\sum_{j=2}^{n} r_j \cdot x_j - x_1}.$$

The remaining conditional probabilities for $j = 2 \ldots n$ are found as follows:

$$p_j^* = (1 - p_1^*) \cdot r_j.$$

Step 3

Find the unconditional probability p of a payoff for each gain by multiplying each conditional probability p_i^* with $(1 - p_l)$ (see Step 1), so that for $i = 1, \ldots, n$,

$$p_i = (1 - p_l) \cdot p_i^*.$$

This process was used to estimate the probabilities shown in Figures 8B, 8C, and 8D in the main text. On average, the absolute deviation between the estimated probabilities and true probabilities in the lottery tickets was .017.