Taming Uncertainty

**Ralph Hertwig, Timothy J. Pleskac, Thorsten Pachur,
and The Center for Adaptive Rationality**

**The MIT Press
Cambridge, Massachusetts
London, England**

3 Using Risk–Reward Structures to Reckon with Uncertainty

Timothy J. Pleskac, Ralph Hertwig, Christina Leuker,
and Larissa Conradt

3.1 The Risk–Reward Hypothesis

Would you stop to pick up a \$1 bill lying on the sidewalk? What about a \$100 bill? There is an old story about an economics professor and a keen-eyed student walking across a busy university quad. “Look,” says the student, pointing to the ground, “a \$100 bill!” As the student stoops to pick it up, the professor mutters, “Don’t bother—if it were really a \$100 bill, somebody would have picked it up already.”

This story is often used to introduce the efficient market hypothesis, which states that share prices on the stock market reflect all available relevant information (Fama, 1965). In other words, it is impossible to beat the market on a consistent basis: investors will not be able to earn above-average returns without taking above-average risks. By the same token, it is unlikely that a real \$100 bill will be lying on the street. We contend, however, that the story can also highlight the use of a heuristic. Instead of physically examining the bill to confirm its authenticity, the professor uses its dollar value to estimate the chance of it being real. Because the dollar value of the bill is pretty high, the professor estimates the likelihood of it being real as fairly low. By ignoring other information and using only the bill’s value to infer the probability that it is real, the professor is using what we refer to as the *risk–reward heuristic* (Pleskac & Hertwig, 2014).

Our hypothesis is that this heuristic is not something that only economic professors use, but rather a rule of thumb that many people rely on in real situations, such as when estimating the chances of winning a specific amount on the lottery, the likelihood that an insurable event will occur, or even perhaps, as we will see, when a farmer estimates the chances of a successful insemination of a dairy cow from the price of the semen. These

situations—when the payoffs or outcomes are known, but the probabilities are not—are typically referred to as *decisions under uncertainty* (Luce & Raiffa, 1957; see also chapter 18). The tricky problem with these decisions is that probabilities have to be estimated. This is where the risk–reward heuristic comes into play. In contexts where payoffs and probabilities are inversely related, the risk–reward heuristic can be used to infer unknown probabilities from observed payoffs. But can this inferential heuristic be accurate? And do people use it accurately?

We address these questions by first summarizing the results of an ecological analysis that identifies conditions under which the risk–reward heuristic can provide accurate estimates of probabilities. This ecological analysis also provides insights into why the relationship occurs. We then review what is known about when people use the heuristic and provide empirical evidence that they use it in adaptive ways. Finally, we illustrate how this ecologically grounded inference can shed light on some basic theoretical problems in the area of judgment and decision making and, more practically, how it can help to evaluate the bets that people might try to sell you. Taken together, we will show that the risk–reward heuristic is a simple tool people use for addressing uncertainty—one that can make good ecological sense to use.

3.2 An Ecological Analysis of the Relationship between Risk and Reward

Time and again, thinkers, writers, and scholars have described the intimate ties between risks and rewards. In considering satisfaction in life, the writer Hunter S. Thompson (1955) stated, “So we shall let the reader answer this question for himself: who is the happier man, he who has braved the storm of life and lived or he who has stayed securely on shore and merely existed?” (p. 5). Similarly, Niccolò Machiavelli, the famous Italian statesmen and political philosopher, once noted that in politics, “Never was anything great achieved without danger.” In the area of finance, a casual search of the Internet or self-help books reveals statements like “understanding the relationship between risk and reward is a key piece in building your personal investment philosophy” (Little, 2018). All of these nuggets of wisdom suggest that the key to success is knowing how to trade risk for reward. Yet trade-offs between risk and reward can be made if only there is a systematic relationship between the two.

3.2.1 A Marketplace of Human-Made Gambles

Against this background, Pleskac and Hertwig (2014) conducted an ecological analysis of the risks and rewards inherent in a diverse set of life’s gambles. Their goal was twofold. First, they sought to determine whether there are indeed reliable risk–reward structures in various environments. Second, they aimed to shed light on possible causes of this relationship. The first set of environments that they examined were actual monetary gambles. Consider, for instance, the game of roulette. A ball is spun around a wheel with 38 unique pockets, and players bet on where the ball will land. Players have a variety of betting options. For instance, they can bet that the ball will land in pocket number 1 or in a red pocket. In total, there are 155 possible bets; of those, 22 are unique in terms of payout. Figure 3.1 plots the relationship between risks and rewards across those 22 bets. It shows a systematic inverse relationship that we have all come to expect, with increasing payoffs becoming less and less likely to occur.

Figure 3.2a also plots the risk–reward relationship in roulette, but now both payoffs and probabilities are on a log scale. Plotting the risk–reward relationship in this way helps show the same inverse relationship extends to other monetary environments, including the bets offered at a horse track (see figure 3.2b) and the conditions of a life insurance policy (see figure 3.2c).

Figure 3.1
Relationship between payoffs and probabilities in the game of roulette (recreated from Pleskac & Hertwig, 2014).

Figure 3.2

Relationship between payoffs and probabilities (plotted in log-log space) in the context of roulette (a), the horse track (b), and life insurance policies (c). The dashed lines in each panel show the fair-bet probabilities p ; see text for details (recreated from Pleskac & Hertwig, 2014).

3.2.2 Fair Bets

At this point, you may feel inclined to protest, “Of course there is a risk–reward relationship in these gambles—that is exactly how they were constructed!” This objection is justified, but it also pinpoints the first driver of the risk–reward relationship: the forces of the marketplace. To see how these forces give rise to the risk–reward relationship, let us consider a seller in possession of a gamble that has the prospect of gaining g if an event

whose probability is p occurs (e.g., the United States beats Germany at soccer). The seller offers a player this gamble at price c . 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 forces push the gambles toward a state in which expected gains correspond to expected losses, such that:

$$p \times g = (1 - p) \times c. \quad (1)$$

Reworking this equation shows that under these conditions the probability of winning will be inversely proportional to the total payoff:

$$p = c / (c + g). \quad (2)$$

These fair-bet probabilities—or what we call risk-reward probabilities—are plotted as a dashed line in each panel of figure 3.2 (setting $c=1$). Comparison of actual (roulette) and estimated probabilities (horse track and life insurance) with the risk-reward probabilities shows a close alignment, supporting the intuition that the risk-reward relationship is engendered by market forces.

There are, however, small but systematic deviations (remember that the plots are in log-log space and on different scales). In roulette, the actual probability of winning deviates from p by an average absolute deviation of 0.93 percentage points. In horse racing, the average absolute deviation is 0.74 percentage points; in purchasing life insurance, it is 1.69 percentage points. Casino rule number 1 offers an answer to the question of why there are deviations: the house always wins. That is, there is a bias built into games of chance, as a result of which the risk-reward probabilities in roulette and at the horse track tend to overestimate the true probabilities. The opposite holds for life insurance, where the risk-reward probabilities underestimate the true probabilities. This happens because in insurance the financial risks are spread out across a large number of contributors, thereby offsetting the risk faced by any one contributor.

3.2.3 Beyond Monetary Gambles

In the casino, at the horse track, and when purchasing life insurance, it is hard to escape the fact that one is entering a wager. But to what extent does the relationship between risk and reward also exist outside the marketplace of human-made monetary gambles? Such real-world wagers are more difficult to identify, but consider the three examples plotted in figure 3.3.

Figure 3.3

Relationship between payoffs and probabilities in farmers' choices of semen for the artificial insemination of dairy cows, scientists' choices of psychology journals when submitting manuscripts, and in the ultimatum game (recreated from Pleskac & Hertwig, 2014).

First, figure 3.3a plots the relationship between risks and rewards for a choice made by dairy farmers on a fairly regular basis: which specimen of bull semen to use to artificially inseminate their cows. To make this choice, farmers typically consult a catalog listing the specimens available from different sires. Various statistics are listed for each specimen, but farmers are generally advised to focus on two factors: (a) the sire conception rate, an estimate of the conception rate of the sire's sperm, and (b) the lifetime net merit of the sample, an

estimate of the average improvement in profit that a dairy farmer can expect from the daughters of a specific bull. These factors are essentially proxies for probabilities and payoffs and here, too, there is an inverse relationship between risks and rewards.

Second, scientists weigh up the odds when submitting a manuscript for publication. Should they aim for a high-impact journal that rejects many of the manuscripts submitted, or should they settle for a lower-impact journal that accepts more manuscripts? Figure 3.3b plots this risk–reward relationship in the field of psychology, using a journal's acceptance rate (a proxy for risk, indicating the proportion of manuscripts accepted for publication by a given journal) and impact factor (a proxy for reward, indicating the average number of times an article in that journal is cited).

Even in these environments, market forces can explain the relationship between risk and reward. Dairy farmers would be unwise to buy the sperm of a bull with a low sire conception rate and low lifetime net merit. Scientists would be ill-advised to publish in journals with low acceptance rates and low impact factors. In other words, choosing dominated options is self-defeating. As a result, market forces will remove such dominated options. Companies only sell semen samples that are high on lifetime net merit, sire conception rate, or both, and remove samples that are low on either factor from the reference class. Such filtering is also likely to occur with scientific journals. However, filtering takes time, and not all dominated options are as easily detected as an unfair bet at the roulette table. Moreover, there are sometimes other reasons for picking an option than the potential payoff. All of these factors imply the relationship will be noisy.

But market forces are not the only driver of risk–reward relationships, as indicated by the regularity of the data points in figure 3.3c. These data show the inverse relationship that emerges in the interactions of a pair of players in the ultimatum game, an economic game used to study bargaining. Interestingly, market forces predict—if anything—the absence of a risk–reward relationship in this game. The reason is as follows. In the ultimatum game, one person, the proposer, suggests how to split a fixed monetary pie (in this case, 100 points, with each point worth real money). The split represents a take-it-or-leave-it offer that the other person, the responder, must either accept or reject. If the responder rejects the offer, both players go away empty-handed. From the perspective of a self-interested utility maximizer, this game is devoid of uncertainty: any responder maximizing their

self-interest should accept any positive payoff, no matter how small. Thus, the proposer should offer the smallest amount possible to maximize their self-interest, a lick of the spoon so to speak. However, responders rarely accept just a lick of the spoon. Instead, responders are typically willing to forgo any payoff at all so as to punish such violations of fairness (Fehr & Schmidt, 1999; Quervain et al., 2004; Rabin, 1993). But because proposers anticipate this threat of punishment, this exchange turns into an uncertain gamble where the amount of money offered trades off with the risk of being punished, thereby creating a risk–reward relationship.

Our ecological analysis has been inductive in nature, but a more deductive approach is possible. Box 3.1 outlines such a deductive approach, showing how competition for finite resources suffices to create an inverse relationship between risks and rewards. Taken together, these analyses suggest that the risk–reward relationship is not a peripheral phenomenon. Instead, it is a frequent and recurrent ecological structure driven often by market forces and competition over limited resources, but not always. The next question is whether people have learned to exploit this risk–reward structure to reckon with uncertainty.

Box 3.1

How competition for limited resources can create a risk–reward relationship.

The environments in the ecological analysis summarized in figures 3.2 and 3.3 share several factors. One is competition for limited resources, suggesting that competition may be an important ingredient for the development of risk–reward relationships. In fact, we can use a theoretical framework from behavioral ecology, known as ideal free distribution (IFD) theory (Fretwell & Lucas, 1969), to show competition is a sufficient condition for a risk–reward relationship (see Pleskac, Conradt, Leuker, & Hertwig, 2018 for more details).

IFD theory predicts that groups of foraging animals will distribute themselves proportionally to the resources in a patch—in the same way as people distribute themselves among grocery store checkout lines according to the speed at which each line is moving. Formally, IFD theory states that the number of competitors n_y in patch y is proportional to the amount of resources r_y in the patch:

$$n_y \propto r_y. \quad (\text{B1})$$

The IFD assumes that competitors are able to detect the patch with the highest rate of consumption (ideal) and are free to move between patches. These assumptions imply that competitors will move to patches with higher

Box 3.1 (continued)

potential consumption rates and that, over time, the entire system will reach an equilibrium state at which the number of competitors in each patch is proportional to the amount of resources in that patch.

For our purposes, it is useful to consider what a new competitor experiences when entering a landscape with an IFD, namely, a risk-reward relationship. To see this, we can break down the total amount of resources in patch r_y into the number of resource items m_y multiplied by the size of each resource s_y : $r_y = m_y \times s_y$. Thus, the IFD (equation B1) can be formally rewritten as:

$$n_y \propto m_y \times s_y. \quad (\text{B2})$$

We assume that the probability p of an individual competitor successfully obtaining a resource increases with the number of resource items in the patch and decreases with the number of competitors in the patch, such that

$$p_y \propto m_y / n_y. \quad (\text{B3})$$

Substituting equation B3 into equation B2 reveals that the probability of successfully obtaining a resource is inversely proportional to the size of the resource in a patch:

$$p_y \propto 1/s_y. \quad (\text{B4})$$

That is, if there is variation in the size of each resource between patches but not within patches (e.g., because different patches consist of different habitat types), then individual competitors will face a choice between patches that trade off resource size for the probability of obtaining an item. Thus, the probability of obtaining a resource item is inversely proportional to resource item size (i.e., reward size) across landscapes, a risk-reward relationship. It is in these environments that it is ecologically rational to use a heuristic that infers the probabilities of obtaining a resource or reward from the size of the resource itself: the risk-reward heuristic.

One benefit of this framework is that it predicts not only when and why a risk-reward relationship will emerge but also when it will not emerge. By extension, the theory suggests conditions under which inferring probabilities from the magnitudes of rewards will fail. One straightforward answer is that a risk-reward relationship will not be present when the system is not in equilibrium. Thus, all other things being equal, a system that either has not reached or cannot reach equilibrium in terms of the distribution of competitors will not reliably evolve a risk-reward relationship. It follows that reliance on the risk-reward heuristic is not ecologically rational in nascent competitive environments.

3.3 The Risk–Reward Heuristic

Our ecological analysis demonstrates that in many environments the probabilities of an event are inversely related to the magnitude of the associated payoffs. This finding supports our general hypothesis that people can harness the risk–reward relationship as a heuristic to gauge the unknown probability of an event occurring from the payoff associated with that event. Our analysis also provides a precise specification of the risk–reward heuristic for monetary gambles with a pay-to-play structure as follows:

For gambles that offer a single payoff g , otherwise nothing, infer the unknown probability p of winning a payoff from the ratio of the entry costs of playing the gamble c to the total amount of possible winnings as $p' = c / (c + g)$.

The probability of not winning and forfeiting the price c of the gamble is $1 - p'$.

Many decisions, including the ones in our ecological analyses, have this pay-to-play structure. There are, of course, other types of gambles, including gambles that deal with losses only or with multiple payoffs. In many of these cases, the specification of the heuristic process can be modified to provide precise predictions. For the purposes of this chapter, which seeks to test the general hypothesis that people use the risk–reward heuristic, we focus on gambles with a pay-to-play structure and the chance of winning a single possible payoff. We now review the empirical evidence we have gathered so far on whether and how people use the risk–reward heuristic.

3.4 Do People Use the Risk–Reward Heuristic?

As a first step in examining whether people use the risk–reward heuristic, we offered participants an opportunity to play a gamble for a fixed prize. We were interested in their estimates of the chances of winning and in how these estimates changed with the size of the prize. We have run the study several times (Pleskac & Hertwig, 2014). Here is what we told a recent sample of 455 volunteers in an online study (Pleskac et al., 2018):

Imagine you have been asked to play the following gamble. The gamble offers the opportunity to win €4, but it costs you €2 to play. If you choose to play you would pay the €2 and, without looking, draw a marble from a basket. In the basket there are 1000 marbles. The marbles are either black or red. If the marble is red you will win €4; otherwise, if the marble is black you will receive nothing. Thus, the

number of red marbles in the basket determines the probability that you will win. You are less likely to win the €4 the lower the number of red marbles in the basket.

Some participants saw the vignette above. Others saw the same vignette, but with a different prize amount. Some had the chance to win €2.5, others €10, others €50, all the way up to €10,000. We asked the volunteers two questions (counterbalanced):

1. How many red, winning marbles do you expect are in the urn?
2. Would you play the bet, yes or no?

Although this is a simple study, there are several competing hypotheses on how people ignorant of the distribution of marbles will respond to the questions. They could employ the risk–reward heuristic, according to which the estimated number of winning marbles will decrease as the value of the payoff increases. Alternatively, they may be influenced by the desirability bias (see Krizan & Windschitl, 2007, for a review) and assess payoffs as more likely to occur as they become more desirable. A final possibility is that they may subscribe to Laplace’s *principle of indifference* and treat each outcome as equally likely (Laplace, 1776, quoted in Hacking, 1975/2006, p. 132; Keynes, 1936/1973a; see also the equiprobable heuristic in chapter 2).

As figure 3.4a shows, consistent with the predictions of the risk–reward heuristic, the estimated number of winning marbles in the basket decreased as the size of the reward increased. However, the observed estimates deviated somewhat from the fair-bet probabilities. Estimates were regressive relative to the risk–reward probabilities (dotted line), overshooting them for the high prizes and undershooting them 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). This perturbation would be sufficient to produce the regressive effect observed in the estimated number of winning marbles; moreover, it is consistent with the idea that people making decisions about risky prospects from description tend to overweight rare events. In other words, the inferential nature of the risk–reward heuristic can lead to what appears to be nonlinear distortion of probabilities in decision making under uncertainty (see chapter 8; Tversky & Fox, 1995). For another possible explanation related to the makeup of the environment, including the abilities of the other competitors in the environment, see Pleskac et al. (2018).

Figure 3.4

(a) Estimated number of red marbles by reward magnitude. Higher rewards were associated with lower estimates. The dashed line depicts the objective number of marbles predicted by the risk–reward heuristic. Each bin contains ~50 responses. (b) Proportion of volunteers who would play the gamble. Those offered higher rewards were more likely to play.

Figure 3.4b shows the proportions of people who said they would pay to play each gamble. Before discussing the choice patterns further, we should emphasize that the risk–reward heuristic is not a choice heuristic but an inferential heuristic. Thus, a choice rule is needed to make predictions about people’s choices. It might be that people use the risk–reward heuristic to infer a probability and then use a process akin to prospect theory (Kahneman & Tversky, 1979) to make a choice; alternatively, they may use the estimated probabilities in a heuristic choice process such as a lexicographic rule (see chapter 2).

For the purposes of this chapter, we assume that participants in our study used the estimate from the risk–reward heuristic to calculate the expected value of the gamble and then chose the option with the larger expected value. Strictly speaking then, an expected-value calculator combined with the risk–reward heuristic predicts indifference between paying to play the gamble or not across all prize conditions. As figure 3.4b shows, this was not observed on the aggregate. Instead, as reward magnitudes increased, there was an increasing propensity to choose to play the gamble. This tendency is more consistent with participants exhibiting a desirability bias when they make choices or even relying on the principle of indifference.

However, as figure 3.4a shows, the estimates were regressive relative to the fair-bet probabilities, with the estimates for low payoffs being lower than expected and the estimates for high payoffs being slightly higher. This regressive property actually implies a preference for not playing the gamble at low payoffs, with an increasing propensity to play the gamble as payoffs increase; this is shown in figure 3.4b. Moreover, at the individual level within each condition, people who estimated a higher number of red marbles were more likely to accept the gamble (see also Pleskac & Hertwig, 2014).

In sum, people's estimates and choices reveal the use of a risk–reward heuristic, although the estimates tend to be less extreme relative to fair-bet probabilities. These results imply that the risk–reward relationship found in the world is reflected in people's minds—but how does the relationship enter the mind in the first place?

3.5 Do People Adapt to Different Risk–Reward Environments?

The risk–reward heuristic makes use of a statistical regularity in the environment. To exploit this regularity, the mind must be cognizant that a relationship exists between risk and reward—not least because the nature of that relationship differs across contexts. Whereas we found a near-perfect (though biased) inverse relationship between payoffs and probabilities in monetary gambles in casinos (see figure 3.2), the risk–reward relationship was less reliable in other environments, such as where to submit a scientific manuscript (see figure 3.3), and it may not yet have evolved in still other environments, such as in newly forming markets (see box 3.1). These differences across contexts present a challenge for users of the risk–reward heuristic—especially as people are rarely explicitly informed about the relationship between risk and reward. Nor does it seem that they deliberately set out to learn it. So, do people learn the risk–reward structure, even when it is not their main goal to do so?

To answer this question, Leuker, Pachur, Hertwig, and Pleskac (2018) ran an experiment with three risk–reward environments with monetary gambles: a negatively correlated, a positively correlated, and an uncorrelated environment (see figure 3.5). Each dot in figure 3.5 is a simple gamble of the form “ p chance of winning x , otherwise nothing.” Participants were assigned to one of the choice environments—without being told which environment they were in or being instructed to learn a risk–reward

Figure 3.5

The three experimentally manipulated risk–reward relationships. Payoffs were in terms of a laboratory currency E\$ to help minimize participants using risk–reward knowledge from other actual currencies (recreated from Leuker et al., 2018).

relationship. Instead, participants were exposed to the environment by evaluating a series of 172 gambles from it. For each gamble, they were asked to imagine they owned the gamble and to state the price at which they would be willing to sell it. To get a feel for how this experiment worked and try your hand at detecting different types of risk–reward environment, please visit interactive element 3.1 (at <https://taming-uncertainty.mpib-berlin.mpg.de/>).

After a short break, participants completed several more tasks, including the following two. In a probability estimation task, participants were

Figure 3.6
(a) Average probability estimates for each of the payoff levels in the probability estimation task. (b) Proportion of times the uncertain option was chosen in the decisions under uncertainty task.

presented with different payoff magnitudes and asked to estimate the chances that they were associated with specific rewards, based on the gambles they had seen earlier. As figure 3.6a shows, the average estimates reflected the risk–reward environment. That is, people learned the risk–reward structure without being given any feedback or having their attention drawn to the relationship.

In order to determine whether, having learned the risk–reward structure, people harness it as a heuristic when making their decisions, participants were asked to make a series of decisions between a gamble that offered a reward with an unknown probability (e.g., E\$100 with $p=?$) and a sure thing that was half the size (e.g., E\$50 with $p=1.0$). As figure 3.6b shows, the risk–reward environments to which participants were exposed impacted their choices in a manner consistent with their using the risk–reward heuristic to estimate the odds of the payoffs. Relative to people in the uncorrelated environment, people in the positive risk–reward environment showed an increasing preference for the gamble as the size of the potential reward increased. In the negative risk–reward environment, this pattern was reversed. However, there was less of a difference between the negative and uncorrelated conditions. In particular, participants in the negative condition

did not choose the low-payoff gamble as often as was warranted by their probability estimates.

To summarize, people can learn risk–reward structures even when their attention is not drawn to the risk–reward relationship. Moreover, they use those structures to help make decisions under uncertainty. As a result, the risk–reward structures to which people were exposed created environment-dependent preferences. We should emphasize that this change in preferences is not some sort of fallacy, but an ecologically rational bet on the structure of the environment. Next, we review three implications of the risk–reward heuristic that illustrate its ability to help make sense of both old and new problems in the area of decision making.

3.6 Implications of the Risk–Reward Heuristic

3.6.1 Ambiguity Aversion

When people are given a choice between two otherwise equivalent options—one in which the probability information is given and one in which it is missing—they tend to avoid the option with missing probability information. This phenomenon—which has become known as *ambiguity aversion*—was first established by Ellsberg (1961), who used it to distinguish between risk and uncertainty (see also Keynes, 1936/1973a, p. 75). Let us consider, for example, the choice between two gambles:

Gamble A: Win €100 if a red marble is drawn from an urn containing 50 red marbles and 50 black marbles.

Gamble B: Win €100 if a red marble is drawn from an urn containing 100 marbles with an unknown proportion of red to black marbles.

People generally prefer the clear gamble A to the ambiguous gamble B (Camerer & Weber, 1992). However, typically it has been assumed that if you ask people whether a red or a black marble is more likely to be drawn, in either gamble, they would say red and black are equally likely. But observing a preference for gamble A, the clear gamble, seems inconsistent with this assumption. Ellsberg (1961) used this inconsistency to argue that uncertainty cannot simply be reduced to a probability, as subjective expected utility theory would suggest (Savage, 1954).

Yet, as we have seen in this chapter, if you assign one color as the winning marble in gamble B, and ask people how likely it is, they probably

would not answer that both colors are equally likely. Instead they may use the risk–reward heuristic and intuit the probabilities from the payoff magnitudes (see figure 3.4a). Thus, it is possible that this so-called ambiguity aversion is actually the result of a reasonable inference based on the structure of the environment. Interestingly enough, this idea was anticipated by Ward Edwards (personal communication cited in Lopes, 1983, p. 141), one of the founders of behavioral decision theory. In considering why people do not adhere to the principles of subjective expected utility theory, but exhibit ambiguity aversion, Edwards observed that “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 the world in which we live.”

To test whether the risk–reward heuristic might help explain ambiguity aversion, Pleskac and Hertwig (2014) created gambles equivalent to gambles A and B and asked two groups of people for their preferences. The first group was told to assume that the payoff for drawing a red marble from either urn was 10 Swiss francs (CHF10). The second group was told to assume that the payoff for drawing a red marble was CHF1,000. Pleskac and Hertwig predicted that, due to the prevalence of an inverse risk–reward relationship in the world, the ambiguous or uncertain option would become increasingly undesirable as the magnitude of the payoff increased. The results support this prediction. In the CHF10 condition, 63% of participants preferred the unambiguous urn (gamble A). In the CHF1,000 condition, by contrast, 81% preferred the unambiguous urn.

The observed increase in ambiguity aversion as the payoff magnitude increased is difficult to explain in terms of the other accounts of ambiguity aversion. The relative variance between the two gambles across the payoff conditions remains constant, meaning that there should be no difference in the comparative ignorance people experience between the two gambles (Fox & Tversky, 1995) or the inclination to avoid variance (Rode, Cosmides, Hell, & Tooby, 1999). Moreover, our ecological analysis of life’s gambles suggests the risk–reward account of ambiguity aversion is not an unjustified suspicion (Binmore, 2007; Frisch & Baron, 1988; Kühberger & Perner, 2003; Yates & Zukowsky, 1976). Instead, one can argue that this so-called aversion merely reflects an ecologically grounded belief that high payoffs have a low probability of occurring.

3.6.2 Bets Leak Information

A second implication of the risk–reward heuristic is that the gambles people are offered convey or leak information (see also McKenzie & Nelson, 2003; Sher & McKenzie, 2006)—in particular, information about the beliefs of the person who constructed the gamble. For instance, let us consider a situation where two people—call them Tim and Ralph—are talking about an upcoming soccer match between Germany and the United States. Tim offers Ralph the following bet: for the cost of €1, he will pay Ralph €4 if Germany wins. We might think this is just an ordinary bet between colleagues. But according to the risk–reward heuristic, Tim has also shared some information with Ralph about his own beliefs about the likelihoods of the different outcomes. A challenge for future work is to determine whether and how Ralph will choose to use this information based on what he knows about Tim.

More generally, the idea that gambles can leak information about the seller's beliefs would seem to challenge the conventional approach in the decision sciences, where gambles are treated as stimuli that allow buyers' preferences and beliefs to be gauged (see, e.g., Newell, Lagnado, & Shanks, 2015; Winterfeldt & Edwards, 1986). Indeed, this point resonates with an early concern raised by Frank Ramsey (1926), the mathematician who anticipated subjective expected utility theory, with regard to using bets to identify people's beliefs: "The old-established way of measuring a person's belief is to propose a bet, and see what are the lowest odds which he will accept. This method I regard as fundamentally sound; but it suffers from being insufficiently general, and from being necessarily inexact. [...] Besides, the proposal of a bet may inevitably alter his state of opinion" (p. 68).

In fact, the problem that the bets people are offered may change their opinion is part of the reason that other methods, such as scoring rules, were developed to measure a person's beliefs (De Finetti, 1981). The risk–reward heuristic, with its ecological approach to decision making, can help to make sense of how people use the gambles they are offered to navigate uncertainty.

3.6.3 Making Money with the Risk–Reward Heuristic

A final implication of the risk–reward heuristic is that people can use it to sift through possible bets to find ones to avoid (much like the economics professor deciding not to pick up the \$100 bill) and ones to exploit. For instance, Kaunitz, Zhong, and Kreiner (2017) collected the closing odds that online bookmakers placed on soccer games from January 2005 to June

2015. They then calculated the average odds assigned to each game. Following a process exactly akin to the risk–reward heuristic, they then calculated the probabilities that those odds implied. As in the ecological analysis described in section 3.2, these intuited probabilities were good estimates of the probability of the games' outcomes.

Kaunitz et al. (2017) took the analysis one step further, noting that although most of the bookmakers' odds hovered close to the average, some odds appeared to deviate markedly from the average, making them an attractive bet.¹ Having made this observation, the authors developed an algorithm to identify these bets. Specifically, they compared the probabilities estimated from the individual odds via the risk–reward heuristic with the probabilities intuited from the average set of odds. They found that when the estimated probability of a bet was 5 percentage points lower than would be expected from the average odds, a gambler could turn a nice profit by placing a bet with that bookmaker. In fact, when they implemented their strategy with real money, placing 265 \$50 bets, they made a profit of \$957.50 (a net return of 8.5%). So apart from helping people make accurate inferences about events, the risk–reward heuristic could help people find options that may look too good to be true—but actually are true (see also Leuker, Pachur, Hertwig, & Pleskac, *in press*).

3.7 Conclusion

Scholars of probability theory have long drawn a connection between common sense and the calculation of probabilities (Daston, 1988/1995). Pierre-Simon Laplace (1814/1902), the father of Bayesian probability theory and the first advocate of the principle of indifference, emphasized the connection when he stated, “the theory of probabilities is at bottom only common sense reduced to calculus; it makes us appreciate with exactitude that which exact minds feel by a sort of instinct without being able oftentimes to give a reason for it” (pp. 196–197).

1. Bookmakers may offer odds above fair value either to attract clients from other bookmakers or to maintain a balanced book. For instance, when gamblers bet too much on one outcome, bookmakers might increase the odds of the complementary event to encourage more people to bet on the complementary event and offset the overbooked outcome.

Sometimes, though, calculus fails to provide a means of assigning probabilities to events. In such cases, mathematicians, statisticians, scientists, and laypeople alike fall back on other commonsense rules to finish their calculations, make inferences, and make decisions. The Laplacian principle of indifference, treating each outcome as equally likely, has been one of the rules used to assign probabilities to events in the absence of information (see also chapter 2). Yet, as we have shown, the payoffs associated with the events can carry information about the likelihood of those events, allowing people to estimate the probabilities from the payoffs alone. We have also shown that the risk–reward heuristic is not just an abstract rule divined by armchair psychologists, but a rule that people use and—given enough information and exposure to the environment—adjust to fit the situation. The risk–reward heuristic is a simple tool that people use to navigate uncertainty—a tool that not only makes good common sense, but good ecological sense as well.