Taming Uncertainty

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

The MIT Press
Cambridge, Massachusetts
London, England

2 The Robust Beauty of Heuristics in Choice under Uncertainty

Ralph Hertwig, Jan K. Woike, Thorsten Pachur, and Eduard Brandstätter

2.1 Axiomatizing Rational Choice—within Two Hours

> April 14, 1942. Today at Johnny's: axiomatization of measurable utility together with the numbers. It developed slowly, more and more quickly, and at the end, after two hours (!) it was nearly completely finished. It gave me great satisfaction, and moved me so much that afterwards I could not think about anything else.... (Oskar Morgenstern, cited in Leonard, 1995, p. 753)

The diary writer is the Austrian economist Oskar Morgenstern; Johnny is the Hungarian-born mathematician John von Neumann. The two men first met in the fall of 1938, by which time they had both left Europe for good and were working at Princeton. The culmination of their collaboration, the book *Theory of Games and Economic Behavior* (von Neumann & Morgenstern, 1944/2007), was an intellectual coup that would thoroughly transform a range of fields. One major step on their route to game theory was to formulate—within "two hours (!)"—an axiomatic foundation of Daniel Bernoulli's (1738/1954) path-breaking expected utility theory. According to this theory, a rational decision maker will choose among risky options in such a way as to maximize expected utility. Von Neumann and Morgenstern derived a set of axioms—such as transitivity, completeness, and independence (see Luce & Raiffa, 1957)—that the preferences and choices of a decision maker obeying expected utility theory would have to satisfy. The axiomatized version of utility theory swiftly became a framework for research in areas as diverse as statistical decision theory, management science, operation research, and the theory of the firm. Within a decade, expected utility theory was generalized from "objective" probabilities (or "risk," to use Knight's, 1921/2002, terminology) to "subjective" probabilities (Savage, 1954), giving rise to what is now called Bayesian decision theory. According to the Bayesian approach, a rational person can translate any uncertainty into numbers, that is, subjective probabilities, which must, first and foremost, be consistent—but not necessarily plausible. Diehard Elvis Presley fans who believe he is living among us, now in his early eighties, may estimate this probability to be 99% and assign 1% to him being dead. Largely unconstrained by facts, such beliefs can nevertheless be coherent.

2.2 The Olympian Model and Its Unrealistic Assumptions

Von Neumann and Morgenstern's (1944/2007) axiomatized utility theory also evoked fierce criticism. One challenge was empirical in nature. French economist and later Nobel laureate Maurice Allais (1953) did not mince his words: "Whatever their attraction might be, none of the fundamental postulates leading to the Bernoulli principle as formulated by the American school can withstand analysis. All are based on false evidence" (p. 505). Another challenge was conceptual. Although Herbert Simon respected utility theory's normative appeal (at least for the domains in which its assumptions hold)[1]—he profoundly criticized its unrealistic assumptions. In his article *A Behavioral Model of Rational Choice* (1955), Simon spelled out the "severe demands" of what he later described as an "Olympian model" (Simon, 1983, p. 19)—an ideal that might work for omniscient gods, but was simply out of place in the real world:

> If we examine closely the "classical" concepts of rationality [...], we see immediately what severe demands they make upon the choosing organism. The organism must be able to attach definite pay-offs (or at least a definite range of pay-offs) to each possible outcome. This, of course, involves also the ability to specify the exact nature of outcomes—there is no room in the scheme for "unanticipated consequences." The pay-offs must be completely ordered—it must always be possible to specify, in a consistent way, that one outcome is better than, as good as, or worse than any other. And, if the certainty or probabilistic rules are employed, either the outcomes of particular alternatives must be known with certainty, or at least it must be possible to attach definite probabilities to outcomes. (pp. 103–104)

In actual human choice, such demands are rarely met. Beyond what Savage (1954) called "small worlds"—highly simplified environments such as

1. In fact, Simon (1945) wrote a glowing review of *Theory of Games and Economic Behavior.*

monetary gambles, where the consequences (e.g., monetary payoffs) and probabilities of all outcomes are known ("decisions from description"; Hertwig, 2015)—it is impossible for real people to live up to the decision-making ideal of specifying all possible outcomes, assigning them probabilities, and then maximizing the expected payoff. Instead, mere mortals often have access to only some of the information, or are unable to integrate that information in the sophisticated way mandated by expected utility theory, and instead rely on simplifying procedures. What do these constraints mean for the quality of people's choices? Simon (1956) conjectured that real organisms' behavior, although adaptive and satisficing, probably *"falls far short* of the ideal of 'maximizing' as postulated in economic theory" (Simon, 1956, p. 129, emphasis in the original). He further suggested that "the environments to which organisms must adapt possess properties that permit further simplification of its choice mechanisms" (p. 129).

Our goal in this chapter is to examine both of Simon's key theses. First, we investigate the price of simplicity: How far short of the ideal of maximization do simple choice strategies that fail to obey the demands of classical rationality fall? Second, we examine which statistical properties of the environment support or impede the performance of such heuristics. We analyze both questions in the time-honored environment of monetary gambles, that is, the very environment from which the concept of mathematical expectation and the classic notion of rational choice emerged (Hacking, 1975/2006; see also chapter 8). We do not, however, implement the monetary gambles as "small worlds" (Savage, 1954), fully described and with all outcomes and probabilities known; instead, we introduce uncertainty into the simulated environment in the form of imperfect knowledge. Before we describe this paradigm in more detail, let us briefly review two investigations that have inspired our own.

2.3 How Short Do Risky Choice Heuristics Fall of the Ideal of Maximization?

This analysis builds on the foundation of two previous investigations: Thorngate's (1980) strategy tournament and J. W. Payne, Bettman, and Johnson's (1988, 1993) influential research on the adaptive decision maker. Both focused on simple choice strategies and analyzed how short they fall of maximization. They explored this question in the world of risk, in which each choice leads to one of a set of possible specified outcomes, and each outcome occurs with a known probability (Luce & Raiffa, 1957, p. 13)—that is, in a world in which there are no surprises.

2.3.1 Efficient Decision Heuristics and Measures of Success

Since the 1970s, many decision scientists have joined the quest to identify systematic biases originating from people's reliance on heuristics (Kahneman, Slovic, & Tversky, 1982). Ignoring the zeitgeist, Thorngate (1980) instead asked how simple choice strategies can be: How much information can they ignore and still permit successful choices? In one of the early computer simulations in the decision sciences, he orchestrated a computer tournament in which 10 choice heuristics competed against one another. The heuristics were tested in a randomly generated choice environment involving choice problems with two, four, or eight options, with each option offering two, four, or eight outcomes (for details, see Thorngate, 1980). Here, we focus on the performance of the two top-performing heuristics in Thorngate's competition—the equiprobable heuristic and the probable heuristic—and the worst-performing heuristic, the least-likely heuristic. The three heuristics' policies are outlined in box 2.1. Each of the heuristics ignores a different aspect of the available information and thus implements a different variant of cognitive simplification (Gigerenzer & Gaissmaier, 2011). The equiprobable heuristic ignores all probabilities, meaning that it does not multiply (weigh) outcomes by their probabilities. Like Dawes' (1979) "improper linear models," it simply calculates the arithmetic mean of all outcomes per option and chooses the option with the highest mean. It thus acts as if each outcome, no matter how small or large, is as probable as any other outcome. The equiprobable heuristic thus embodies the "principle of indifference," a coinage attributed to Keynes (1921/1973b). It states that whenever there is no evidence favoring one possibility over another, they have the same probability (see also chapter 5). The probable heuristic, in contrast, considers probabilities, but only to classify outcomes into two sets (probable vs. improbable outcomes); it then removes all improbable outcomes from consideration. The least-likely heuristic considers only one kind of outcome per option, namely, the worst possible outcome, and chooses the option with the smallest probability that this worst outcome will occur.

How high is the price that the heuristics pay for straying from the ideal of maximization and ignoring some or much of the available information?

Box 2.1
Heuristics for choices under risk and uncertainty.

We illustrate each heuristic's policy and choice prediction with reference to the following choice problem with four options:

A –50 with a probability of .2 and 250 with a probability of .8.
B –200 with a probability of .4 and 600 with a probability of .6.
C –400 with a probability of .1 and 500 with a probability of .9.
D 100 with a probability of .7 and 400 with a probability of .3.

The *equiprobable heuristic* calculates the arithmetic mean of all outcomes within each option and chooses the option with the highest mean. It chooses option *D*, because its mean (250) is higher than that of *A* (100), *B* (200), or *C* (50).

The *least-likely heuristic* identifies each option's worst outcome and selects the option with the lowest probability of the worst outcome. It chooses option *C*, where the probability of the worst outcome (–400) is .1, lower than in *A* (.2), *B* (.4), or *D* (.7).

The *lexicographic heuristic* determines the most likely outcome of each option and selects the option with the highest most likely outcome. If two or more outcomes are equal, it determines the second most likely outcome of each option and selects the option with the highest second most likely outcome. The process is continued until a decision is reached. It chooses option *B*, because it offers the highest outcome (600) among all options' most likely outcomes.

The *probable heuristic* categorizes each option's outcomes as "probable" (i.e., $p \geq .50$ for a two-outcome option and $p \geq .25$ for a four-outcome option) or "improbable" and ignores all "improbable" outcomes. It then calculates the arithmetic mean of the remaining outcomes and selects the option with the highest average outcome. It chooses option *B*, because its probable outcome (600) is higher than that of *A* (250), *C* (500), or *D* (100).

The *natural-mean heuristic* calculates the average of all outcomes sampled per option and divides the average by the number of sampled outcomes. It chooses the payoff distribution with the highest average outcome. This choice depends on the sample of experience. For instance, the heuristic may sample five times from each of the four options and encounter the following sequences:

A 250, 250, 250, –50, 250
B 600, 600, –200, –200, 600
C 500, 500, 500, 500, 500
D 100, 100, 400, 100, 400

The heuristic chooses option *C*, where the average of all outcomes (500) is highest.

Thorngate (1980) used the proportion of times a heuristic selected the option with the highest expected value (henceforth, the "best" option) as benchmark for its performance. Expected value is defined as

$$E(x) = \Sigma p_i x_i, \quad (1)$$

where p_i and x_i are the probability and the amount of money associated with each outcome ($i = 1, \dots, n$) of an option. Measured against this benchmark, all heuristics will, ignoring exceptional circumstances, perform worse than expected value theory. In Thorngate's tournament, all heuristics performed better than chance, but two heuristics stood out: averaged across all variants of the gambling environment, the equiprobable heuristic and the probable heuristic chose the best option in 75.4% and 75.2% of cases, respectively. From this, Thorngate concluded that a "wide variety of decision heuristics will usually produce optimal, or close to optimal, choice and can thus be termed relatively efficient" (p. 223). Given that the best heuristics' performance lagged behind optimal performance (i.e., expected value theory) by a hefty 25 percentage points, this conclusion seems somewhat overstated. Paradoxically, however, due to his choice of performance metric, Thorngate may in fact have underrated the heuristics' efficiency.

To examine this possibility, we replicated Thorngate's (1980) simulation (with 100,000 instead of 200 gambles) but used a different performance metric. Specifically, how costly the failure to choose the best alternative will be depends on the discrepancy between the expected values of the option chosen and the best option. If the difference is small, the costs of deviating from optimality will be minor; if the difference is large, they will be consequential. In our simulation, we expressed performance on the 0%–100% range defined by the sum of the expected values of the worst possible options and the sum of the expected values of the best possible options. On this metric, expected value theory will, by definition, score perfectly (100%), whereas random choice will perform at around 50%. In our replication, the equiprobable heuristic scored 94.1% and the probable heuristic, 94.6%. The performance gap between these heuristics and expected value theory was thus reduced to about 5–6 percentage points. In other words, in this simulation, heuristics evidently tend to err in cases where the consequences are relatively benign. Thorngate's measure therefore underestimated the heuristics' performance (even the worst-performing heuristic, least-likely, jumped from a meager 41.4% on Thorngate's metric to 65.7% on ours).

In sum, although Thorngate (1980) applauded the efficiency of the heuristics, they fell noticeably short of the ideal of expected value maximization on his performance measure (which used the "best" option as a benchmark). In contrast, on a performance measure that gauges how costly it is to fail to choose the best alternative when choosing heuristically, we found that simple choice policies fared substantially better. Furthermore, Thorngate's analysis did not address the role of environmental properties in fostering or impeding the heuristics' performance. For instance, how robust will the top performer, the equiprobable heuristic, prove to be if the environment's probability distribution is skewed, causing high variance in probability information and potentially rendering the assumption of equal weights dangerously inaccurate?

2.3.2 The Adaptive Decision Maker

The interaction between environmental structures and the information-processing architecture of heuristics was a focus of J. W. Payne et al.'s (1993) research program. Its premise was that people can select from a multitude of available strategies. Each strategy combines attractive properties (e.g., simplicity, low cognitive effort, accuracy) with unattractive ones (e.g., higher cognitive effort, lack of accuracy). The impact of these properties varies across choice environments and conditions such as time pressure and cognitive abilities. An adaptive decision maker considering candidate strategies for a task will select the one that affords the best trade-off between anticipated accuracy and effort. The assumption here is that there is an inescapable and law-like *accuracy–effort trade-off*, meaning that the less information, computation, or time a strategy requires and the decision maker invests, the less accurate (rational or optimized) the ensuing behavior will be.

J. W. Payne et al. (1988) used computer simulations to analyze this accuracy–effort trade-off in risky and riskless choice in a range of environments and conditions. We focus on the simulation of risky choice and on a single environmental condition: variance in probabilities. For example, one option may have four possible outcomes, with probabilities of .28, .25, .25, and .22, respectively. In this case, variance in probabilities is low. The outcomes of another option may have probabilities of .7, .2, .08, and .02; in this case, variance in probabilities is high. How did heuristics perform under high and low probability variance? Here, we focus on the equiprobable heuristic (or "equal-weight heuristic," to use the terminology of J. W. Payne et al., 1988), the best performer in the Thorngate (1980) tournament, and the lexicographic heuristic (LEX), the best performer in the J. W. Payne et al. tournament. Like the other heuristics, LEX ignores part of the information (see box 2.1). It is a noncompensatory strategy; note that the environmental cue that in this heuristic overrides all other cues and determines the choice depends on the properties of the choice problem at hand. Details of the simulations are provided in J. W. Payne et al. (1988).

Three key results emerged (J. W. Payne et al., 1988). First, heuristics can be highly competitive in one environment but fail in another. The lexicographic heuristic, for instance, was very successful in the high-variance environment, but its performance dropped by more than 20 percentage points in the low-variance environment. No generalist heuristic emerged that was able to perform consistently well in all environments. Second and relatedly, variance in probabilities affected processing policies differently. Whereas high probability variance fostered the performance of the lexicographic heuristic, it undermined the success of the equiprobable heuristic; low probability variance had the opposite effect. Third, when averaged across all environments, the equiprobable heuristic and the lexicographic heuristic chose the best option in 79% and 56% of cases, respectively (the performance criterion in this simulation was equivalent to the expected value criterion employed by Thorngate, 1980). That is, the two heuristics paid a substantial premium for simplicity, similar in magnitude to that observed in Thorngate's analysis: the top choice performance was more than 20 percentage points below that of the expected value benchmark.

Let us summarize the findings so far. The most successful choice heuristics in Thorngate (1980) and J. W. Payne et al. (1988) represent two very different paths to cognitive simplification (see box 2.1). Whereas the equiprobable heuristic considers all outcomes per option but neglects probabilities altogether, the lexicographic heuristic examines one cue at a time, ignoring all others. Furthermore, low variance in outcome probabilities supports a focus on outcomes, whereas high variance in probabilities is more compatible with a noncompensatory choice policy. These insights were obtained in the world of risk, where the probability and outcome space are known. What will be the trade-off between performance and simplicity—or, to use J. W. Payne et al.'s terminology, accuracy and effort—when knowledge is imperfect and surprises can happen? In other words, how do the heuristics perform in a world of incomplete knowledge and uncertainty? The bias–variance framework, originally developed in machine learning (Geman, Bienenstock, & Doursat, 1992; T. Hastie, Tibshirani, & Friedman, 2001) offers a conceptual approach for understanding the impact of uncertainty on prediction models.

2.4 The Bias–Variance Dilemma

Bounds on people's knowledge about the environment can be hugely consequential, to the extent that limited knowledge intensifies the *bias–variance dilemma* (e.g., Gigerenzer & Brighton, 2009; Katsikopoulos, Schooler, & Hertwig, 2010). To introduce this concept, which is relevant for any kind of prediction model, we offer an example. Bias and variance both contribute to the total error committed by any prediction model. Let us imagine that a prediction model is attempting to learn an underlying (true) function from a sample of (potentially noisy) data that was generated by this function. Averaged across all possible data samples of a given size, the bias of the algorithm is defined as the difference between the underlying function and the mean function computed by the algorithm from these data samples. Consequently, if this mean function is the same as the underlying function, bias will be zero. Variance reflects the sensitivity of the prediction model to different samples drawn from the same environment. High variance implies that the predictions of a model may differ greatly depending on the specific properties of the observed samples. This type of variance increases with model flexibility. For example, the more flexible the model, the more likely it will capture not only the true structure (assuming that there is a true structure and that the models are complex enough to capture it) but also unsystematic patterns, such as noise. Bias and variance both depend on the structure to be predicted (e.g., daily temperature across a year in a specific location; daily stock market fluctuation across a year of a specific index), and at least the variance also depends on how many sampled observations are available for this structure. Therefore, it will not always be adaptive to seek low bias in a prediction model by including as many adjustable components as possible to flexibly capture patterns in the sampled data. Model flexibility can itself become a curse when there is a high risk of increasing error through variance. From this, it follows that a model should be complex enough to avoid excessive bias, but simple enough to avoid overfitting idiosyncratic noise in the limited samples on which the estimates are made (e.g., Pitt, Myung, & Zhang, 2002).

The principles underlying the bias–variance dilemma can be applied to decisions about monetary lotteries under uncertainty—that is, when the task is to predict the value of the lotteries. Let us assume that choice strategies do not enjoy perfect knowledge of outcomes and probabilities but instead gauge them from samples drawn from the environment. Consequently, all strategies face two sources of error—and the possible trade-off between them. One source is error through variance. For illustration, expected value theory assumes that outcomes are weighted (multiplicatively) by their exact probabilities and then summed and maximized. Under conditions of imperfect information, other, simpler forms of information integration—for example, additive rather than multiplicative integration (Juslin, Nilsson, & Winman, 2009)—or forgoing integration altogether (e.g., lexicographic heuristic) may be more robust and less likely to suffer from overfitting. Yet simplifications can go too far, causing a substantial bias in the choice strategy and, consequently, prompting performance to deteriorate. How will these two sources of error shape the performance of choice heuristics under uncertainty? We examined this question by conducting a new set of simulations.

2.5 What Is the Price of Cognitive Simplicity in Choice under Uncertainty?

The benchmark used in our simulations is the performance of the expected value model under perfect knowledge. We call this model the omniscient expected value model. In addition, we used the same performance metric as in our reanalysis of the Thorngate (1980) competition. On this metric, 100% represents the sum of the expected values across choice problems in the case that a choice strategy always selects the option with the highest expected value (as the omniscient expected value model does); 0% represents the sum of the expected values across choice problems in the case that a strategy always selects the option with the lowest expected value. All heuristics were tested in environments in which the available information on the options' outcomes and probabilities was incomplete. Information was acquired through repeated sampling of monetary outcomes. We define a sample as the draw of a single monetary outcome from each option in a choice problem (with replacement). Each new draw thus offered information about the gambles' possible outcomes and their relative frequency of occurrence (i.e., probability). Furthermore, after each new sample, each heuristic rendered a choice, allowing us to analyze how the heuristics' performance changed as knowledge increased. In this simulated environment, the uncertainty the decision maker faced concerned both the outcome space (i.e., at any given point in time, the decision maker did not know for certain whether they were aware of the full outcome space) and the probabilities (i.e., the decision maker estimated the probabilities from the sequences of the encountered outcomes, meaning that the probabilities of possible outcomes that had not been encountered were unknown and, for known outcomes, the probabilities could only be estimated on the basis of the experienced sample of draws).

2.5.1 Competitors

We tested the equiprobable, the probable, and the lexicographic heuristics—the three top performers in the simulations by Thorngate (1980) and J. W. Payne et al. (1988)—as well as the least-likely heuristic, the worst-performing heuristic in Thorngate's competition. In addition, we tested the natural-mean heuristic (Hertwig & Pleskac, 2008, 2010), whose policy is described in box 2.1. This heuristic has some interesting characteristics. For one, it predicts the same choice as expected value theory if the latter also bases its choices on samples of experience rather than on perfect knowledge. It thus defines the level of accuracy of a sampling-based expected value theory (without Bayesian priors). However, the heuristic rests on a much simpler processing policy than expected value theory. Instead of multiplying each sampled outcome by its inferred (sample-based) probability and summing up the products, the heuristic simply totals up all experienced outcomes per lottery and then divides this sum by the sample size per option. In other words, it replaces the multiplicative core of expected value theory by simple summing and division, thus requiring no explicit representation of probabilities.

2.5.2 Environments

We implemented 20 choice environments, designed by combining five outcome distributions with four ways of constructing the associated probabilities (see also chapter 3). The outcome distributions consisted of a rectangular, a normal, an exponential, a Cauchy, and a lognormal distribution. These distributions permitted us to implement varying degrees of outcome variance. All five outcome distributions were symmetrical (for details, see figure S2.1 and table S2.1 in the online supplement at https://taming-uncertainty.mpib-berlin.mpg.de/). The four construction mechanisms producing the probabilities (i.e., rectangular, U-shaped, exponential, and skewed), henceforth *P-generators*, were chosen with the goal of obtaining different degrees of variance among probabilities (see J. W. Payne et al., 1993). As figure S2.2 in the online supplement illustrates, the four P-generators yielded markedly different probability distributions depending on the number of outcomes in the gamble. Table S2.2 in the online supplement quantifies the degrees of probability variance within options, giving a more direct measure of probability variance. Finally, each choice problem consisted of two, four, or eight options, with each option having the same number of outcomes (two, four, or eight). Here, we focus on the condition with two options per choice problem; this resulted in 60 sets of choice problems (i.e., 20 environments consisting of two, four, or eight outcomes).

2.5.3 Learning by Sampling

All heuristics learned about the properties of the choice problem in question by sequentially taking one draw at a time from each of the two options. Based on this information, the heuristics chose what they inferred to be the best option after each sample. The heuristics thus advanced from complete ignorance to progressively more knowledge and less uncertainty with each round of sampling. To experience the consequences of this type of sampling and heuristic choice for yourself, please visit interactive element 2.1 (at https://taming-uncertainty.mpib-berlin.mpg.de/). In our simulation, learning stopped after 50 rounds (resulting in 50 sampled outcomes for each option and 100 sampled outcomes for each choice problem with two options). For each of the 20 environments, 6,000 choice problems were randomly generated (2,000 each for choice problems containing options with two, four, and eight outcomes). Each heuristic therefore made $50 \times 6{,}000$ choices, amounting to 300,000 choices per environment (when a heuristic was unable to reach a choice, a random choice between the options was implemented). How did uncertainty and the successive reduction of uncertainty through learning affect the heuristics' performance? We first turn to how the heuristics fared in the environment that Thorngate (1980) analyzed (i.e., rectangular outcome and probability distributions); we then consider their performance across all 20 environments and, finally, examine the interaction between the specific simplifying assumptions made by each heuristic and its performance in specific environmental structures.

2.5.4 How Do Heuristics Fare When Uncertainty Rules?

Figure 2.1 plots the performance of the five competitors as a function of learning for choice problems involving two options. The figure plots the average expected value of the option chosen by each heuristic, relative to chance level (50%) and to the performance of omniscient expected value theory (100%). Let us highlight four observations. First, there was substantial variability in performance, with the least-likely heuristic again (as in Thorngate's, 1980, tournament) lagging far behind. Second, in some cases, performance increased little with learning—or even decreased (e.g., the lexicographic heuristic); in others, it increased substantially (e.g., the probable heuristic). Third, the natural-mean and the equiprobable heuristics clearly outperformed the lexicographic heuristic and, by a smaller margin, the probable heuristic. Fourth, when learning samples were small, the equiprobable heuristic performed as well as or even slightly better than the natural-mean

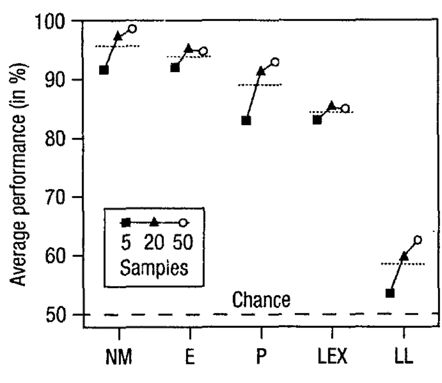

Figure 2.1
Average performance for each of the five heuristics (see box 2.1): the natural-mean (NM), equiprobable (E), probable (P), lexicographic (LEX), and least-likely (LL) heuristics in the environment that Thorngate (1980) implemented (rectangular outcome and probability distributions) as a function of learning (5, 20, and 50 samples per option) for choice problems involving two options (and averaged across two, four and eight outcomes). The performance benchmark for omniscient expected value theory is 100%. The dotted line represents the average performance across the three sample sizes.

heuristic; once more knowledge became available, the latter had the edge. Averaged across the number of outcomes, and with performance expressed relative to the range of best possible to worst possible performance (as in our reanalysis of the Thorngate environment), the equiprobable heuristic scored 92.1% (5 samples), 95.1% (20 samples), 94.8% (50 samples), with a mean of 94.0%; the natural-mean heuristic scored 91.7% (5 samples), 97.3% (20 samples), 98.6% (50 samples), with a mean of 95.9%. The latter represents the performance of expected value theory, assuming it lacks omniscience and has to make do with sampled slices permitting a noisy glimpse of the environment. Thus, the advantage of expected value theory decreases from 5–6 percentage points under conditions of perfect knowledge in our reanalysis of the simulation by Thorngate (1980) to 1.9 percentage points under conditions of imperfect knowledge. Will these findings generalize across environments beyond the specific environment that Thorngate chose for his simulations?

2.5.5 Generalization: Does the Performance Gap Close More Generally?

We next examined the heuristics across all 20 choice environments implemented (see section 2.5.2). We again focused on choice problems involving two options; the results for problems with more options were not qualitatively different. Figure 2.2 plots the heuristics' performance as a function of learning. As before, the natural-mean and the equiprobable heuristics were the frontrunners. Averaged across 5, 20, and 50 samples, the former scored 97.2% and the latter 93.8%—resulting in a difference of 3.4 percentage points, almost half the size of the gap (6 percentage points) found in our reanalysis of Thorngate's (1980) data. In particular, under high uncertainty (5 samples), the performance of the two heuristics was nearly identical: 93.2% vs. 93.9%. The next-best heuristic was again the probable heuristic (average score: 90.7%), followed by LEX (87.4%). These findings show that the equiprobable heuristic, which ignores all probabilities, lagged behind the natural-mean heuristic by just 3.4 percentage points when averaged across all sample sizes; when uncertainty was pronounced, their performance was nearly indistinguishable. As mentioned in section 2.5.1, the latter heuristic is equivalent to expected value theory without the gift of omniscience. In sum, across two simulations, we found that the advantage of expected value theory over simple heuristics shrank substantially, or even reversed slightly,

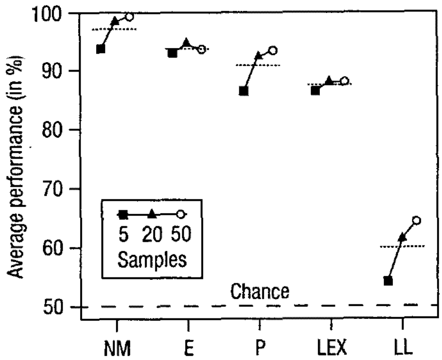

Figure 2.2
Average performance for each of the five heuristics (see box 2.1): the natural-mean (NM), equiprobable (E), probable (P), lexicographic (LEX), and least-likely (LL) heuristics across all 20 environments. The results are plotted as a function of learning (5, 20, and 50 samples per option) and are limited to choice problems with two options (and averaged across two, four and eight outcomes). The dotted lines represent the heuristics' performance averaged across 5, 20, and 50 samples.

when all strategies were tested under realistic assumptions of imperfect knowledge relative to perfect knowledge.

We next turn to Simon's (1956) second proposal: "The environments to which organisms must adapt possess properties that permit further simplification of its choice mechanisms" (p. 129). Taking advantage of the various environments we implemented, we are now, like J. W. Payne et al. (1993), able to examine the impact of distinct environmental structures on the heuristics' performance.

2.5.6 What Are Heuristics' Environmental Allies?

According to J. W. Payne et al.'s (1988) observations, variance in probability information is an environmental ally of the lexicographic heuristic but a foe of the equiprobable heuristic. Our multiple combinations of outcome distributions and probability mechanisms permitted us to systematically analyze the impact of variance (see table S2.2), separately for the dimensions of probability and outcome. In the following analyses, we focus on choice problems with two options and eight possible outcomes per option, making it easier to detect possible dependencies between environmental structures and heuristic policies. Moreover, we limit our analysis to two heuristics with different policies—lexicographic and equiprobable—and use the natural-mean heuristic as a benchmark; again, the latter is comparable to expected value theory in the context of sampling, where knowledge is limited.

Our first analysis combined the rectangular outcome distribution with the probability distributions obtained from the four P-generators: rectangular, U-shaped, exponential, and skewed (see figure S2.2 and details in the online supplement). Figure 2.3 plots the three heuristics' scores as a

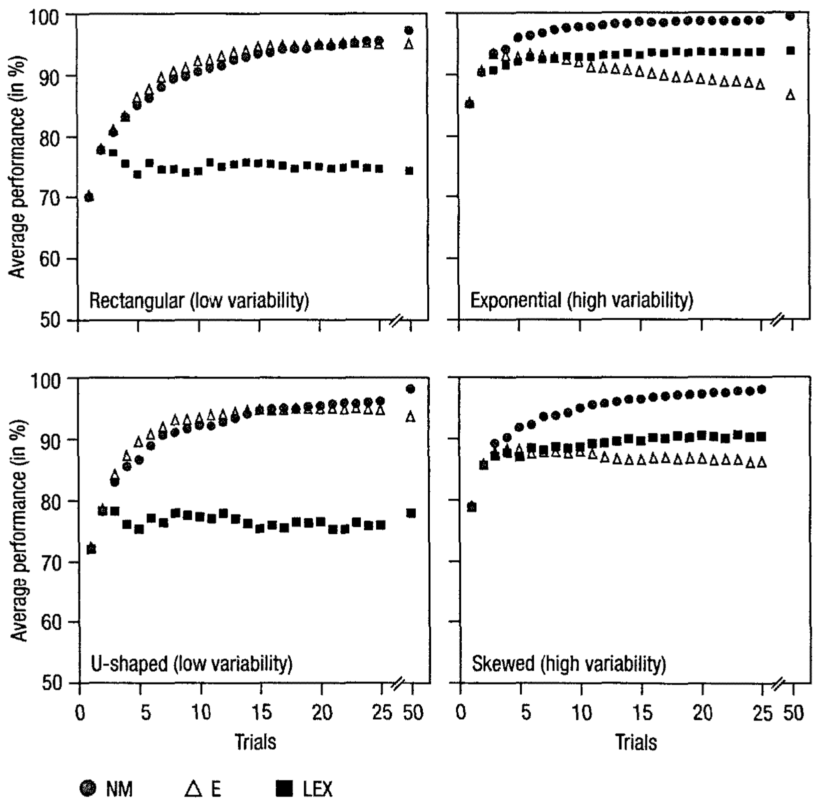

Figure 2.3
Average expected score for the natural-mean (NM), equiprobable (E), and lexicographic (LEX) heuristics as a function of learning ($n = 1-50$ samples per option) and of four construction principles for the probability information (rectangular, exponential, U-shaped, and skewed). In each case, the distribution of the outcome information is rectangular. The results refer to choice problems with two options and eight outcomes.

function of learning, by probability distribution. When variance in probabilities was relatively small—as was the case for the rectangular and the U-shaped P-generators—the equiprobable heuristic outpaced the natural-mean heuristic for small to medium sample sizes; in addition, the performance of the lexicographic heuristic did not improve with learning. The reason for this lack of learning is that the "most likely" outcome (LEX's top priority) tends to be experienced early in the learning process; more sampling tends not to yield previously unknown outcomes. When variance in probabilities was high, as was the case for the exponential and skewed P-generators, the natural-mean heuristic clearly outstripped the equiprobable and lexicographic heuristics. Under this condition, the accuracy of the equiprobable heuristic decreased with learning. This phenomenon is explained by the disproportional impact of rare outcomes. Specifically, a heuristic that keeps learning will eventually encounter relatively rare outcomes. This causes a problem for the equiprobable heuristic: no matter how rare an outcome is, the equiprobable heuristic will keep assigning it the same weight as it assigns to much more common outcomes. Consequently, the equiprobable heuristic "overweights" rare outcomes (see chapter 8). The lexicographic heuristic does not fall into the same trap, because its top priority is the most likely outcome.

Variance in probabilities is thus an environmental property that clearly impacts the performance of heuristics. The same does not hold for variance in outcomes. Following the same logic as in the previous simulation, in a second analysis we combined the rectangular P-generator with five outcome distributions: the rectangular, normal, exponential, Cauchy, and lognormal distributions (see figure S2.1 in the online supplement). The equiprobable heuristic outpaced the natural-mean heuristic across all distributions (and degrees of outcome variance) when uncertainty was pronounced (i.e., 20 samples or fewer). It was only after substantial learning that the natural-mean heuristic gained the upper hand. The lexicographic heuristic performed on a much lower level; in addition, it failed to benefit from learning. The reason is that, even in very small samples, LEX settles on the most likely outcome—and the most likely outcome will occur early on, irrespective of how much return it offers. Consequently, rare but (very) large outcomes that are experienced with more sampling are unlikely to lead to a different choice for LEX because they cannot replace the "most likely outcome," to which LEX gives priority.

Taken together, these results help to reveal the relative match (or lack thereof) of choice heuristics to environmental structures: we can thus identify environmental properties that support or hamper the heuristics' performance under uncertainty. First, when variance in probability information is relatively high, boldly ignoring all probabilities—that is, applying the equiprobable heuristic—is detrimental to performance (see figure 2.3). Under these conditions, the equiprobable heuristic risks attaching too much weight to rare events by giving each outcome the same weight. In contrast, noncompensatory heuristics such as the lexicographic heuristic, which shields itself against overweighting low-probability events, learn and perform better than a unit-weight policy. When, however, variance in probability information is relatively low, the equiprobable heuristic, which sums across all outcomes, offers the better simplification policy. For small samples, it even eclipses the natural-mean heuristic. In this environment, the noncompensatory policy of the lexicographic heuristic falls behind. Variance can also occur on the dimension of outcomes. The equiprobable heuristic proves utterly robust against this source of variance. Regardless of whether outcome variance is low or high, the policy of simply summing all (nonweighted) outcomes excels when probability variance is low. In addition, when uncertainty is high to medium (small samples), the equiprobable heuristic even eclipses the natural-mean heuristic.

2.6 Theoretical Realism and Imperfect Knowledge

As Nobel laureate Kenneth Arrow (1951) stressed, incomplete knowledge is the key property and condition of many real-world choices (see chapter 18). People have to choose between actions without being fully aware of the consequences. What theories of choice respect the limits of time, knowledge, and computational power under these ubiquitous environmental and psychological conditions? In Simon's (1956, 1983) view, the classical economic framework makes unrealistic demands on people. Instead, he suggested that most tasks are mastered by "approximate methods" (Simon, 1990, p. 6). However, there are costs to using these methods rather than the optimal model. When Simon first proposed the notion of bounded rationality, he stated that approximate methods are likely to fall far short of the ideal of maximizing as postulated in economic theory (Simon, 1956). He thus appears to have assumed what J. W. Payne et al. (1988, 1993) later portrayed as an accuracy–effort trade-off. The less information, computation, or time a decision maker invests, the less accurate their behavior (choice, inferences, and so on) will be. This appears to be an inescapable and law-like truth about the mind's decision-making machinery. Indeed, both Thorngate (1980) and J. W. Payne et al. (1993) observed evidence in support of this trade-off, notwithstanding the relatively high level of accuracy that some heuristics achieved in their computer tournaments. These tournaments did not, however, implement imperfect knowledge. Instead, the analyses were focused on "small worlds" (Savage, 1954), decisions from descriptions (Hertwig & Erev, 2009), and decisions under risk (as defined by Knight, 1921/2002, and Luce & Raiffa, 1957), where the possible outcomes and their probabilities are precisely known. Without doubt, such choice situations exist, but they may be the exception rather than the rule.

How pronounced is the accuracy–effort trade-off in environments characterized by uncertainty, imperfect information, and the need for search? We addressed this question in the context of monetary gambles, that portray decision making in terms of just two basic dimensions: outcomes and probabilities. Admittedly, this environment is somewhat odd, divorced as it is from any real content and context. Yet it would be difficult to overstate the role that monetary gambles have played in the development of normative and descriptive theories of choice (e.g., Allais, 1953; Bernoulli, 1738/1954; Kahneman & Tversky, 1979). Moreover, it is the context in which the ideal of rational decision making in terms of weighting and summing all pieces of information was originally conceived. From expected value theory and expected utility theory to the other descriptive neo-Bernoullian choice theories such as cumulative prospect theory (Tversky & Kahneman, 1992)—all suggest, once interpreted as information-processing theories, that people have to weight and sum. Heuristic policies that forgo weighting or summing can lead to surprisingly accurate performance; moreover, our simulation results (see figure 2.2) suggest that—to the extent that an accuracy–effort trade-off exists—it is less severe under incomplete knowledge (uncertainty) than under perfect knowledge (risk).

2.7 The Strategy Selection Problem and Uncertainty

Our results suggest an interesting twist to a key question: How does the mind select a heuristic from the adaptive toolbox? The simulations of choice heuristics by Thorngate (1980) and J. W. Payne et al. (1988), as well as those in this chapter, demonstrate one thing: there is great variability in the performance of heuristics, and choosing the wrong one in a given environment may prove costly. How does the mind figure out which heuristic to select? Selection could be guided by individual reinforcement learning (e.g., Rieskamp & Otto, 2006), by teaching (e.g., physicians, firefighters, and pilots are taught which cues to consult, in which order, and how to process them), by social learning (e.g., when procedures and strategies are copied from peers or other models), or by meta-inductive strategies that consider heuristics' past successes (Schurz & Thorn, 2016). The process might also involve systematic change and adaptation on an evolutionary time scale, as with rules of thumb for predation and mate search in animal species (Hutchinson & Gigerenzer, 2005). There is, however, another answer. The good average performance of the equiprobable heuristic under conditions of highly limited knowledge (see figure 2.2)—a result that held across environmental structures—suggests that some heuristics may offer good fallback options when an informed strategy selection is not possible. In the domain of inference, people have indeed been found to rely more on such unit-weighting strategies under conditions of greater uncertainty (i.e., when the cue hierarchy is unknown) than under conditions of lesser uncertainty (e.g., Pachur & Marinello, 2013).

2.8 Objectives beyond Expected Value Maximization

Our results from the preferential domain of monetary gambles demonstrate that the dynamic of the accuracy–effort trade-off differs, depending on whether performance is tested under "small world" conditions or under realistic conditions of incomplete knowledge and uncertainty. Similar results have been found in the inference domain (Gigerenzer & Brighton, 2009; Gigerenzer, Hertwig, & Pachur, 2011). However, let us emphasize one important point. Simplification in the choice process can go too far. Heuristics with a strong bias (see discussion of the bias–variance dilemma in section 2.4)—for instance, the least-likely heuristic (see figures 2.1 and 2.2)—fall far behind the best-performing heuristics. One could, however, argue that this heuristic achieves a very different objective than expected value theory—namely, it minimizes the risk of ending up with a bad outcome, or even the worst possible outcome. In future analyses, it will therefore be important to evaluate these heuristics against diverse objectives, including risk aversion. Future simulations can also examine the extent to which imperfection in the process of learning outcomes and probabilities, as well as the forgetting of experiences, affects the performance of normative and heuristic models of choice. We suspect that such factors might compromise the more complex normative models to a greater extent than the simpler heuristics, thus bringing them even closer together.

2.9 Do People Use the Simplification Policies Studied?

Models of heuristics are realistic insofar as they respect constraints such as imperfect knowledge, time pressure, and limited computational resources. But do people actually use these heuristics? Relatively little is known about the extent to which people employ the equiprobable, lexicographic, and natural-mean heuristics in choices under uncertainty (e.g., decisions from experience; Hertwig & Erev, 2009; for choices under risk, see Brandstätter, Gigerenzer, & Hertwig, 2006; Glöckner & Pachur, 2012; Pachur, Hertwig, Gigerenzer, & Brandstätter, 2013). However, the equiprobable and lexicographic heuristics have been observed to be used in the domain of inference (e.g., Bröder, 2011; Pachur & Marinello, 2013). Drawing on a meta-analysis of over 45,000 sampling sequences and subsequent decisions from experience (Wulff, Mergenthaler-Canseco, & Hertwig, 2018), Hertwig, Wulff, and Mata (2018) analyzed the ability of various heuristics to predict people's choices. Of the heuristics implemented here, the natural-mean heuristic was best at predicting people's choices, outperforming, for instance, cumulative prospect theory (Tversky & Kahneman, 1992; assuming existing sets of parameter values); on average, it correctly predicted 73.2% of choices (Hau, Pleskac, Kiefer, & Hertwig, 2008). This does not imply that the people in question used this or other heuristics to make those choices. But it suggests that heuristics represent a class of models that should always be taken into account when the aim is to provide a realistic account of how people make choices under uncertainty.

2.10 Revisiting Simon's Expectation

To conclude, the results of the pioneering computer tournaments by Thorngate (1980) and J. W. Payne et al. (1988, 1993) supported Simon's expectation that relying on simple decision strategies necessarily comes at a price: their choices will fall short of the ideal of maximization postulated in economic theory (Simon, 1956). Our simulations substantially qualify both Simon's expectation and the conclusions drawn from the past computer tournaments: in the context of more realistic choices, where the state of knowledge is not perfect, the performance gap between simple choice heuristics and maximization is much smaller than previously thought. It is considerably smaller under uncertainty than under risk, and under some environmental conditions heuristics even take the lead. This is an important result. Heuristics are applied not only because the mind's cognitive resources are inevitably limited when measured against the world's complexity. Equally important, the environment often deprives decision makers of the information required by computationally and informationally more complex strategies. The Olympian models—such as the expected utility theory as proposed by Bernoulli (1738/1954) and axiomatized by von Neumann and Morgenstern (1944/2007)—require perfect knowledge of probabilities. When objective probabilities are unknown or only imprecisely known, the decision maker therefore needs to estimate subjective ones, as in Savage's (1954) subjective expected utility theory. Or—and this is the alternative that emerges from our analyses, as exemplified by the good performance of both the natural-mean heuristic and the equiprobable heuristic—they can do without knowledge of probabilities. In many environments, doing without may be a better bet than running the risk of making consequential errors when estimating subjective probabilities out of thin air.