Chapter 8
How Representations of Knowledge Shape Actions

Ralph Hertwig and Renato Frey

In 2009 the world found itself in the midst of the worst recession since the Great Depression. Events thought of as extremely unlikely, such as the burst of the U.S. housing boom, the meltdown of the financial system, and the bankruptcy of colossal companies, happened in breathtakingly fast succession. Why was the world so badly prepared for these improbabilities? One explanation is that the crisis of the financial industry preceding the economic recession occurred because the industry's supposedly optimal risk-management models failed to reckon with "black swans" (Taleb, 2007)—unexpected and unpredictable rare events that carry an enormous impact. Of course, modern risk-management paradigms were not alone in failing to take the black-swan event into account—so did individual players, such as many homeowners who could no longer afford their mortgages. Can psychological theories and findings account for such blind spots?

At first glance, the answer is no. Influential studies in behavioral decision research consistently suggest the opposite propensity: People are oversensitive to rare events. For example, they overestimate the chance of getting food poisoning or of contracting lung cancer from smoking (Lichtenstein, Slovic, Fischhoff, Layman, & Combs, 1978; Viscusi, 2002). Moreover, people are depicted as remembering past experiences by how they felt at their peak (rare moment) and end (Redelmeier & Kahneman, 1996). Such oversensitivity is not only empirically observed but also theoretically suggested. According to the most influential descriptive theory of risky choice, people overweight low-probability events (Tversky & Kahneman, 1992). In fact, cumulative prospect theory explains the puzzling co-occurrence of two

R. Hertwig (✉)
Center of Adaptive Rationality (ARC), Max Planck Institute for Human Development, Berlin, Germany
e-mail: Hertwig@mpib-berlin.mpg.de

R. Frey
Center for Cognitive and Decision Sciences (CDS), Department of Psychology, University of Basel, Basel, Switzerland

P. Meusburger et al. (eds.), *Knowledge and Action*, Knowledge and Space 9,
DOI 10.1007/978-3-319-44588-5_8

behaviors—that the same people who purchase lottery tickets promising tiny chances of winning (thus being risk-seeking) also take out insurance against tiny chances of damage (thus being risk averse; Friedman & Savage, 1948)—on the assumption that small probabilities receive too much weight.

In light of people's ostensible oversensitivity to rare events, why did so many people, financial experts and laypeople alike, behave as though they were not cognizant of the rare events that triggered what some observers called a bona-fide depression (Posner, 2009)? Analyses have highlighted a variety of enabling factors, ranging from purportedly rational bankers who acted on strong incentives to take maximum risks in their lending (Posner, 2009) to humans' "animal spirits" (Akerlof & Shiller, 2009). However, there is another possibly enabling condition. The customary portrayal of humans as being oversensitive to rare events obscures the evidence that people, when recruiting their experience sampled across time to make risky decisions, tend to accord rare events (such as the burst of housing bubbles) less weight than they deserve according to their objective probabilities.

The Description–Experience Gap

Just as biologists use the *Drosophila* as one model organism, behavioral-decision researchers have used choice between monetary gambles as a model for risky choice, assuming that many real-world options have the same properties as gambles, namely, *n* outcomes and associated probabilities (Lopes, 1983). Moreover, many researchers have grown accustomed to presenting their respondents with one particular genus of the fruit fly: gambles in which all outcomes and their probabilities are stated and respondents make a single choice. Figure 8.1 illustrates typical description-based decision-making problems.

In everyday life, however, people can rarely peruse such descriptions of probability distributions—although there are a few exceptions, such as media weather forecasts stating probabilities of precipitation (Gigerenzer, Hertwig, Van Den Broek, Fasolo, & Katsikopoulos, 2005). When people decide whether to take out a loan or contemplate the success of a first date, there are no risk tables to consult. Instead, people need to rely on whatever experience they have had with these options, making decisions based on experience rather than on description (Hertwig, Barron, Weber, & Erev, 2004). Both kinds of decision can be understood as opposite poles on a continuum of uncertainty about what one is choosing between. In Knight's (1921) terminology, decisions from descriptions involve *a priori probabilities*, whereas decisions from experience involve *statistical probabilities*, which one must assess "if at all, by tabulating the results of experience" (p. 215), so they invariably fall short of the standards of accuracy set by a priori probabilities (Hau, Pleskac, & Hertwig, 2010).

In the 1950s and early 1960s, before modern behavioral-decision research, scientists who studied decision-making investigated decisions from experience. They examined, for example, whether and how people learn the probability structure of an outcome distribution through trial-by-trial feedback (for a review see Luce &

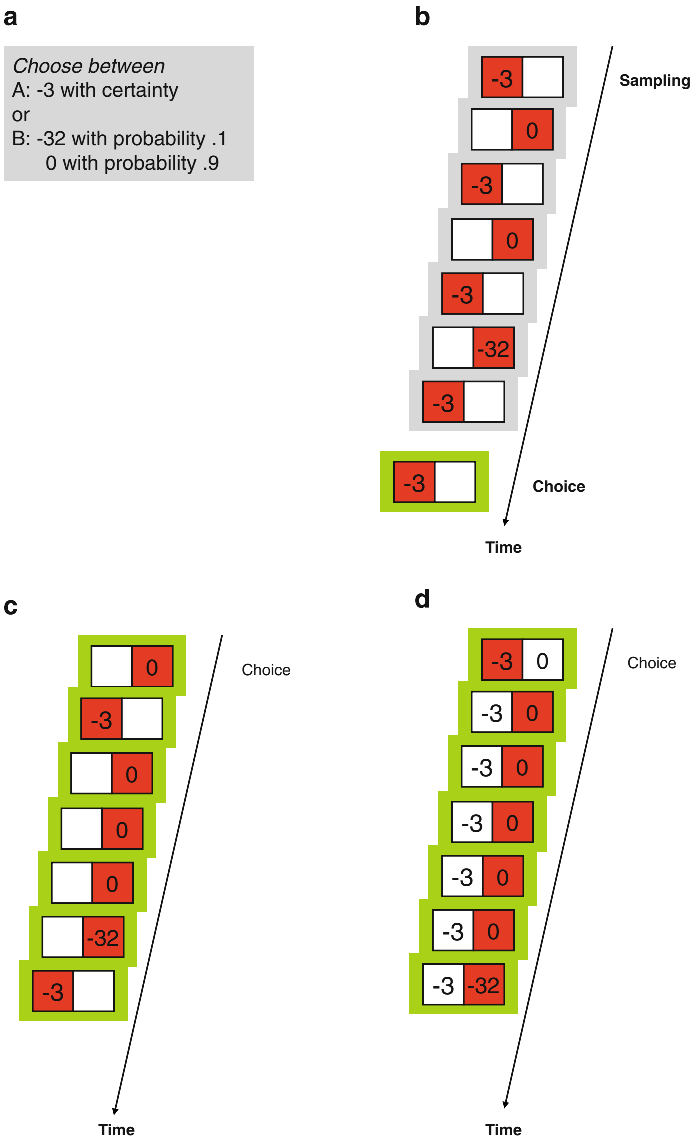

Fig. 8.1 How to study decisions from descriptions and experience. The choice task in decisions from description typically consists of two monetary gambles with explicitly stated outcomes and their probabilities (**a**). In decisions from experience, three paradigms have been employed. The *sampling paradigm* (**b**) consists of an initial sampling stage (here represented by seven fictitious draws) in which a person explores two payoff distributions without costs by clicking on one of the two buttons on the computer screen, followed by an outcome drawn from the respective distribution. The buttons chosen by a participant are marked in *red*. After terminating sampling, the person sees a choice screen (*green screen*) and is asked to select the button to draw once for real. The *partial-feedback paradigm* (**c**) combines sampling and choice, thus each draw represents both an act of exploration and an act of exploitation. The respondent receives feedback regarding the obtained payoff after each draw from the chosen button (*red box*). The *full-feedback paradigm* (**d**) is identical to the partial-feedback paradigm, except that it also provides feedback concerning the forgone payoff (i.e., the payoff that the person would have received had she chosen the other option; *white box*) (Reprinted from Hertwig and Erev (2009, p. 518) with permission from Elsevier)

Suppes, 1965). The impracticality of the research designs—purportedly hundreds of trials are needed before behavior stabilizes—may have been the reason that modern behavioral-decision research turned away from the transients of learning (for an exception see, for example, Busemeyer, 1985). Moreover, with the increasing importance of expected utility theory, the study of anomalies became pertinent, which required the conveying of perfect information about the probabilities of relevant events (Fig. 8.1a). Interest in issues of learning and experience-based decisions, however, remained alive in other fields, such as operation research (see literature on multiarmed bandit problems; Sutton & Barto, 1998).

Modern decision-making researchers' interest in decisions from experience has been rekindled by the recent observation of systematic and robust differences between them and decisions from description. Research on decisions from experience has come with a simple experimental tool, a "computerized money machine." Respondents see two buttons on a computer screen, each one representing an initially unknown payoff distribution. Clicking a button results in a random draw from the respective distribution. Three variations of this experimental tool have been employed. In the *sampling paradigm* (Fig. 8.1b), people first sample as many outcomes as they wish and only then decide from which distribution to make a single draw for real (Hertwig et al., 2004; Weber, Shafir, & Blais, 2004). In the *full-feedback paradigm* (Fig. 8.1d), there is a limited number of draws (typically 100), each of which contributes to people's earnings, and they receive draw-by-draw feedback on the obtained and the forgone payoffs (i.e., payoff received had the other option been selected; Yechiam & Busemeyer, 2006). The *partial-feedback paradigm* (Fig. 8.1c) is identical to the full-feedback paradigm, except that people learn about the obtained payoffs only (Barron & Erev, 2003; Erev & Barron, 2005). Unlike the first two paradigms, the partial-feedback paradigm presents respondents with an exploitation–exploration trade-off. Exploitation and exploration represent two alternative goals associated with every choice, namely, to obtain a desired outcome (exploitation) or to gather new information about other, perhaps better, actions (exploration; Cohen, McClure, & Yu, 2007).

Across all three experiential paradigms, a robust and systematic description–experience gap has emerged in numerous studies. Figure 8.2 illustrates this gap in six decision-making problems (Erev et al., 2010). Each one offers a choice between a risky option with two outcomes and a safe option. In the risky options, either the desirable outcome or the less desirable outcome occurs with low probability (.1 or less). In all three experiential paradigms, respondents tend to select the risky option when the desirable outcome occurs with high probability, and they select the safe option when the desirable outcome occurs with low probability. This tendency is reversed in decisions from description. The general pattern can be summarized as follows: In decisions from experience, people behave as if the rare events have less impact than they deserve according to their objective probabilities, whereas in decisions from description people behave as if the rare events have more impact than they deserve (consistent with cumulative prospect theory).

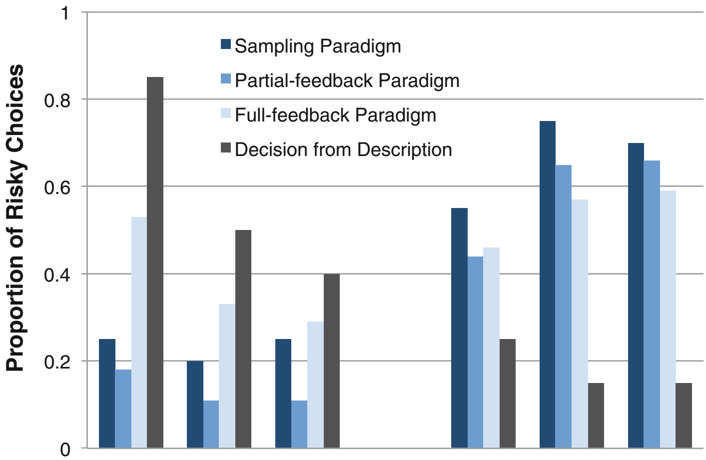

Problem	1	2	3	4	5	6
Safe option	7	-4.1	11.5	2.2	6.8	11
Risky option						
Maximum	16.5	1.3	25.6	3	7.3	11.4
p (maximum)	.01	.05	.10	.93	.96	.97
Minimum	6.9	-4.3	8.1	-7.2	-8.5	1.9
Expected value	7	-4.0	9.9	2.3	6.7	11.1

Fig. 8.2 Proportion of choices of the risky option as a function of the probability of the more desirable outcome in 6 of 120 problems studied in Erev et al. (2010). Each problem presents a choice between a risky option and a safe option. The decision-making problems and the expected values (EV) of the risky options are displayed below the plot. Each problem was studied using the four paradigms listed in Fig. 8.1 (Erev et al., 2010; the data from the full-feedback paradigm are unpublished). Participants (20 per paradigm) were paid (in shekels) for one of their choices, randomly selected. The partial- and full-feedback paradigms involved 100 choices per problem, and the reported proportions are the means over these choices and participants (Reprinted from Hertwig and Erev (2009, p. 519) with permission from Elsevier)

What Causes the Description–Experience Gap?

Several causes may be contributing to the description–experience gap.

Small Samples

Two classes of factors have been identified as shaping the search process in the sampling paradigm: properties of the decision-making problems (e.g., the magnitude of the incentives, see Hau, Pleskac, Kiefer, & Hertwig, 2008; and whether the outcomes are gains or losses, see Lejarraga, Hertwig, & Gonzalez, 2012) and individual characteristics, such as people's emotional state (Frey, Hertwig, & Rieskamp, 2014) or age (Frey et al. 2015). However, across numerous studies (reviewed in Hau et al., 2010), respondents typically proved restrained in their information search, with a median number of samples per choice problem typically ranging between 11 and 19. These results suggest that reliance on small samples is one factor that contributes to the attenuated impact of rare events (Hertwig et al., 2004). For small samples the chances are that a person does not even experience the rare events. More generally, one is more likely to undersample than oversample the rare event, for the binomial distribution of the number of times a particular outcome will be observed in n independent trials is markedly skewed when p is small (i.e., the event is rare) and n is small (i.e., few outcomes are sampled). Interestingly, reliance on small samples has also been discussed as a potential explanation for bumblebees' underweighting of rare events: Studying foraging decisions by bees in a spatial arrangement of flowers that promise with varying probabilities different amounts of nectar, Real (1991) concluded that "bumblebees underperceive rare events and overperceive common events" (p. 985). He explained this distortion in bees' probability perception as a consequence of their sampling behavior—"bees frame their decisions on the basis of only a few visits" (Real, 1992, p. 133)—and suggested that such reliance on small samples can be adaptive.

> Short-term optimization may be adaptive when there is a high degree of spatial autocorrelation in the distribution of floral rewards. In most field situations, there is intense local competition among pollinators for floral resources. When "hot" and "cold" spots in fields of flowers are created through pollinator activity, then such activity will generate a high degree of spatial autocorrelation in nectar rewards. If information about individual flowers is pooled, then the spatial structure of reward distributions will be lost, and foraging over the entire field will be less efficient. In spatially autocorrelated environments ("rugged landscapes"), averaging obscures the true nature of the environment. (p. 135)

Could there be any advantage to frugal sampling in experience-based decisions by humans? Hertwig and Pleskac (2008, 2010) proposed one possible advantage that rests on the notion of amplification. Unlike Real (1992), however, they argued that amplification proffers a cognitive rather than an evolutionary benefit. Through mathematical analysis and computer simulation, Hertwig and Pleskac (2010) showed that small samples amplify the difference between the options' average rewards. That is, drawing small samples from payoff distributions results in experienced differences of sample means that are larger than the objective difference. Such amplified absolute differences simplify the choice between gambles and thereby explain the frugal sampling behavior observed in investigations of decisions from experience—a conjecture for which Hertwig and Pleskac (2010) found empirical evidence.

The explanation of the description–experience gap in terms of small samples has prompted a critical response (Fox & Hadar, 2006) and has led to an ongoing debate. What appears to be underweighting of rare events in decisions from experience could be consistent with overweighting of low probabilities as assumed in cumulative prospect theory. When the probability experienced in a sample is smaller than the event's objective probability, people may still overweight this sample probability.

Despite this overweighting, the erroneous impression of underweighting would emerge if the *overweighting* did not fully compensate for the *underestimation* that results from the skew in small samples. In this view the description–experience gap is statistical (sampling error) rather than psychological in nature.

Several approaches have been taken to examine whether the gap observed in the sampling paradigm can indeed be reduced to sampling error. If sampling error was the sole culprit, then reducing the error by extending the sample should attenuate and eventually eliminate the gap. Increasing sample sizes substantially (up to 50 and 100 draws per choice problem) reduced but did not eliminate the gap (Hau et al., 2008, 2010). If sampling error caused the gap, then removing the error by aligning the sample's experienced probabilities to the objective probabilities should eliminate it. It did not (Ungemach, Chater, & Stewart, 2009). If sampling error was the sole root of the gap, then presenting respondents in the description condition the same information that others experienced (*yoking*) should eliminate the gap. In one study it did (Rakow, Demes, & Newell, 2008); in another it did for small samples but not for large ones (Hau et al., 2010; see these authors' discussion of trivial choices as one possible explanation for the mixed results obtained). The gap persisted even when people were presented both descriptions and experience rather than descriptions only (Jessup, Bishara, & Busemeyer, 2008).

In summary, the reality of the description–experience gap across the three experiential paradigms is unchallenged—its cause, however, is disputed. Some researchers have argued that the gap in the sampling paradigm is statistical in nature (Fox & Hadar, 2006; Hadar & Fox, 2009; Rakow et al., 2008); others have proposed that the sampling error is not the sole cause (Hau et al., 2008, 2010; Hertwig et al., 2004; Ungemach et al., 2009). Regardless of how this debate will advance, it is informative to go beyond the sampling paradigm. Reliance on small samples, for example, cannot be the reason behind the description–experience gap in the full-feedback paradigm (Fig. 8.1d) paradigm, in which the impact of rare events is attenuated even after a hundred trials with perfect feedback. Beyond sampling error, what psychological factors may be in play?

Recency

A psychological factor that may contribute to the description–experience gap is *recency* (Hertwig et al., 2004). Ubiquitously observed in memory, belief updating, and judgments (Hogarth & Einhorn, 1992), recency refers to the phenomenon that observations made late in a sequence receive more weight than they deserve (i.e., more than $1/n$). Recency is closely related to reliance on small samples: The small sample of recent events can reintroduce the aforementioned skew into large samples of experience. Although the original finding was that people give more weight to recent than to previous outcomes in the flow of their experience (Hertwig et al., 2004), little or no impact of recency was observed in later studies (Hau et al., 2010; Rakow et al., 2008; Ungemach et al., 2009).

Estimation Error

In theory, the description–experience gap could also be the result of a systematic estimation error (Fox & Hadar, 2006), with people systematically underestimating the frequencies of the rare event experienced in the sample. Studies of frequency and probability assessments, however, commonly report overestimation of rare events (Hertwig, Pachur, & Kurzenhäuser, 2005; Lichtenstein et al., 1978). Moreover, studies recording people's estimates of rare events in the sampling paradigm found them to be well calibrated or a little too high relative to the experienced frequency (Hau et al., 2008; Ungemach et al., 2009). That is, people do not systematically estimate rare things to be even rarer than they statistically are.

Contingent Sampling

Still another factor that could underlie the description–experience gap, especially in the feedback paradigm, is the notion that people inform their decisions by recruiting recent and past experiences garnered in similar situations (for related notions see Gilboa & Schmeidler, 1995; Gonzalez, Lerch, & Lebiere, 2003). Such contingent sampling is likely to be ubiquitous in the wild (Klein, 1999). For example, when firefighters need to predict the behavior of a fire, they appear to retrieve from memory similar instances from the past. Contingent sampling implies recency and reliance on small sampling to the extent that similarity decreases with time. Furthermore, in dynamic environments (e.g., the restless bandit problem; Whittle, 1988), reliance on similar experiences is an efficient heuristic (Biele, Erev, & Ert, 2009). Below, we turn to the manner in which the process of contingent sampling can be modeled.

Spatial Search Policies

Like any organism, humans can sample information in at least two very different ways from payoff distributions (e.g., flowers, ponds, other people, and gambles). Figure 8.3 depicts two paradigmatic sequential-sampling strategies based on two assumed options. In piecewise sampling, the searcher oscillates between options, each time drawing, in the most extreme case, the smallest possible sample. In comprehensive sampling, by contrast, the searcher samples extensively from one option and then turns to the other option to explore it thoroughly.

Taking these two sampling strategies as a starting point, Hills and Hertwig (2010) suggested that this spatial way of sampling foreshadows how people make their final decision. Specifically, they proposed that a person who samples piecewise will tend to make decisions as would a judge who scores each round of a boxing match: She determines which option yields the better reward in each round of sampling and ultimately picks the one that wins the most rounds. By contrast, a person using a comprehensive-sampling strategy will tend to gauge the average reward for each

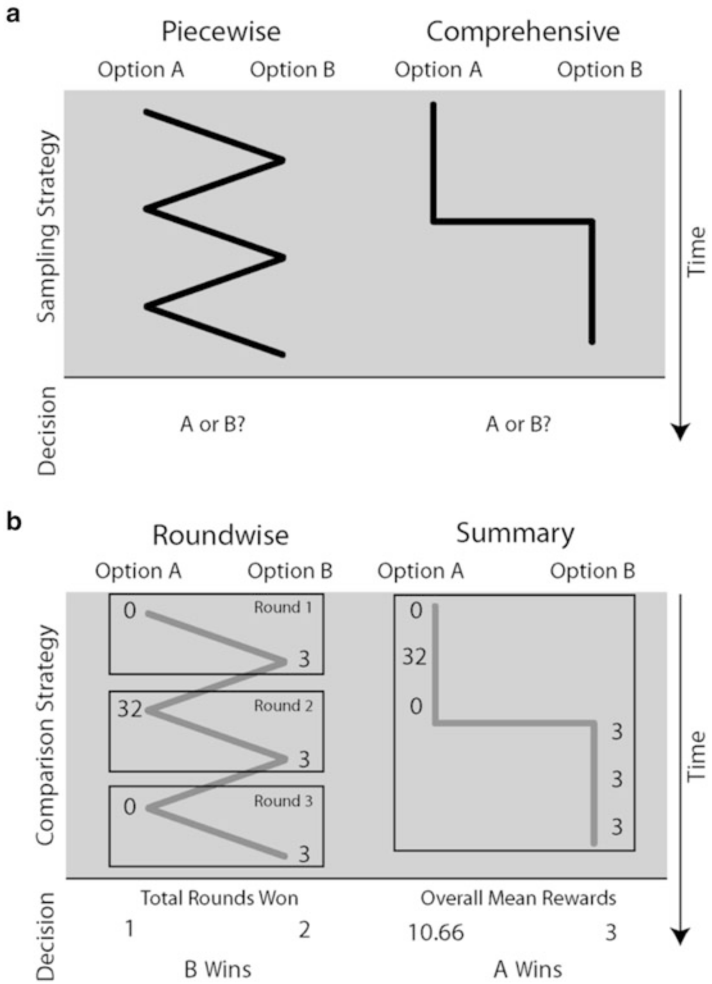

Fig. 8.3 (**a**) Representations of the sampling patterns associated with piecewise- and comprehensive-sampling strategies. Piecewise strategies repeatedly alternate back and forth between options. Comprehensive-sampling strategies take one large sample from each option. Following the sample phase, the participants make a decision about which option they prefer. (**b**) Representations of the comparison strategies associated with roundwise and summary strategies for a set of hypothetical outcomes. Roundwise strategies compare outcomes over repeated rounds and choose options that win the most rounds. Summary strategies compare final values (here, the overall expected value) and choose options with the better final value (Reprinted from Hills and Hertwig (2010, p. 1788) with permission from Associations for Psychological Science)

option and then choose the one promising the larger reward harvest. The reason for this dependency of the decision strategies on search is that the piecewise- and comprehensive-sampling strategy foster comparisons across different scales of information: rounds vs. summaries, respectively. Determining a winner who is ahead in most rounds and determining the one yielding the largest expected reward can lead to different choices even when both decisions-makers experience the same information. The reason is that the person using the former decision strategy weighs each round equally, ignores the magnitude of wins and losses, and thus acts as if it underweights rare, but consequential, outcomes. That link between sampling strategy and decision strategy is exactly what Hills and Hertwig (2010) found. Individuals who frequently oscillated between options were more likely to choose the round-wise winning options and to make choices as if they underweighted rare events than were individuals who switched options rarely.

In summary, modern behavioral decision research has been strongly focused on people's responses to descriptions of events. In recent years three experiential paradigms have been used to study how experience affects risky choice. A consistent picture has emerged. When rare events are involved, description-based and experience-based decisions can drastically diverge. We now turn to different ways of modeling decisions from experience.

Cognitive Strategies in Decisions from Experience

In attempting to capture the information search (learning) and decision-making processes in decisions from experience, researchers have proposed models that can be grouped into three classes. The first class—neo-Bernoullian models—rests on the premise that respondents form a mental representation of the relative frequency (probability) with which events occur in the process of sampling outcomes. Combined with outcome information, these probabilities then enter the evaluation of the two gambles' desirability. But do decisions from experience inevitably give rise to an explicit representation of probabilities? The second and the third class of models—associative learning models and heuristics—reflect the assumption that decision-makers can and will do without probabilities. In this section we discuss the three classes of models.

Neo-Bernoullian Models

Expected utility theory postulates that one can, or should, model human choice by assuming that people behave as if they have multiplied some function of probability and value and then have maximized it. Applied to decisions from experience, expected utility theory and related models require explicit representation of probabilities. An example is the "two-stage model" (Tversky & Fox, 1995, p. 279) of decision under uncertainty, in which it is assumed that decision-makers first estimate the probability p of an uncertain event A and then make a choice. The psychological impact of the event A with its associated (estimated) probability p is then measured in terms of cumulative prospect theory's probability weighting function π (Fox & Tversky, 1998).

Associative Learning Models

In this class of theories, human choice is conceptualized as a learning process (Busemeyer & Myung, 1992; Bush & Mosteller, 1955). Learning consists in changing the propensity to select a gamble according to the experienced outcomes. Good experiences boost the propensity of choosing the gamble associated with them, and bad experiences diminish it (e.g., Barron & Erev, 2003; Denrell, 2007; Erev & Barron, 2005; March, 1996). Two associative-learning models that have been proposed to capture decisions from experience are the value-updating model (Hertwig, Barron, Weber, & Erev,2006) and the instance-based learning (IBL) model (Gonzalez & Dutt, 2011).

The value-updating model stipulates that learners update their estimates of the value of the gamble after each new draw from it. Specifically, the model computes the weighted average of the previously estimated value and the value of the most recently experienced outcome. The model includes two parameters, namely, the number of draws and a recency parameter. The former parameter is determined empirically; the second is adjustable (i.e., fitted to the data). Importantly, the model does not necessitate representation of probabilities. Furthermore, the best fitting parameter in a model competition indeed suggested a substantial recency effect (Hau et al., 2008).

The IBL model also stipulates a learning process but goes beyond the relatively simple assumptions of the value-updating model: It is assumed that a choice (given that it is not automatically reproduced) represents the selection of the option with the higher utility (blended value). An option's blended value is a function of its associated outcomes and the probability of retrieving corresponding instances from memory (contingent sampling). Memory retrieval depends on memory activation, which, in turn, is a function of the recency and frequency of the experience. Activation is specified by the mechanism originally proposed in Adaptive Control of Thought—Rational (ACT-R; Anderson & Lebiere, 1998), a cognitive architecture used by cognitive psychologists to model problem-solving, learning, and memory. The IBL model is particularly attractive because it "predicts not only the final consequential choice but also the sequence of sampling selection" (Gonzalez & Dutt, 2011, p. 529; but see Hills & Hertwig, 2012) and because it offers a single learning mechanism (leading up to an instance's activation) across all experiential designs (Fig. 8.1b–d). Indeed, in a quantitative comparison of models, Gonzalez and Dutt (2011) were able to show that the IBL model predicts final experience-based decisions as well as or better than any other proposed model (including, for instance, the value-updating model and cumulative prospect theory).

Heuristics

Another class of models designed to describe both the process and outcome of choice are cognitive choice heuristics (see Brandstätter, Gigerenzer, & Hertwig,2006). Heuristics can be separated into two classes: those that use solely outcome information and exclude probabilities (outcome heuristics), and those that use at least rudimentary probability information (dual heuristics). Outcome heuristics such as maximax and minimax (Luce & Raïffa, 1957; Savage, 1954) were originally proposed as models for decision-making under ignorance in which people have no information whatsoever about probabilities.

Another cognitive heuristic that focuses on outcomes is the natural-mean heuristic (Hertwig & Pleskac, 2008). It works in two steps:

Step 1. Calculate the natural mean of outcomes for both gambles by summing, separately for each gamble, all *n*-experienced outcomes and then dividing by *n*.
Step 2. Choose the gamble with the larger natural mean (i.e., the gamble that had the best average outcome in the sampling phase).

The natural-mean heuristic was originally proposed in the context of *n*-armed bandit problems (Sutton & Barto, 1998) as a simple method for estimating the values of actions (e.g., the play of one of a slot machine's levers) and for using the estimates to select between actions: "The true value of an action is the mean reward received when the action is selected. One natural way to estimate this is by averaging the rewards actually received when the action was selected" (p. 27). The natural-mean heuristic totes up all experienced rewards (or losses) per gamble and then divides this sum by the sample size per gamble to arrive at the *natural mean*. One interpretation of the natural-mean heuristic is that in decisions from experience it is a simple and psychologically plausible instantiation of the expected-value calculus—particularly in continuous outcome distributions. Indeed, the natural-mean heuristic was not inferior to the more complex models described above and predicted a comparable number of correct predictions in decisions from experience (Hau et al., 2008).

In light of these models that do not require explicit representations of probabilities, we return to the question of what the possible codeterminants of the gap between description and experience are. The two associative-learning models and the natural-mean heuristic are format dependent. That is, they cannot capture decisions from description, for the input into these models consists of a sequence of outcomes that get integrated into one summary measure. They have no conceptual parameters with which to take probability information into account, and, in fact, probabilities are not directly apparent in decisions from experience. In decisions based on description, however, probabilities are made explicit to decisions-makers. Differences in description- and experience-based choices could therefore arise partly because different formats of mathematically equivalent information trigger different cognitive strategies (see Gigerenzer & Hoffrage, 1995, for a related argument in Bayesian reasoning).

Decisions from Experience: A Key to Otherwise Puzzling Human Behavior

The most famous eruption of Mount Vesuvius occurred in 79 AD, destroying many neighboring towns, among them Pompeii, the luxurious resort of wealthy Romans and now the most renowned still life of volcanic doom. This eruption, however, was not the most devastating one. As recent volcanological and archaeoanthropological studies have revealed, an earlier, Bronze Age eruption (around 3780 BC) covered the surrounding area as far as 25 km away, burying land and villages, causing a global climatic disturbance and the abandonment of the entire area for centuries. The loss of life and property was less extensive in the Bronze Age cataclysm than in the eruption of AD 79, but researchers recently discovered evidence of a mass exodus: a huge number of human and animal footprints pressed into the ash bed and all leading away from the volcano (Mastrolorenzo, Petrone, Pappalardo, & Sheridan,2006).

At present, at least three million people live within the area that was destroyed by the Bronze Age eruption. In fact, the periphery of Mount Vesuvius, which includes a significant chunk of the Naples metropolitan area, is among the most populated of any active volcano (Bruni, 2003). According to simulations by Mastrolorenzo et al. (2006), an eruption comparable in magnitude to the Bronze Age eruption would cause total devastation and mortality within a radius of at least 12 km (7½ miles). In addition, great quantities of fine ash in more distant zones might cause severe respiratory-tract injuries and fatalities due to acute asphyxia. Although it is impossible to predict the exact probability of such a catastrophe happening, volcanologists such as Michael Sheridan have argued that roughly 2000 years have passed since Pompeii's destruction and that "with each year, the statistical probability increases that there will be another violent eruption" of Vesuvius (Wilford, 2006). In light of these dire forecasts, one might expect that local residents would be keen to move away from the danger zone. On the contrary, relocating residents has proven extremely difficult, despite considerable incentives offered by the regional authorities. "In the shadow of Vesuvius, those residents have cultivated a remarkable optimism, a transcendent fatalism and a form of denial as deep as the earth's molten core" (Bruni, 2003, par. 12).

How can one explain the willingness of residents to defy fate? Perhaps it has become clear by now why the distinction between description-based and experience-based decisions may be key to understanding this and other puzzling risk-taking behavior. Personal experience tells residents in the vicinity of Mount Vesuvius that violent eruptions are extremely rare; in fact, in most people's lifetime, they have been nonexistent. Unless catastrophes have occurred recently, the relative indifference with which citizens and politicians often consider rare, but high-consequence, events like bursting levies, catastrophic earthquakes, and eruptions of volcanoes may be owed to the experience of their rarity (Weber, 2012). Just as residents in the vicinity of Mount Vesuvius have ignored incentives to relocate, people living in flood plains who make decisions about insurance based on their personal experience with floods—a rare event—have tended to turn down even federally subsidized flood insurance (Kunreuther, 1984).

At the same time, experiencing a rare, but highly consequential, event in reality can also have a lasting psychological impact. This possibility brings the discussion full circle. Generations growing up in a period of low stock returns appear to take an unusually cautious approach to investing, even decades later. In other words, young people who experienced the dramatic economic slump of 2008–2009 may enter the stock and housing market much more cautiously than their parents did.

Conclusion

Modern behavioral decision research has commonly focused on decisions from description. The observations stemming from this research suggest that humans overestimate and overweight rare events. Recent research on risky choice that takes into account the role of experience has found that people behave as if rare events are accorded less weight than they deserve relative to their objective probabilities. These observations are not contradictory; they describe how the mind functions in two different informational environments. In other words, research on description-based behavior and research on experience-based behavior should not be played against each other—their contrast is enlightening. However, to improve the understanding of how people make decisions with incomplete and uncertain information "in the wild" and how people respond to events that are rare but highly consequential, it is necessary to study the psychology and rationality of people's decisions from experience.

Acknowledgments This text is based largely on Hertwig and Erev (2009) and Hertwig (2012). We thank Elsevier, Sage, and Springer for granting the permission rights, and the Swiss National Science Foundation for a grant to the first author (100014–126558).

References

Akerlof, G. A., & Shiller, R. J. (2009). *Animal spirits: How human psychology drives the economy, and why it matters for global capitalism*. Princeton: Princeton University Press.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah: Lawrence Erlbaum.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16,* 215–233. doi:10.1002/bdm.443

Biele, G., Erev, I., & Ert, E. (2009). Learning, risk attitude and hot stoves in restless bandit problems. *Journal of Mathematical Psychology, 53,* 155–167. doi:10.1016/j.jmp.2008.05.006

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

Bruni, F. (2003, August 26). San Giuseppe Vesuviano Journal: Who's afraid of Vesuvius? (Pompeii is history). *New York Times*. Retrieved from http://www.nytimes.com/2003/08/26/world/san-giuseppe-vesuviano-journal-who-s-afraid-of-vesuvius-pompeii-is-history.html

Busemeyer, J. R. (1985). Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 538–564. doi:10.1037/0278-7393.11.3.538

Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General, 121,* 177–194. doi:10.1037/0096-3445.121.2.177

Bush, R. R., & Mosteller, F. (1955). *Stochastic models for learning.* Oxford: John Wiley & Sons.

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences, 362,* 933–942. doi:10.1098/rstb.2007.2098

Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review, 114,* 177–187. doi:10.1037/0033-295X.114.1.177

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112,* 912–931. doi:10.1037/0033-295X.112.4.912

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R.,…Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making, 23,* 15–47. doi:10.1002/bdm.683

Fox, C. R., & Hadar, L. (2006). "Decisions from experience" = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making, 1,* 159–161.

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

Frey, R., Hertwig, R., & Rieskamp, J. (2014). Fear shapes information acquisition in decisions from experience. *Cognition, 132,* 90–99. doi:10.1016/j.cognition.2014.03.009

Frey, R., Mata, R., & Hertwig, R. (2015). The role of cognitive abilities in decisions from experience: Age differences emerge as a function of choice set size. *Cognition, 142,* 60–80. doi:10.1016/j.cognition.2015.05.004

Friedman, M., & Savage, L. J. (1948). The utility analysis of choices involving risk. *Journal of Political Economy, 56,* 279–304. doi:10.1086/256692

Gigerenzer, G., Hertwig, R., Van Den Broek, E., Fasolo, B., & Katsikopoulos, K. V. (2005). "A 30% chance of rain tomorrow": How does the public understand probabilistic weather forecasts? *Risk Analysis, 25,* 623–629. doi:10.1111/j.1539-6924.2005.00608.x

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review, 102,* 684–704. doi:10.1037/0033-295X.102.4.684

Gilboa, I., & Schmeidler, D. (1995). Case-based decision theory. *Quarterly Journal of Economics, 110,* 605–639. doi:10.2307/2946694

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review, 118,* 523–551. doi:10.1037/a0024558

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science, 27,* 591–635. doi:10.1016/S0364-0213(03)00031-4

Hadar, L., & Fox, C. R. (2009). Information asymmetry in decision from description versus decision from experience. *Judgment and Decision Making, 4,* 317–325.

Hau, R., Pleskac, T. J., & Hertwig, R. (2010). Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *Journal of Behavioral Decision Making, 23,* 48–68. doi:10.1002/bdm.665

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description–experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making, 21,* 493–518. doi:10.1002/bdm.598

Hertwig, R. (2012). The psychology and rationality of decisions from experience. *Synthese, 187,* 269–292. doi:10.1007/s11229-011-0024-4

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science, 15,* 534–539. doi:10.1111/j.0956-7976.2004.00715.x

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2006). Decisions from experience: Sampling and updating of information. In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition* (pp. 71–91). New York: Cambridge University Press.

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

Hertwig, R., Pachur, T., & Kurzenhäuser, S. (2005). Judgments of risk frequencies: Tests of possible cognitive mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35,* 621–642. doi:10.1037/0278-7393.31.4.621

Hertwig, R., & Pleskac, T. J. (2008). The game of life: How small samples render choice simpler. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 209–235). Oxford: Oxford University Press.

Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition, 115,* 225–237. doi:10.1016/j.cognition.2009.12.009

Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science, 21,* 1787–1792. doi:10.1177/0956797610387443

Hills, T. T., & Hertwig, R. (2012). Two distinct exploratory behaviors in decisions from experience: Comment on Gonzalez & Dutt, 2011. *Psychological Review, 119,* 888–892. doi:10.1037/a0028004

Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology, 24,* 1–55. doi:10.1016/0010-0285(92)90002-J

Jessup, R. K., Bishara, A. J., & Busemeyer, J. R. (2008). Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science, 19,* 1015–1022. doi:10.1111/j.1467-9280.2008.02193.x

Klein, G. A. (1999). *Sources of power: How people make decisions*. Cambridge: MIT Press.

Knight, F. H. (1921). *Risk, uncertainty, and profit*. Boston: Houghton Mifflin.

Kunreuther, H. (1984). Causes of underinsurance against natural disasters. *The Geneva Papers on Risk and Insurance, 9,* 206–220.

Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition, 124,* 334–342. doi:10.1016/j.cognition.2012.06.002

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., & Combs, B. (1978). Judged frequency of lethal events. *Journal of Experimental Psychology: Human Learning and Memory, 4,* 551–578. doi:10.1037/0278-7393.4.6.551

Lopes, L. L. (1983). Some thoughts on the psychological concept of risk. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 137–144. doi:10.1037/0096-1523.9.1.137

Luce, R. D., & Raïffa, H. (1957). *Games and decisions: Introduction and critical survey*. New York: Dover.

Luce, R. D., & Suppes, P. (1965). Preference, utility, and subjective probability. In R. Duncan, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology, Vol. III* (pp. 249–410). New York: John Wiley & Sons.

March, J. G. (1996). Learning to be risk averse. *Psychological Review, 103,* 309–319. doi:10.1037/0033-295X.103.2.309

Mastrolorenzo, G., Petrone, P., Pappalardo, L., & Sheridan, M. F. (2006). The Avellino 3780-yr-BP catastrophe as a worst-case scenario for a future eruption at Vesuvius. *Proceedings of the National Academy of Sciences of the United States of America, 103,* 4366–4370. doi:10.1073/pnas.0508697103

Posner, R. A. (2009). *A failure of capitalism: The crisis of '08 and the descent into depression*. Cambridge: Harvard University Press.

Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior and Human Decision Processes, 106,* 168–179. doi:10.1016/j.obhdp.2008.02.001

Real, L. A. (1991). Animal choice behavior and the evolution of cognitive architecture. *Science, 253,* 980–986. doi:10.1126/science.1887231

Real, L. A. (1992). Information processing and the evolutionary ecology of cognitive architecture. *American Naturalist, 140,* 108–145. doi:10.1086/285399

Redelmeier, D. A., & Kahneman, D. (1996). Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain, 116,* 3–8. doi:10.1016/0304-3959(96)02994-6

Savage, L. J. (1954). *The foundations of statistics* (2nd ed.). New York: Dover Publications.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge: MIT Press.

Taleb, N. N. (2007). *The black swan: The impact of the highly improbable*. New York: Random House.

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

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

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science, 20,* 473–479. doi:10.1111/j.1467-9280.2009.02319.x

Viscusi, W. K. (2002). *Smoke-filled rooms: A postmortem on the tobacco deal*. Chicago: University of Chicago Press.

Weber, E. U. (2012). Doing the right thing willingly: Using the insights of behavioral decision research for better environmental decisions. In E. Shafir (Ed.), *The behavioral foundations of public policy* (pp. 380–397). Princeton: Princeton University Press.

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

Whittle, P. (1988). Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability, 25,* 287–298. doi:10.2307/3214163

Wilford, J. N. (2006, March 7). Long before burying Pompeii, Vesuvius vented its wrath. *New York Times*. Retrieved from http://www.nytimes.com/2006/03/07/science/07volc.html

Yechiam, E., & Busemeyer, J. R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making, 19,* 1–16. doi:10.1002/bdm.509