

The Role of Information Sampling in Risky Choice

Ralph Hertwig, Greg Barron, Elke U. Weber,
and Ido Erev

Life is a gamble. True to the cliché, we can rarely be certain of the consequences of our everyday actions. Presumably because of life's lack of certainty, much of the psychological research on decision making under risk is devoted to the study of choices between monetary gambles or prospects. With few exceptions, the outcomes and outcome probabilities of such gambles are explicitly described to the decision maker before he or she makes a choice. How often are life's gambles described in this way? When we decide whether to back up our computer's hard drive, cross a busy street, or go out on a date, we often do not know the complete range of the possible outcomes, let alone their probabilities. Yet we routinely make such decisions, usually without difficulty. We do so by drawing upon our experience with the relevant prospects, for instance, by recalling the outcomes of previous choices. Though effective, sampling from experience brings with it the potential to generate a skewed picture of the risky prospects we face. As we will argue here, this potential is greatest when the outcome probabilities are small and relatively few outcomes have been experienced by the decision maker. Under these circumstances, decisions derived from experience are likely to be systematically different from those made in full knowledge of the outcomes and outcome probabilities.

In this chapter, we investigate risky decisions made on the basis of sampled experience. Three findings deserve emphasis. First, when outcome information is sampled, the likelihood of rare events is often underestimated. Second, when the sampling process is sequential, more recent outcomes tend to receive more weight than do earlier outcomes. Third, in decisions from experience, the option that appears best in light of the experienced sample tends to be the one selected. To account for these effects, we propose a learning mechanism that models how people update their estimates of a risky prospect's value in light of newly sampled outcomes.

INTRODUCTION

CALL ME ISHMAEL. Some years ago – never mind how long precisely – having little or no money in my purse, and nothing particular to interest me on shore, I thought I would sail about a little and see the watery part of the world. (Melville, 1851/1993, p. 1)

Thus begins *Moby Dick*, and with it the voyage of the Pequod and her captain, Ahab – a voyage of which Ishmael will be the sole survivor. With its myriad metaphors, symbols, and characters, Herman Melville's great novel has been interpreted as everything from a saga about Promethean determination and undying hatred to allegories about the Gold Rush, man's quest for knowledge, or man's will to master nature. Yet *Moby Dick* is also a fascinating factual chronicle of nineteenth-century whaling as revealed in the hardships, routines, and adventures of the men aboard the Pequod. As Melville (1851/1993, pp. 321–322) portrayed it, whale hunting was a high-stakes gamble:

But here be it premised, that owing to the unwearied activity with which of late they have been hunted over all four oceans, the Sperm Whales, instead of almost invariably sailing in small detached companies, as in former times, are now frequently met with in extensive herds, sometimes embracing so great a multitude, that it would almost seem as if numerous nations of them had sworn solemn league and covenant for mutual assistance and protection. To this aggregation of the Sperm Whale into such immense caravans, may be imputed the circumstance that even in the best cruising grounds, you may now sometimes sail for weeks and months together, without being greeted by a single spout; and then be suddenly saluted by what sometimes seems thousands on thousands.

The whalers aboard the Pequod were forced to make decisions about, for instance, the course of their ship under conditions that can be described in decision-theoretic terms as ignorance of the options' outcomes and probabilities. Although they hoped that sailing in known cruising grounds would increase their chance of being at the right place at the right time, they knew neither how large the herds they came across would be nor the probability of such encounters. Instead of guessing, however, the whalers most likely relied on their experience from previous hunting expeditions. Moreover, they were able to update their expectations in light of new samples of information. In this view, each day on which they cruised waters that the sperm whale was believed to frequent could be seen as a draw from a payoff distribution. After sequentially accumulating experience about a specific cruising ground, they would decide to stay or to abandon it.

Of course, decisions in which knowledge about the outcomes and outcome probabilities is scant are not unique to nineteenth-century whaling. In fact, ignorance or at least partial ignorance of outcomes and their probabilities may be the rule rather than the exception in many everyday decisions.

Like Melville's whalers, decision makers then can rely on past experience and continue to sample information about such prospects before making a decision. We have referred to this category of decisions as *decisions from experience* (Hertwig et al., 2004).

As the metaphor of life as a gamble implies, decisions from experience represent a bet on the future that is based on one's necessarily limited experience of the past. Indeed, the gambling metaphor has inspired a good deal of the psychological research on decision making. On the basis of a review of the early literature such as Savage's (1954) *The Foundations of Statistics*, Goldstein and Weber (1997, p. 569) concluded the following:

Gambling decisions were believed to be prototypical of virtually *all* decisions. Almost any contemplated action, for example, where to go to school, what job to take, or whom to marry, will have consequences that cannot be predicted with certainty but which vary in their likelihood and desirability. In this sense, then, real-life decisions have the same structure as gambles. Moreover, because this applies to virtually all decisions, *life is a gamble*.

Given this background, it is not surprising that monetary gambles became the fruit fly of decision research (Lopes, 1983). Just as biologists use the *Drosophila* as a model organism to study, for instance, genetic inheritance, so psychologists use the monetary gamble as a model task to study, for instance, decision making under risk. A monetary gamble is a well-defined alternative consisting of n outcomes ($n = 1, 2, 3, \dots, k$) and those outcomes' probabilities. In the typical notation, the gamble "\$32, .1; 0 otherwise" yields an outcome of \$32 with a 10% chance, and an outcome of 0 with a 90% chance. A dilemma is introduced by presenting people with two such gambles and asking them to say which one they prefer. For instance, they may be asked to choose between the following two gambles:

\$32, .1; 0 otherwise \$3 for sure.

If people maximize expected value, they will decide against the sure gamble because the two-outcome gamble offers the higher expected value, \$3.20.¹

Now compare this choice with the choices faced by the whalers aboard the Pequod or, to turn to a present-day example, with those faced by overweight people considering whether to go on a diet. Although the diet has the potential to help them reduce their weight, it may have undesirable side effects, the nature and probability of which are relatively unknown. Consider the Atkins diet. Its recommendation to shun carbohydrates and eat any amount of fat contradicts the predominant view in the field of nutrition

¹ The principle of choosing the option with the highest expected value (EV) is expressed as $EV = \sum p_i x_i$, where p_i and x_i are the probability and the amount of money, respectively, of each possible outcome ($i = 1, \dots, n$) of a gamble.

that a weight-loss diet should be low in fat and high in starch. Two recent studies published in the *Annals of Internal Medicine*, however, appear to validate the Atkins approach (Stern et al., 2004; Yancy et al., 2004). These studies found that the diet results not only in weight loss but also in cardiovascular health benefits and improved blood sugar, prompting some nutritionists to warn of the diet's unknown health effects in the long term (Big News, 2004). Thus, even when rigorous research results are available, life is indeed a gamble. Because risky prospects do not necessarily come with convenient descriptions of all possible outcomes and outcome probabilities, people must often rely on their experience to choose among them.

What do we know about the psychology underlying such experience-based decisions? Almost nothing. This is because researchers have focused almost exclusively on what we have called *decisions from description* (Hertwig et al., 2004). The goal of this chapter is to explore the important class of decisions from experience. As we demonstrate shortly, investigating decisions from experience does not require researchers to abandon their *Drosophila*. Choices among monetary gambles can be constructed such that people sample information from the payoff distributions – that is, garner experience of the outcomes and outcome probabilities – before they select among them.

DECISIONS FROM EXPERIENCE IN MONETARY GAMBLES

How much information do decision makers sample from payoff distributions before they feel ready to choose among them? Will their choices differ from those based on complete descriptions of the gambles? To address these questions, Hertwig et al. (2004) presented students at the Technion (Haifa, Israel) with the six risky choice problems displayed in Table 4.1. The two gambles in each problem differed from one another with regard to both expected value and variability. Of the six problems, four required choosing among positive prospects (gains), and two required choosing among negative prospects (losses).

Unlike in most studies of risky choice, respondents were not given the gambles' outcomes and probabilities. Instead, they were afforded the opportunity to sample information about the properties of the two gambles in each problem. Specifically, each respondent was shown two boxes on a computer screen and was told that each box represented a payoff distribution. Clicking on a given box triggered random sampling (with replacement) of an outcome from its distribution. Respondents could sample outcomes from the boxes as often as and in whatever order they desired. They were encouraged, however, to sample until they felt confident enough to decide from which of the two boxes they would prefer to draw given that their next draw would have real monetary consequences. After they had stopped sampling and indicated their choice, they proceeded to the

TABLE 4.1. Summary of the Choice Problems and Results

Choice Problem	Gambles ^a		Expected Value		Percentage Choosing H		Prediction for H Choices ^b		Difference Between Groups ^c
	H	L	H	L	Description Group	Experience Group	Rare Event	H Choices	
1	<u>4, .8</u>	3	3.2	3	36	88	0, .2	Higher	-52(z = 3.79; p = .000)
2	<u>4, .2</u>	3, .25	.8	.75	64	44	4, 2	Lower	+20(z = 1.42, p = .176)
3	-3	<u>-32, .1</u>	-3	-3.2	64	28	-32, .1	Lower	-36(z = 2.55, p = .005)
4	-3	<u>-4, .8</u>	-3	-3.2	28	56	0, .2	Higher	+28(z = 2.01, p = .022)
5	<u>32, .1</u>	3	3.2	3	48	20	32, .1	Lower	+28(z = 2.09, p = .018)
6	<u>32, .025</u>	3, .25	.8	.75	64	12	32, .025	Lower	-52(z = 3.79, p = .000)

^a Underlining indicates the gamble including the rare event. H, gamble with the higher expected value; L, gamble with the lower expected value.

^b The entries in this column identify the rare event and indicate whether the percentage of respondents choosing the H gamble was expected to be higher or lower in the experience group than in the description group, assuming underweighting of rare events in the experience group.

^c This column shows the percentage of H choices in the experience group minus the percentage of H choices in the description group, along with the z statistic testing whether the difference between the two sample proportions is significantly different from zero.

Source: "Decisions from experience and the effect of rare events in risky choice," by R. Hertwig, G. Barron, E.U. Weber and I. Erev, 2004, Psychological Science, 15, p. 536. Copyright 2004 by the American Psychological Society. Reprinted with permission.

next problem. After making all six choices, they played out their preferred gambles and received the associated payoffs. In both groups, participants received a \$4.50 show-up fee and 2¢ for each point won (e.g., the outcome 32 was worth 64¢).

Let us turn to Problem 1 in Table 4.1 to illustrate the procedure. This problem offers a choice between two payoff distributions: a sure-thing distribution in which every draw yields 3, and a risky-prospect distribution in which a draw results in 4 or 0, with the former outcome being four times as likely as the latter outcome (.8 versus .2). One respondent in Hertwig et al.'s (2004) study sampled seven times from the sure-thing distribution, each time encountering 3, and sampled eight times from the risky-prospect distribution, finding 0 once and 4 seven times. Which one did she choose? She selected the risky prospect, as did 88% of respondents.

How does this context compare with one in which people do not need to develop their own representations of the gambles? Hertwig et al.'s (2004) study addressed this question by including a second group of respondents. Whereas those in the aforementioned *experience* group ($n = 50$) learned about each gamble's outcomes and outcome probabilities by *sampling* information from the payoff distributions, respondents in the *description* group ($n = 50$) received complete information (using the notation introduced in the foregoing) about each gamble's outcomes and outcome probabilities. Table 4.1, which shows the choices by group for each of the six problems, illustrates the striking differences between decisions from descriptions and decisions from experience. In Problems 2, 3, and 6 as well as in Problem 5, respondents in the description group were more likely to select the gamble with the higher expected value, H , than were respondents in the experience group; in Problems 1 and 4, the pattern was reversed. Except in Problem 2, all differences in the proportion of H choices were statistically significant.

Hertwig et al. (2004) proposed that the crucial difference between decisions from description and decisions from experience resides in the psychological impact of rare events on each type of decision. According to the dominant psychological theory of decision making under risk, *prospect theory* (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), people choose between risky prospects as if small-probability events receive more weight than they deserve given their objective probabilities and as if large-probability events receive less weight than they deserve given their objective probabilities.² Prospect theory's weighting function, however, is

² Prospect theory's decision-weight function plots weights that range from 0 to 1 against objective probabilities: Points above the 45° diagonal signal overweighting; that is, they represent weights that exceed the outcomes' objective probability of occurrence. Points below the 45° diagonal signal the opposite pattern. Note that such decision weights are not assumed to reflect any explicit judgment of the subjective probability of outcomes. Instead, weights are inferred from choices and provide a measure of an outcome's impact on a decision.

based exclusively on studies in which respondents received explicit descriptions of the gambles' outcomes and probabilities. Consequently, Hertwig et al. proposed that in decisions from description people choose between risky gambles as if they *overweight* low-probability events relative to their objective probability of occurrence, consistent with prospect theory, whereas in decisions from experience they choose as if they *underweight* low-probability events.³

To see how this pattern of weighting of rare events is consistent with the observed choices, look again at Problem 1, in which the rare event (i.e., 0) occurs with a probability of .2. If in decisions from description the rare outcome has more than its due impact on the attractiveness of the option, then respondents will tend to prefer the sure outcome 3 over the risky prospect 4 with probability .8. Indeed, 64% of respondents in the description group chose the sure thing in Problem 1 (Table 4.1). If in decisions from experience the rare outcome has less than its due impact on the attractiveness of the option, then respondents will tend to prefer the risky prospect to the sure thing. Indeed, the large majority of respondents in the experience group selected the risky prospect in Problem 1 (Table 4.1).

More generally, Table 4.1 shows that, in each of the six problems, the pattern of *H* choices in the description and experience groups was consistent with that expected assuming that people choose as if they underweight rare events in decisions from experience and overweight rare events in decisions from description (see the Prediction for *H* choices column in Table 4.1). Comparison of Problems 1 and 2 highlights this difference in the weighting of rare events. Linear weighting of probabilities implies that a given person will have the same preference in Problem 1 as in Problem 2 (i.e., the person will prefer *H* or *L* in both problems, depending on his or her utility function). In contrast, overweighting of rare events ($p < .25$) implies more *H* choices in Problem 2 than in Problem 1, and underweighting of rare events implies the opposite preference. Kahneman and Tversky (1979) observed 65% *H* choices in Problem 2 and 20% *H* choices in Problem 1. They interpreted this result, a violation of expected utility theory known as the Allais paradox, as an instance of overweighting of rare events (Kahneman & Tversky, 1979). [For other violations that prospect theory explains in terms of overweighting of rare events, see Camerer (2000) and Tversky & Kahneman (1992).] As Table 4.1 shows, the majority choices in the description group replicate those observed by Kahneman and Tversky, whereas those in the experience group show the reverse pattern.

Although people in the experience and description groups responded to structurally identical choice problems, their choices were dramatically

³ Here we define a low-probability, rare event in a preliminary and somewhat arbitrary fashion as an event with a probability less than or equal to .2.

different. Respondents in the experience group sequentially sampled outcomes from the respective payoff distributions and chose between options as if they underweighted rare events. Respondents in the description group received complete outcome and probability information and chose between options as if they overweighted rare events. Hertwig et al.'s (2004) findings corroborate those observed in a small set of other studies that investigated decisions from experience (e.g., Barkan, Zohar, & Erev, 1998; Barron & Erev, 2003; Weber, Shafir, & Blais, 2004). The question is, why do people choose as if they underweight rare events in decisions from experience?

INFORMATION SEARCH IN DECISIONS FROM EXPERIENCE

Decisions from experience depend on what outcomes are sampled before the decision is made. Therefore, any account of the experience–description distinction ought to consider how much information people sample from the gambles' payoff distributions and how they combine the sampled information into a judgment of the gambles' attractiveness. We first address the issue of sample size and then investigate the impact of recency on the updating of sampled information.

Sample Size Matters

Figure 4.1 shows the median number of draws per problem in the experience group. Two observations are noteworthy. First, people sampled almost equally often from each payoff distribution in each problem. Second, the total number of draws per problem was relatively small, with a median of 15 draws.

Whatever the reasons for people's limited search effort,⁴ it has an obvious consequence: The smaller the number of draws, the larger the probability that a decision maker will never come across rare events, remaining ignorant of their existence. Table 4.2 presents the number of respondents who never encountered the rare event in the experience group. For illustration, consider Problem 5. Here the median respondent sampled seven cards from the payoff distribution that offered 32 with probability .1 (and 0 otherwise). As a consequence, 68% of respondents (17 out of 25) never encountered the good but rare outcome 32. Across the six problems, the rare event was not encountered in 44% of all sampling sequences.

⁴ One explanation involves short-term memory limits that provide a natural stopping rule for information acquisition (Kareev, 2000). In fact, more than half of respondents in Hertwig et al.'s (2004) experience group sampled exclusively from one payoff distribution before switching to the other one. The median number of draws from each option was around seven, a number often associated with the capacity of short-term memory.




FIGURE 4.1. Median number of draws in the experience group in each of the six choice problems (see Table 4.1). *H* and *L* represent the gambles with the higher and lower expected value, respectively. From "Decisions from experience and the effect of rare events in risky choice," by R. Hertwig, G. Barron, E. U. Weber and I. Erev, 2004, *Psychological Science*, 15, p. 537. Copyright 2004 by the American Psychological Society. Reprinted with permission.

Ignorance of the rare event appears to have a clear impact on the subsequent choice (see the Choosing Rare-Event Gamble columns in Table 4.2). When the rare event was "good" relative to other outcomes, as in Problems 2, 4, 5, and 6, not encountering it virtually assured that respondents did not select the gamble involving the rare event, whereas encountering it at least once raised the chance of choosing the rare-event gamble to about 50%. When the rare event was "bad" relative to other outcomes, as in Problems 1 and 3, the opposite was true. When respondents never encountered the rare event, they always selected the gamble involving the rare event. In contrast, encountering it at least once reduced the chance of choosing the rare-event gamble to about 66%.

As well as increasing the risk of never encountering the rare event, drawing a small sample makes it more probable that one will encounter the rare event less frequently than expected given its objective probability. This is because the binomial distribution for 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). For such distributions, one is more likely to encounter the rare event less frequently than expected (np) than more frequently than expected. For illustration, let us assume that 1,000 people sample from a distribution in which the critical event has a probability of .1 and estimate the event's probability to be the proportion in whatever sample they

TABLE 4.2. *Sample History and Choice*

Choice Problem	Gambles ^a		Rare Event Never Seen	Choosing Rare-Event Gamble	Rare Event Encountered (at Least Once)	Choosing Rare-Event Gamble
	<i>H</i>	<i>L</i>				
1	4, .8 <u>0, .2</u> <i>bad</i>	3	11	11	14	11
2	<u>4, .2</u> <i>good</i> 0, .8	3, .25 0, .75	7	0	18	11
3	-3	<u>-32, .1</u> <i>bad</i> 0, .9	10	10	15	8
4	-3	-4, .8 <u>0, .2</u> <i>good</i>	2	0	23	11
5	<u>32, .1</u> <i>good</i> 0, .9	3	17	1	8	4
6	32, .025 <i>good</i> 0, .985	3, .25 0, .75	19	1	6	2

^a Underlining indicates the rare event. Each rare event is indicated as either good or bad relative to other outcomes. *H*, gamble with the higher expected value; *L*, gamble with the lower expected value.

observe. Each person samples 20 times. Of the 1,000 people, 285 will observe the critical event twice and thus are likely to estimate its probability accurately. Another 392 will never observe the critical event or will observe it only once and thus will probably underestimate p . The remaining people (323) will encounter the crucial event 3, 4, 5, . . . , or 20 times and thus are likely to overestimate its probability. Note that, averaged across the 1,000 people, the estimated probability of the rare event will equal its probability in the population (i.e., .1) because the sample proportion is an unbiased estimator of the proportion in the population. However, for small samples, more people will encounter the rare event less frequently than expected rather than more frequently than expected.

Hertwig et al. (2004) averaged the number of times that respondents encountered the rare event across all six problems in the experience group and found that, consistent with the analysis of the binomial distribution, 78% of respondents encountered the rare event less frequently than expected (i.e., fewer than np times) whereas 22% of respondents encountered the rare event as frequently or more frequently than expected. In addition, the experienced frequency of the critical event had a clear impact on choices. When the rare event was "good" (e.g., 32 in Problem 5), the option involving it was selected in only 23% of cases in which it was encountered less frequently than expected. The same option was selected in 58% of cases in which it was encountered as frequently or more frequently than expected. Similarly, when the rare event was "bad" (e.g., 0 in Problem 1), the option involving it was selected in 92% of cases in which it was encountered less frequently than expected but in only 50% of cases in which it was encountered as frequently or more frequently than expected.

These results show that, in the small samples drawn by respondents in Hertwig et al.'s (2004) study, the experienced relative frequencies of rare events were on average smaller than the rare events' objective probabilities. It is therefore plausible that, even if people correctly recalled all the experienced outcomes and computed the probabilities based on the outcomes' frequencies in the sample, they tended to *underestimate* the probability of rare events. In this context, underestimation does not necessarily imply psychological distortion but rather a kind of information sampling (i.e., drawing small samples) that gives rise to a distorted picture of the world.

Given that many of the respondents received input that would give them a skewed picture of the objective probabilities, how good were their choices? To address this question, we computed the expected value of each gamble for each respondent *on the basis of the outcomes that he or she observed*. Viewed in light of the samples they drew, respondents generally acted like expected value maximizers: Across the six problems, respondents chose the gamble that according to the sample promised the higher expected value

in 74% of cases. Only 44% of choices, however, maximized the expected value as calculated on the basis of the objective probabilities.

Sampling Order Matters

Decisions from experience require people to update their impression of the desirability of a gamble by combining newly sampled outcomes with previously sampled outcomes. If all observations received an equal weight – that is, a weight of $1/n$, where n is the number of observations sampled from a payoff distribution – then the order in which observations are sampled would not matter. Both the expected value calculus and the expected utility calculus are impervious to order effects. Research on memory (e.g., Atkinson & Shiffrin, 1968) and belief-updating (e.g., Hogarth & Einhorn, 1992), however, show that the order in which evidence is presented to people matters even in tasks where it should be irrelevant (e.g., in free recall of the content of lists or in overall impressions of the value of choice options when outcomes are randomly sampled).

One order phenomenon is the *recency* effect, according to which observations made late in a sequence of observations receive more weight than they deserve (i.e., more than $1/n$). Such an effect would result in choices consistent with underweighting of rare events. To see why, let us assume the most extreme form of recency, in which the attractiveness of a gamble is determined solely by the most recently sampled outcome. Let us further assume that each of 100 people sample from two distributions: distribution *A*, in which the rare event, 32, has a probability of .1 and the common event, 0, has a probability of .9, and distribution *B*, in which 3 occurs for certain. After any number of draws, only 10 of the 100 people, on average, will have encountered 32 in their most recent draw from distribution *A*; the rest will have encountered 0 in that draw. Moreover, all of them will have encountered 3 in the most recent draw from distribution *B*. If only the most recent draw per distribution matters to a gamble's attractiveness, it is likely that the majority of participants will prefer distribution *B* over *A*, resulting in a modal choice that is consistent with underweighting of rare events. In other words, in decisions resulting from continuous updating of outcome probabilities, the recency effect can produce choices that make people look as if they underweight rare events.

To examine whether recency affected decisions in the experience group, Hertwig et al. (2004) split, for each choice problem and each respondent, the outcomes sampled from each option according to whether they fell in the first half or the second half of the sample sequence. After computing the options' average payoffs according to the information observed in the first half and the second half of the sequence, respectively, the authors then predicted each person's choice in each problem on the basis of these average

payoffs. The predictive power of the payoffs computed from the second half of the sequence clearly outperformed that of those computed from the first half. Whereas the first half predicted, on average, 59% of the choices, the second half predicted 75% of the choices; the difference between the two proportions was statistically significant ($t(49) = -3.1, p = .003$; two-tailed).

Hertwig et al.'s findings suggest the existence of two phenomena. First, people may underestimate the objective probabilities of rare events in small samples, and, second, they may overweight the impact of recent events. As demonstrated, both phenomena can contribute independently to choices consistent with underweighting of rare events. Let us clarify, however, that we do not assume that people explicitly estimate outcome probabilities or explicitly combine outcomes and probability estimates in a fashion suggested by expected value theory, expected utility theory or prospect theory. Instead, we propose that people go through a cognitively far less taxing process of evaluating and updating the value of a gamble, resulting in choices that appear to reflect underweighting of rare events. In what follows, we propose a model of such a process.

INFORMATION INTEGRATION IN DECISIONS FROM EXPERIENCE: THE VALUE-UPDATING MODEL

How do people integrate sequentially sampled outcomes? Many learning studies point to a weighted adjustment process in which a previous impression of the value of an option is combined with a newly sampled outcome (see, e.g., Bush & Mosteller, 1955; March, 1996). We propose a specific variant of a weighted adjustment mechanism and apply it to value updating in decisions from experience. In the *value-updating model*, a learner is assumed to update his or her estimate of the value of a gamble after each new draw from the payoff distribution by computing a weighted average of the previously estimated value and the value of the most recently experienced outcome. The model has three building blocks:

Sampling rule: During sampling, each trial consists of a random draw from one of the two payoff distributions. The person stops searching after a total of s draws (with each payoff distribution being sampled $s/2$ times).

Choice rule: When making the choice, the person selects the gamble with the highest value as determined on the basis of the sampled outcomes.

Updating rule: In algebraic terms, the updating mechanism can be written $A_j(t) = (1 - \omega_t)A_j(t - 1) + (\omega_t)v(x_t)$, where $A_j(t)$ represents the value of gamble j after t samples are drawn. This value is the weighted average of the value of the gamble based on the $t - 1$ previously




FIGURE 4.2. The proportion of people who selected *H*, the gamble with the higher expected value, in each choice problem and the corresponding proportion predicted by the value-adjustment model.

drawn outcomes, $A_j(t - 1)$, and the value of the outcome obtained in the most recently drawn outcome, x_t . The weight accorded to the new outcome is $\omega_t = (1/t)^\varphi$, where φ is a recency parameter. When $\varphi = 1$, all outcomes are equally weighted. When $\varphi < 1$, more recent outcomes receive more weight, implying a recency effect. When $\varphi > 1$, earlier outcomes receive more weight, implying a *primacy* effect. Finally, $v(\cdot)$ is prospect theory's value function, which is assumed to be concave for gains and convex for losses (Kahneman & Tversky, 1979):

$$v(x_i) = \begin{cases} x_i^\alpha & \text{if } x_i \geq 0, \\ \lambda |x_i|^\alpha & \text{if } x_i < 0. \end{cases}$$

Although the value-updating model has four parameters, only the recency parameter, φ , needs to be estimated from the data. The parameters α from the gain value function and λ from the loss value function are assumed to take the values estimated by Tversky and Kahneman (1992; $\alpha = .88; \lambda = 2.25$), and the number of draws, s , is taken to be the empirically observed mean sample size for each choice problem (see Figure 4.1).

We fit the model to the choices observed in Hertwig et al.'s study (2004). Figure 4.2 shows the actual and predicted (best mean-squared fit) choice proportions. The mean-squared deviation score between the proportions was .012, and the correlation between the actual and the predicted proportions was .91. With the exception of Problem 1, the predictions closely tracked the observations. Moreover, the φ parameter was estimated to be 0.29. As expected, this value implies a recency effect.

In conclusion, the value-updating model suggests one way in which sampled outcomes are integrated into an impression of the value of a

gamble in the context of decisions from experience. In future studies of decisions from experience, this mechanism is flexible enough to take different degrees of sampling (i.e., sample sizes) into account (through s), and it allows for quantification of the recency effect.

CONCLUSIONS

In a review of all published studies examining risky choices made by human respondents between a sure option and a two-outcome gamble of equal expected value, Weber et al. (2004) assessed the prevalence of decisions from description in the laboratory. Every single choice involved a decision from description. There is no doubt that important situations outside the laboratory call for decisions from description. Consider, for example, parents who must decide whether their child should receive the diphtheria, tetanus, and pertussis (DTaP) vaccine. By researching the DTaP vaccine at the National Immunization Program Web site, for instance, parents will learn that the possible side effects are high fever (a temperature of more than 105°F) and seizures and that up to 1 child in 1,000 suffers from the former and about 1 child in 14,000 suffers from the latter as a result of immunization.

Yet circumstances often deny us access to thorough descriptions of options' outcomes and their probabilities. One way to fill the void is to sequentially sample information and update one's mental representation of the situation. According to an article in *The New Yorker* (Dubner, 2004), this is how a thief named Blane Nordahl managed to steal millions of dollars worth of silver from scores of stately homes on the east coast of the United States. Burglary is fraught with risks, most of which cannot be estimated by consulting the Internet. Arguably the most important risk for a burglar to assess is that of getting caught. Nordahl learned from experience that it was dangerous to have partners in crime because there was a high probability that, if they got caught, they would turn him in. So he worked alone. Experience also taught him to steal silver – which, unlike jewelry, for example, tends to be stored on the ground floor of a home – and to work during the night rather than the day. Of course, the samples of experience on which risky choices are based can be vicarious as well as first-hand. Quoted in the same article in *The New Yorker*, Malcolm X's recollections of his early career as a burglar provide an example of decisions from vicarious experience:

I learned from some of the pros, and from my experience, how important it was to be careful and plan. . . . Burglary, properly executed, though it had its dangers, offered the maximum chances of success with the minimum risk. . . . And if through some slip-up you were caught, later, by the police, there was never a positive eyewitness (Dubner, 2004, p. 76).

In other contexts (e.g., financial investment), decision makers have access both to comprehensive statistics and to personal experience. Expert and novice decision makers often differ in the degree to which they rely on one type of information versus the other. For instance, in the domain of insurance, experts (i.e., actuaries) rely almost exclusively on summary statistics when pricing policies, whereas novices (i.e., insurance buyers) often use personal experience with the hazard to decide whether a price is fair. This can make people reluctant to buy actuarially fair or advantageous insurance policies for rare hazards like floods (Kunreuther et al., 1978). The difference between decisions from experience and decisions from description is one of a number of possible contributors to expert–novice differences in judgment and choice (for others, see Slovic & Weber, 2002).

The experience–description distinction can also shed light on some striking similarities between the behavior of humans and other animals (see also Weber et al., 2004). Because animals do not share the human ability to process symbolic representations of risky prospects, all their decisions are decisions from experience. In a study of the foraging decisions made by bees, Real (1991) observed that “bumblebees underperceive rare events and overperceive common events” (p. 985). To explain why bees’ “probability bias” diverges from that postulated in humans (in prospect theory’s weighting function), Real cited, among other factors, the fact that bees’ samples from payoff distributions (i.e., foraging patches) are truncated by memory constraints. When humans are placed in situations where they must make decisions based on experience and their experience is limited (by memory constraints or by a small number of encounters), they, like bumblebees, behave as if they underestimate the probability of rare events.

A Mere Mention Lends Weight

Why people choose as if they overweight rare events in decisions from description remains an open question in the literature (for a discussion, see Gonzalez & Wu, 1999), and we have no counterpart to the value-updating model to offer as an explanation. We suspect, however, that the overweighting of rare events is one of several phenomena in which merely presenting a proposition appears to increase its subjective truth and psychological weight. For instance, merely asking people whether a particular person is a crook makes them more likely to believe that such a person is one (e.g., Wegner et al., 1981); merely considering a proposition enhances its subjective truth (Fiedler, 2000; Fiedler et al., 1996), and merely imagining the occurrence of positive events (e.g., winning a free trip) or negative events (e.g., being arrested) increases the subjective likelihood of those events (Gregory, Cialdini, & Carpenter, 1982). We hypothesize that merely presenting a rare event may increase its psychological weight in a similar

way. The propositional representations of gambles in decisions from description – for instance, “32 with probability .1; 0 otherwise” – put more equal emphasis on the two possible outcomes than the discrepancy between their actual probabilities of occurrence warrants. If attention translates into decision weight as some research suggests (Weber & Kirsner, 1996), then, other things being equal, the psychological weights of rare and common events will indeed be closer to one another than they should be. Thus, the effect of mere presentation might explain underweighting of common events as well as overweighting of rare events in decisions from description.

The mere presentation account raises interesting questions. For example, would overweighting of rare events in decisions from description decrease or even vanish if the gambles were described in an analogical way? The structure of analogical representations reflects more the structure of what is represented than do propositional representations. For instance, the option “32 with probability .1; 0 otherwise” can be expressed as follows: “a random draw from the population {0, 0, 0, 0, 0, 32, 0, 0, 0, 0}.” In this representation, information regarding the relative frequency of the option’s outcomes can be read off directly. Moreover, to the extent that more attention needs to be allocated to the processing of the frequent event than the rare event, the resulting decision weights may more accurately reflect the objective probabilities.

Small Samples Show Less Variability

Reliance on small samples not only plays a role in decisions from experience but also contributes to the fact that people perceive the world as less variable than it actually is (Kareev, Arnon & Horwitz-Zeliger, 2002). In fact, underestimating the variance of populations is equivalent to underweighting rare events. For instance, in Problem 6, about two-thirds of respondents encountered only the outcome 0 when they drew from the payoff distribution involving the rare event (Table 4.2). If they had used sample variability or lack thereof as an estimate of population variability without correcting for sample size, then they would have underestimated the true variability in that option’s payoff distribution. Consistent with this conclusion, Kareev et al. showed that people tend to perceive sample variability as smaller than it is and that reliance on small samples (the size of which is often related to working-memory capacity) causes the underestimation.

Kareev et al. (2002) suggested that time constraints, limits in memory capacity, and lack of available data are key conditions under which people perceive variability as smaller than it is (see also Kareev et al. in this volume). These may also be the conditions under which decisions from experience and decisions from description diverge. We think that even if

decision makers derive their decisions from large samples of outcomes, however, the psychological impact of a rare event may still deviate from its objective probability because recency could still amplify the impact of recently sampled outcomes.

Primacy and Recency Join Forces

In our analysis based on the value-updating model, we found evidence only for recency. However, whether it is later or earlier observations in a sequential updating process that get more weight is a subject of dispute. Some researchers have found evidence of primacy (Nisbett & Ross, 1980, p. 172), others of recency (Davis, 1984), and still others of both effects (Anderson, 1981). In an effort to make sense out of this heterogeneous set of findings, Hogarth and Einhorn (1992) proposed a belief-adjustment model that predicts different types of order effects as a function of the nature of the task (e.g., judging the value of an option after each outcome is drawn or only after sampling is complete), the complexity of the outcome information [e.g., a single value, as in Hertwig et al.'s (2004) study, or a page of text describing a person], and the length of sampling. The key point here is that both primacy and recency effects will cause rare events to receive less weight than they deserve given their objective probability of occurrence. Because of the very rarity of rare events, most people are less likely to encounter rare events than frequent events both at the beginning and at the end of a sequence.

EPILOGUE

Let us return to the decks of the Pequod. On the second day of their three-day chase of Moby Dick, the chief mate, Starbuck, implored Ahab:

Never, never wilt thou capture him, old man – In Jesus' name no more of this, that's worse than devil's madness. . . . Shall we keep chasing this murderous fish till he swamps the last man? Shall we be dragged by him to the bottom of the sea? Shall we be towed by him to the infernal world? Oh, oh, – Impiety and blasphemy to hunt more! (Melville, 1851/1993, p. 467)

Ahab ignored the plea. In Melville's portrayal, to hunt down and kill Moby Dick was nothing less than Ahab's destiny. He had no choice but to continue the quest, destroying himself, his ship, and his crew in the process. Seen from the perspective of decisions from experience, one may speculate that Ahab sampled the hunt and the kill of "this murderous fish" in his imagination innumerable times. And, in our imaginations at least, every one of us is master of the odds.

References

- Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.
- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (Vol 2., 89–195). New York: Academic Press.
- Barkan, R., Zohar, D., & Erev, I. (1998). Accidents and decision making under uncertainty: A comparison of four models. *Organizational Behavior and Human Decision Processes*, *74*, 118–144.
- 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.
- Big News. (2004, May 22). *The Economist*, *371*, 87.
- Bush, R., & Mosteller, F. (1955). *Stochastic models for learning*. New York: Wiley.
- Camerer, C. F. (2000). Prospect theory in the wild: Evidence from the field. In D. Kahneman & A. Tversky (Eds.), *Choices, values, and frames* (pp. 288–300). Cambridge, UK: Cambridge University Press.
- Davis, J. H. (1984). Order in the courtroom. In D. J. Miller, D. G. Blackman, & A. J. Chapman (Eds.), *Perspectives in psychology and law*. New York: Wiley.
- Dubner, S. J. (2004, May 17). The silver thief. *The New Yorker*, 74–85.
- Fiedler, K. (2000). On mere considering: The subjective experience of truth. In H. Bless & J. P. Forgas (Eds.), *The message within: The role of subjective experience in social cognition and behavior* (pp. 13–36).
- Fiedler, K., Armbruster, T., Nickel, S., Walther, E., & Asbeck, J. (1996). Constructive biases in social judgment: Experiments on the self-verification of question contents. *Journal of Personality and Social Psychology*, *71*, 861–873.
- Goldstein, W. M., & Weber, E. U. (1997). Content and discontent: Indications and implications of domain specificity in preferential decision making. In W. M. Goldstein & R. M. Hogarth (Eds.), *Research on judgment and decision making: Currents, connections, and controversies* (pp. 566–617). New York: Cambridge University Press.
- Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology*, *38*, 129–166.
- Gregory, W. L., Cialdini, R. B., & Carpenter, K. M. (1982). Self-relevant scenarios as mediators of likelihood estimates and compliance: Does imagining make it so? *Journal of Personality and Social Psychology*, *43*, 89–99.
- 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.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, *24*, 1–55.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.
- Kareev, Y. (2000). Seven (indeed, plus or minus two) and the detection of correlations. *Psychological Review*, *107*, 397–402.

- Kareev, Y., Arnon, S., & Horwitz-Zeliger, R. (2002). On the misperception of variability. *Journal of Experimental Psychology: General*, *131*, 287–297.
- Kunreuther, H., Ginsberg, R., Miller, L., Sagi, P., Slovic, P., Borkin, B., & Katz, N. (1978). *Disaster insurance protection: Public policy lessons*. New York: Wiley.
- Lopes, L. L. (1983). Some thoughts on the psychological concept of risk. *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 137–144.
- March, J. G. (1996). Learning to be risk averse. *Psychological Review*, *103*, 309–319.
- Melville, H. (1993). *Moby Dick*. New York: Barnes & Noble Books. (Original work published 1851.)
- Nisbett, R., & Ross, L. (1980). *Human inference: Strategies and shortcomings of human judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- Real, L. A. (1991). Animal choice behavior and the evolution of cognitive architecture. *Science*, *253*, 980–986.
- Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.
- Slovic, P., & Weber, E. U. (2002). *Perception of risk posed by extreme events*. White paper for conference "Risk Management Strategies in an Uncertain World," Palisades, NY, April 12, 2002.
- Stern, L., Iqbal, N., Seshadri, P., Chicano, K. L., Daily, D. A., McGrory, J., Williams, M., Gracely, E. J., Samaha, F. F. (2004). The effects of low-carbohydrate versus conventional weight loss diets in severely obese adults: One-year follow-up of a randomized trial. *Annals of Internal Medicine*, *140*, 778–785.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.
- Weber, E. U., & Kirsner, B. (1996). Reasons for rank-dependent utility evaluation. *Journal of Risk and Uncertainty*, *14*, 41–61.
- 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.
- Wegner, D. M., Wenzlaff, R., Kerker, R. M., & Beattie, A. E. (1981). Incrimination through innuendo: Can media questions become public answers? *Journal of Personality and Social Psychology*, *40*, 822–832.
- Yancy, W. S., Jr., Olsen, M. K., Guyton, J. R., Bakst, R. P., Westman, E. C. (2004). A low-carbohydrate, ketogenic diet versus a low-fat diet to treat obesity and hyperlipidemia. *Annals of Internal Medicine*, *140*, 769–777.