

The description–experience gap in risky choice

Ralph Hertwig¹ and Ido Erev²

¹ University of Basel, Department of Psychology, Missionsstrasse 60/62, CH-4055 Basel, Switzerland

² Technion, Max Wertheimer Minerva Center for Cognitive Studies, Faculty of Industrial Engineering and Management, Haifa 32000, Israel

According to a common conception in behavioral decision research, two cognitive processes—overestimation and overweighting—operate to increase the impact of rare events on people’s choices. Supportive findings stem primarily from investigations in which people learn about options via descriptions thereof. Recently, a number of researchers have begun to investigate risky choice in settings in which people learn about options by experiential sampling over time. This article reviews work across three experiential paradigms. Converging findings show that when people make decisions based on experience, rare events tend to have less impact than they deserve according to their objective probabilities. Striking similarities in human and animal experience-based choices, ways of modeling these choices, and their implications for risk and precautionary behavior are discussed.

Beware of black swans

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 bust 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 unlikely events? 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’—unexpected and unpredictable rare events that carry an enormous impact [1]. Of course, not only modern risk management paradigms failed 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 food poisoning or of lung cancer resulting from smoking [2,3]. Moreover, people are depicted as remembering past experiences by how they felt at the peak (rare moment) of their experience and how they felt when the experience came to an end [4]. 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 [5]. In fact, cumulative prospect theory explains the puzzling co-occurrence of two behaviors—that the same people who purchase lottery tickets that promise tiny chances of winning also take out insurance against tiny chances of damage [6]—by 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 private citizens alike, behave as if they were not cognizant of the rare events that triggered what some called a bona-fide depression [7]? 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 [7] to humans’ ‘animal spirits’ [8]. However, there is another possibly enabling condition. The customary portrayal of people as being oversensitive to rare events neglects that when people recruit their experience sampled across time to make risky decisions, chances are that rare events (such as the burst of housing bubbles) receive less impact than they deserve according to their objective probabilities.

The description–experience gap

Just as biologists use the *Drosophila* (fruit fly) 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 [9]. 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 1 illustrates a typical description-based decision problem.

In everyday life, however, people can rarely peruse such descriptions of probability distributions (there are a few exceptions such as media weather forecasts stating probabilities of precipitation [10]). When people decide whether to take out a mortgage loan or contemplate the success of a first date there are no tabulated risks to consult. Instead, they need to rely on their experiences—if existent—with these options, thus making *decisions from experience* rather than *decisions from description* [11]. Both kinds of decisions can be understood as opposite entries on a continuum of uncertainty about what one is choosing between. Using Knight’s terminology [12], decisions from descriptions involve *a priori probabilities*, whereas decisions from experience involve *statistical probabilities*, which must be assessed

Corresponding author: Hertwig, R. (Ralph.Hertwig@unibas.ch)

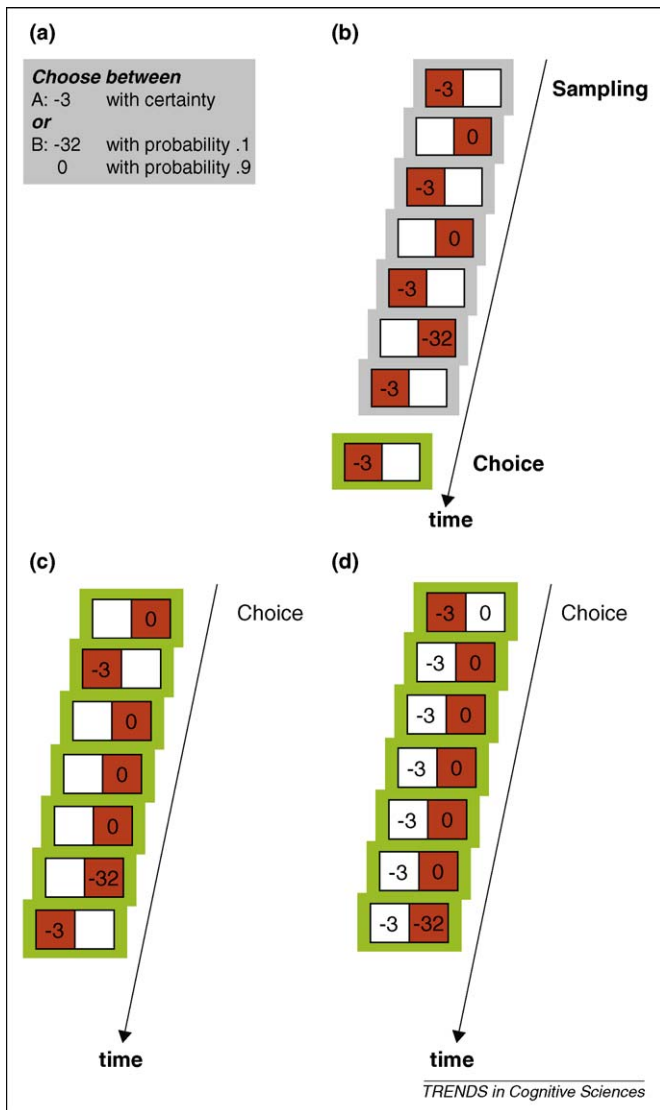


Figure 1. How to study decisions from description and experience? (a) The choice task in decisions from description typically consists of two monetary gambles with explicitly stated outcomes and their probabilities. In decisions from experience, three paradigms have been employed: (b) The *sampling paradigm* 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. (c) The *partial-feedback paradigm* collapses 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). (d) The *full-feedback paradigm* 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).

‘if at all, by tabulating the results of experience’ (p. 215), and therefore they invariably fall short of the standards of accuracy set by a priori probabilities [13].

In the 1950s and early 1960s, prior to the advent of modern behavioral decision research, decision scientists investigated decisions from experience. They examined, for example whether and how people learn the probability structure over outcomes through trial-by-trial feedback (for a review see [14]). Perhaps, because of the designs’ impracticality—purportedly hundreds of trials are needed before behavior stabilizes—modern behavioral decision

researchers turned away from the transients of learning (for an exception see, for example [15]). Moreover, with the increasing importance of expected utility theory, the study of anomalies became pertinent; those required the conveying of perfect information about the probabilities of relevant events (Figure 1). Interest in issues of learning and experience-based decisions, however, remained alive in other fields such as operation research (see literature on multi-armed bandit problems [16]).

Modern decision scientists’ interest in decisions from experience has been rekindled by the recent observation of systematic and robust differences between decisions based on experience and 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 (Figure 1). Clicking a button results in a random draw from those distributions. Three variations of this experimental tool have been employed. In the *sampling paradigm*, people first sample as many outcomes as they wish and only then decide from which distribution to make a single draw for real [11,17]. In the *full-feedback paradigm*, each draw 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) [18]. The *partial-feedback paradigm* is identical to the full-feedback paradigm, except that people only learn about the obtained payoffs [19,20]. Here, unlike in the first two paradigms, respondents face an exploitation–exploration tradeoff. Exploitation and exploration represent two goals associated with every choice, namely, to obtain a desired outcome (exploitation) or to gather new information about other, perhaps better, actions (exploration) [21].

Across all three experiential paradigms, a robust and systematic *description–experience gap* has emerged in numerous studies. Figure 2 illustrates this gap in six decision problems [22]. 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 (probability of 0.1 or less). In all three experiential paradigms, respondents tend to select the risky option when the desirable outcome occurs with high probability, and select the safe option when the desirable outcome occurs with low probability. This tendency is reversed in decisions from description, and 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).

What causes the description–experience gap?

Several causes have been proposed as contributing to the description–experience gap. These are reviewed in the following section.

Small samples

Based on people’s search in the sampling paradigm, reliance on small samples has been proposed as one factor




Figure 2. The description–experience gap. Proportion of choices of the risky option (i.e. risky choice) as a function of the probability of the more desirable outcome in 6 of 120 problems studied in [22]. Each presents a choice between a risky option and a safe option. The decision problems and the expected values (EV) of the risky options are displayed below. Each problem was studied using the four paradigms listed in Figure 1 [22; the data from the full-feedback are reported in Nevo, I, and Erev, I, 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.

that contributes to the attenuated impact of rare events [11]. Across numerous studies, respondents typically proved restrained in their information search, with a median number of samples per choice problem typically ranging between 11 and 19 across studies (reviewed in [13]). For such samples, chances are that a person does not even experience the rare event. More generally, one is more likely to undersample than to oversample the rare event. 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). Reliance on small samples has also been discussed as a potential explanation for bumblebees' underweighting of rare events (Box 1).

This explanation in terms of small samples has prompted a critical response [26] 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 might 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* because of 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 (why people appear to rely on small samples has received scant attention but see Ref. [15,27]). If sampling error were 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) could reduce but did not eliminate the gap [13,28]. If sampling error caused the gap, then removing the error by aligning experienced probabilities in the sample to the objective probabilities should eliminate it; it did not [29]. If sampling error were the sole root of the gap, then presenting respondents in the description condition with the exact same information that others experienced ('yoking') should eliminate the gap. In one study it did [30]; in another it did for small but not for large samples [13; see these authors' discussion of 'trivial choices' as one possible explanation for the mixed results obtained].

Box 1. The common characteristics of human and other animals' decisions from experience

Humans and animals alike face uncertain environments. By studying decisions from experience, researchers enjoy the surplus benefit of being able to better understand human choice by discovering how it resembles or differs from choices made by other animals. There is some indication that an attenuated impact of rare events characterizes both humans' and animals' decisions from experience. Investigating the foraging behavior of bees across floral rewards distributions, Real concluded that 'bumblebees underperceive rare events and overperceive common events' (p. 985) [23]. Moreover, he observed that 'bees frame their decisions on the basis of only a few visits' (p. S133), and suggested that such reliance on small samples might be adaptive when there is a high degree of spatial autocorrelation in the distributions of floral rewards [24]. Placing people into situations in which they render decisions from experience also led to the discovery that risk preference in humans, birds and insects is better predicted in terms of the coefficient of variation (CV), a measure of risk per unit of return, than outcome variance, the measure typically used in normative models such as expected utility theory [17].

A striking difference between humans' and other animals' decisions from experience could be resolved by taking properties of reward into account [25]. In studies of animal choice, 'participants' such as rats receive rewards in the form of cups of water and thus discern differences in rewards less accurately than humans who receive the rewards in the form of digital numbers. When the perceptual noise in the presentation of rewards is increased, however, humans exhibit the same pattern of choices as animals. Similarly, when perceptual noise is decreased and discrimination becomes easier, honeybees display the same choices as humans. In general, in experience-based choices humans and honeybees alike appear to deviate from maximization in favor of the alternative perceived to lead to better outcomes most of the time (thereby underweighting rare events).

Finally, the gap persisted even when people were presented both descriptions and experience, relative to just descriptions [31].

At this point, the reality of the description–experience gap across the three experiential paradigms is unchallenged—its cause, however, is debated. Some researchers have argued that the gap in the sampling paradigm is statistical in nature [26,30,32]; others have proposed that the sampling error is not the sole cause [11,13,28,29]. 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, in which rare events' impact is attenuated even after a hundred of trials with perfect feedback. Beyond sampling error, what psychological factors might be in play?

Recency

A psychological factor proposed to contribute to the description–experience gap is *recency* [11]. Ubiquitously observed in memory, belief updating and judgments [33], 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 outcomes that occurred recently in the flow of their experience relative to previous outcomes

[11], no or little impact of recency was observed in later studies [13,29,30].

Estimation error

In theory, the description–experience gap could also be the consequence of a systematic estimation error [26], 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 [2,34]. 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 [28,29]. That is, people do not systematically estimate rare things to be even rarer.

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 ([35,36] for related notions). Such *contingent sampling* is likely to be ubiquitous 'in the wild' [37]. 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. restless bandit problem [38]), reliance on similar experiences is an efficient heuristic [39]. Admittedly, one weakness of the notion of contingent sampling is the fuzzy term 'similarity.' However, defining similarity within a formal model is not impossible, as the result of a recent model competition demonstrated (Box 2).

Information format and cognitive algorithms

Finally, the description–experience gap might be partly caused by different formats of statistical information triggering different cognitive algorithms [13,28]. Most investigations of decisions from description represent the likelihood of outcomes in terms of single-event probabilities (or percentages). In experiential paradigms, in contrast, people come across sequences of events, without an explicit synopsis of the events' probabilities. This difference echoes the distinction between single-event probabilities and natural frequencies in Bayesian reasoning, with the different format giving rise to different cognitive algorithms and degrees of Bayesian reasoning [44,45]. By analogy, the sequential experience of events might trigger different cognitive algorithms than those triggered by described single-event probabilities.

In sum, modern behavioral decision research was 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. Where rare events are implicated, description-based and experience-based decisions can drastically diverge. Next to sampling error, a number of psychological factors have been proposed. One important question for the future is to find which of these

Box 2. How to account for decisions from experience: model competitions

Decisions from description can differ systematically from decisions from experience. Experience-based decisions, for example, appear to reverse the fourfold pattern of risk attitudes that has been found in decisions from description [40], provoking a call for two different theories of risky choice [11,17,41]. Others, however, have responded that cumulative prospect theory, initially designed for decisions from description, could also account for decisions from experience [26]. To evaluate how well a model predicts people's choices, its performance needs to be compared to those of other models. Several model competitions involving decisions from experience have been conducted ([28,42]; the most extensive one was the choice prediction competition held at the Technion [22]).

The Technion competition consisted of three contests involving the description, the sampling and the partial-feedback paradigms (Figure 1). The objective in all three contests was to predict the proportion of people who prefer a safe option with a medium payoff to a risky option yielding two payoffs, one high, the other one low. Each contest was based on two sets of choice data, a fitting and a competition set. In order to generate both datasets, the same method and participant pool was used, and each dataset consisted of 60 randomly selected problems with the aforementioned properties. The fitting set, together with the fit of several baseline models, was posted on a website, and enabled researchers to fit the model that they thought would perform best in predicting the competition set. Fourteen teams responded to the challenge, representing a wide range of theoretical approaches (e.g. heuristics, algebraic models, logit-regression model and a model based on the ACT-framework).

Again, the empirical data revealed a substantial description–experience gap, thus replicating previous results. Notwithstanding the limited psychological plausibility of some of the leading models, the description–experience gap is also reflected in the models that performed well in the three contests. In decisions from description, the best-performing model was a stochastic version of cumulative prospect theory with decision weight estimates reflecting overweighting of rare events. All leading models in the two experience contests share the assumption that prior to choice the decision maker recalls a small set of about five experiences with each option, and tends to select that option resulting in better outcomes in this sample. In addition, the winning model in the partial feedback paradigm implemented the notion of contingent sampling (see text). Note that this model competition is featured in a special edition of the *Journal of Behavioral Decision Making* devoted to the topic of 'decisions from experience' [43].

factors operate across all three experiential paradigms (see Box 3).

Three more issues merit brief discussion. First, underweighting of rare events as measured in terms of the parameters of the decision-weighting function of cumulative prospect theory [5] is not a *necessary* condition for the description–experience gap. A gap can emerge if the parameters fitted to experience-based choices reflect linear weighting or underweighting for rare events [28,29]. Second, experience and description can be combined in myriad ways, resulting in a rich set of results [46,47]. For example, re-entering experience into decisions from description—people responded repeatedly to the same choice problem and received trial-by-trial feedback—moved individuals' decision weights toward objective probability weighting [31]. Re-entering description of outcomes in the sampling paradigm, in turn, appeared to increase the impact of rare events—the mere presentation effect [47]. This effect might also be at work in studies that have observed that rare events are overestimated or overweighted in decisions from descriptions [2,3,5].

Box 3. Outstanding questions

- Statistical and psychological factors have been suggested as causes for the description–experience gap. Which of these factors operate across all three experiential paradigms (Figure 1), and will ultimately be part of a comprehensive framework of the description–experience gap?
- Why do people rely on small samples (and recency) in decisions from experience? Does this reliance reflect memory constraints or lack of motivation? Is it an adaptive response to environmental properties (e.g. a fast-changing environment) that is generalized to static environments? Do people terminate search early because small samples may render choice easier by amplifying the difference between options?
- Which models or which class of models will best predict *search* and *choice* behavior across the three experiential paradigms of decisions from experience: models that assume that people form an explicit representation of sampled probabilities, or models that forgo such explicit representation of uncertainty such as some heuristics (e.g. minimax, natural-mean heuristic) or associative learning models?
- Decisions from experience have typically been studied in stable environments in which the underlying probability distribution does not change. Many real-world environments, however, are dynamic (e.g. a depleting patch of food). How do people make decisions from experience in such environments?
- Merely mentioning or presenting an event or focusing people's attention on one event seems to increase its psychological impact or increase its estimated frequency or probability, or both. What explains the psychological power of mere descriptions?
- When rare events *do* occur, their psychological impact often exceeds the impact deserved in light of their likelihoods. Beyond probability weighting functions, how can the heightened impact of *experienced* rare events be modeled? How long does their impact last?
- Many domains present people with experience, with descriptions, or with both. Established research paradigms, however, often focus on just one way of presenting information at the expense of others. For example, in social games or in deductive and inductive reasoning tasks people receive descriptions. Would a description–experience gap emerge? For example, could it be that a completely described social game nudges responders to focus on fair and efficient outcomes, whereas incomplete and experienced information reduces the impact of these concerns?

A final issue concerns the potential redundancy of the three experiential paradigms (Figure 1). They are not redundant. The partial-feedback paradigm, for example yields more safe choices than the full-feedback paradigm [18,48]. This difference might be a product of the hot-stove effect. As Mark Twain observed, if a cat happens to jump on a stove that is hot, it will never jump on one again, thus never learning to distinguish between hot and cold stoves. When feedback is limited to the obtained payoff, as in the partial-feedback paradigm, people are more likely to behave risk averse, thus reducing the probability of sampling alternatives with poor past financial and social outcomes [49,50]. This brings the discussion full circle to where it began. Generations growing up in a period of low stock returns appear to take an unusually cautious approach to investing, even decades later (although recent experience matters, too) [Malmendier, U. and Nagel, S., unpublished]. In other words, today's teenagers who have experienced this dramatic economic slump might years later enter the stock and housing market much more cautiously than their parents did. Perhaps, it is also they who will experience other global changes first-hand and behave differently than their parents did (Box 4).

Box 4. Decisions from experience: implications for risk behavior and precautions

How could one design a safer human environment? According to one perspective, one could aim to make people well aware of and tuned to the risks or hazards that they face by informing them about their statistical probabilities. For example, before each broadcast of a Soccer World Cup game, a warning could be aired telling viewers that watching a stressful soccer match more than doubles the risk of an acute cardiovascular event [51]. Such safety warnings could be extended to any activity that people might embark on, from having unprotected sex to consuming a caffeinated energy drink [52]. Thus, equipped with a complete *description* of the world's pitfalls, people could make rational decisions.

Admittedly, this fully described world is both unattainable—because of lack of information—and a caricature, although the ever-increasing amount of information on food products (e.g. health claims, content information), for example, suggests that some agencies believe that endowing consumers with comprehensive information is the best way to nudge them to healthy behaviors. This design perspective, however, neglects that people's risk perception and behavior are also shaped by personal experiences. If risks are rare, then chances are that a person mostly experiences their non-occurrence, thus generating less concern than the risks deserve according to their probabilities (now or in the future), and more concern in those rare cases in which they do occur. This simple regularity in people's experience could explain why warnings of the serious consequences of global warming do not make us behave much 'greener'—most of us have been hitherto spared personal experiences with those consequences [53]. It also explains why many Israeli drivers 'learned' to stop using a safety device that they bought for good money (a car radio with a detachable panel to protect from theft) within a year—nothing much had happened recently [54]. Moreover, it accounts for why providing a warning *before* rather than *after* having had a typically profitable or enjoyable experience with risk taking increases the impact of the warning [55].

How would a designer construct a safer environment, taking people's experience into account? Take the medical environment for illustration. One precaution that medical professionals are supposed to follow is to wear gloves when taking blood. Many doctors, however, do not comply, probably relying on their personal experiences, and downplaying the importance of this and other precautions [56]. How can one change their behavior? Severe sanctions administered with low probabilities are not likely to be effective. Although doctors possibly overweight the sanctions when they are announced, they might come to ignore them with the experience that noncompliance rarely gets punished. Alternatively, members of the medical team could be authorized to stop their colleagues if they saw them skipping a precaution, gently reminding them that this behavior puts them at risk. Such a design implies 'mild' punishment administered with high probability, and has been shown to yield stable increases in precautionary behavior [Erev et al., unpublished].

Concluding Remarks

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 that enters experience in investigations of risky choice find that people behave as if rare events receive less impact than they deserve, relative to their objective probabilities. These observations are not contradictory but describe how the mind functions in two different informational environments. In other words, one should not play off research on description-based and experience-based behavior—their contrast is enlightening. However, to better understand how people make

decisions with incomplete and uncertain information 'in the wild', there is a need to study experiential choices that are often representative of people's actual choices [57], and to learn from the rich repertoire of previous research into experience-based decisions and behaviors [14,58].

Acknowledgements

We are grateful to Robin Hau and Ben Newell for their helpful comments. We also thank Laura Wiles for editing the manuscript, and the Swiss National Science Foundation for a grant to the first author (100014-118283).

References

- 1 Taleb, N.N. (2007) *The Black Swan: The Impact of the Highly Improbable*, Random House
- 2 Lichtenstein, S. et al. (1978) Judged frequency of lethal events. *J. Exp. Psychol. Hum. Learn.* 4, 551–578
- 3 Viscusi, W.K. (2002) *Smoke-filled Rooms: A Postmortem on the Tobacco Deal*, University of Chicago Press
- 4 Redelmeier, D. and Kahneman, D. (1996) Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain* 116, 3–8
- 5 Tversky, A. and Kahneman, D. (1992) Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297–323
- 6 Friedman, M. and Savage, L. (1948) Utility analysis of choices involving risk. *J. Polit. Econ.* 56, 279–304
- 7 Posner, R.A. (2009) *A Failure of Capitalism: The Crisis of 08 and the Descent into Depression*, Harvard University Press
- 8 Akerlof, G.A. and Shiller, R.J. (2009) *Animal Spirits: How Human Psychology Drives the Economy and Why it Matters for Global Capitalism*, Princeton University Press
- 9 Lopes, L.L. (1983) Some thoughts on the psychological concept of risk. *J. Exp. Psychol. Hum. Percept. Perform.* 8, 137–144
- 10 Gigerenzer, G. et al. (2005) "A 30% chance of rain tomorrow": How does the public understand probabilistic weather forecasts? *Risk Anal.* 25, 623–629
- 11 Hertwig, R. et al. (2004) Decisions from experience and the effect of rare events in risky choice. *Psychol. Sci.* 15, 534–539
- 12 Knight, F.H. (1921) *Risk, Uncertainty, and Profit*, Houghton Mifflin Company
- 13 Hau, R. et al. (2009) Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *J. Behav. Decis. Making* DOI: 10.1002/bdm.665
- 14 Luce, R.D. and Suppes, P. (1965) Preference, utility, and subjective probability. In *Handbook of Mathematical Psychology (Vol. III)* (Luce, R.D. et al., eds), pp. 249–410, Wiley
- 15 Bussemeyer, J.R. (1985) Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *J. Exp. Psychol. Learn. Mem. Cogn.* 11, 538–564
- 16 Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, The MIT Press
- 17 Weber, E.U. et al. (2004) Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychol. Rev.* 111, 430–445
- 18 Yechiam, E. and Bussemeyer, J.R. (2006) The effect of foregone payoffs on underweighting small probability events. *J. Behav. Decis. Making* 19, 1–16
- 19 Barron, G. and Erev, I. (2003) Small feedback-based decisions and their limited correspondence to description-based decisions. *J. Behav. Decis. Making* 16, 215–233
- 20 Erev, I. and Barron, G. (2005) On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychol. Rev.* 112, 912–931
- 21 Cohen, J.D. (2007) Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Phil. Trans. R. Soc. B* 362, 933–942
- 22 Erev, I. et al. (2009) A choice prediction competition, for choices from experience and from description. *J. Behav. Decis. Making* DOI: 10.1002/bdm.683

- 23 Real, L.A. (1991) Animal choice behavior and the evolution of cognitive architecture. *Science* 253, 980–986
- 24 Real, L.A. (1992) Information processing and the evolutionary ecology of cognitive architecture. *Am. Nat.* 140, S108–S145
- 25 Shafir, S. *et al.* (2008) Perceptual accuracy and conflicting effects of certainty on risk-taking behavior. *Nature* 453, 917–921
- 26 Fox, C.R. and Hadar, L. (2006) “Decisions from experience” = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber, & Erev (2004). *Judgment Decis. Making* 1, 159–161
- 27 Hertwig, R. and Pleskac, T.J. (2008) The game of life: How small samples render choice simpler. In *The Probabilistic Mind: Prospects for Rational Models of Cognition* (Chater, N. and Oaksford, M., eds), pp. 209–236, Oxford University Press
- 28 Hau, R. *et al.* (2008) The description–experience gap in risky choice: The influence of sample size and experienced probabilities. *J. Behav. Decis. Making* 21, 493–518
- 29 Ungemach, C. *et al.* (2009) Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychol. Sci.* 20, 473–479
- 30 Rakow, T. *et al.* (2008) Biased samples not mode of presentation: Reexamining the apparent underweighting of rare events in experience-based choice. *Organ. Behav. Hum. Decis. Process.* 106, 168–179
- 31 Jessup, R.K. *et al.* (2008) Feedback produces divergence from prospect theory in descriptive choice. *Psychol. Sci.* 19, 1015–1022
- 32 Hadar, L. and Fox, C.R. (2009) Information asymmetry in decision from description versus decision from experience. *Judgment Decis. Making* 4, 317–325
- 33 Hogarth, R.M. and Einhorn, H.J. (1992) Order effects in belief updating: The belief adjustment model. *Cogn. Psychol.* 24, 1–55
- 34 Hertwig, R. *et al.* (2005) Judgments of risk frequencies: Tests of possible cognitive mechanisms. *J. Exp. Psychol. Learn. Mem. Cogn.* 35, 621–642
- 35 Gilboa, I. and Schmeidler, D. (1995) Case-based decision theory. *Q. J. Econ.* 30, 605–639
- 36 Gonzalez, C. *et al.* (2003) Instance-based learning in real-time dynamic decision making. *Cogn. Sci.* 27, 591–635
- 37 Klein, G. (1998) *Sources of Power: How People Make Decisions*, The MIT Press
- 38 Whittle, P. (1988) Activity allocation in a changing world. *J. Appl. Probab.* 25, 287–298
- 39 Biele, G. *et al.* (2009) Learning, risk attitude and hot stoves in restless bandit problems. *J. Maths. Psychol.* 53, 3,155–167
- 40 Hertwig, R. (in press) The psychology and rationality of decisions from experience, Synthese
- 41 Hertwig, R. *et al.* (2006) The role of information sampling in risky choice. In *Information Sampling and Adaptive Cognition* (Fiedler, K. and Juslin, P., eds), pp. 75–91, Cambridge University Press
- 42 Yechiam, E. and Busemeyer, J.R. (2005) Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychon. Bull. Rev.* 12, 387–402
- 43 Rakow, T. and Newell, B.R. (in press) Degrees of uncertainty: An overview and framework for research on experience-based choice. *J. Behav. Decis. Making* DOI: 10.1002/bdm.681
- 44 Gigerenzer, G. and Hoffrage, U. (1995) How to improve Bayesian reasoning without instructions: Frequency formats. *Psychol. Rev.* 102, 684–704
- 45 Hoffrage, U. *et al.* (2000) Communicating statistical information. *Science* 290, 2261–2262
- 46 Newell, B.R. and Rakow, T. (2007) The role of experience in decisions from description. *Psychon. Bull. Rev.* 14, 1133–1139
- 47 Erev, I. *et al.* (2008) What impacts the impact of rare events. *J. Risk Uncertain.* 36, 153–177
- 48 Grosskopf, B. *et al.* (2006) Foregone with the wind: Indirect payoff information and its implications for choice. *Int. J. Game Theory* 34, 285–302
- 49 Denrell, J. (2007) Adaptive learning and risk taking. *Psychol. Rev.* 114, 177–187
- 50 Denrell, J. (2005) Why most people disapprove of me: Experience sampling in impression formation. *Psychol. Rev.* 112, 951–978
- 51 Wilpert-Lampen, U. *et al.* (2008) Cardiovascular events during world cup soccer. *N. Engl. J. Med.* 358, 475–483
- 52 Reissig, C.J. *et al.* (2009) Caffeinated energy drinks—a growing problem. *Drug Alcohol Depend.* 99, 1–10
- 53 Weber, E. (2006) Experience-based and description-based perceptions of long-term risk: Why global warming does not scare us (yet). *Clim. Change* 77, 103–120
- 54 Yechiam, E. *et al.* (2006) The effect of experience on using a safety device. *Safety Sci.* 44, 515–522
- 55 Barron, G. *et al.* (2008) The effect of safe experience on a warnings’ impact: Sex, drugs, and rock-n-roll. *Organ. Behav. Hum. Decis. Process.* 106, 125–142
- 56 Stein, A.D. *et al.* (2003) A survey of doctors’ and nurses’ knowledge, attitudes and compliance with infection control guidelines in Birmingham teaching hospitals. *J. Hosp. Infect.* 54, 68–73
- 57 Dhami, M. *et al.* (2004) The role of representative design in an ecological approach to cognition. *Psychol. Bull.* 130, 959–988
- 58 Erev, I. and Haruvy, E. (in press) Learning and the economics of small decisions. In *The Handbook of Experimental Economics* (Kagel, J.H. and Roth, A.E., eds.), Princeton University Press