Cognition 115 (2010) 225–237

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Decisions from experience: Why small samples?

Ralph Hertwig [a,*], Timothy J. Pleskac [b]

[a] *University of Basel, Department of Psychology, Missionsstrasse 60/64, 4055 Basel, Switzerland*
[b] *Michigan State University, Department of Psychology, 231 Psychology Building, East Lansing, MI 48824, USA*

ARTICLE INFO

Article history:
Received 4 July 2007
Revised 14 December 2009
Accepted 16 December 2009

Keywords:
Sampling
Decisions from experience
Choice
Small samples
Amplification effect

ABSTRACT

In many decisions we cannot consult explicit statistics telling us about the risks involved in our actions. In lieu of such data, we can arrive at an understanding of our dicey options by sampling from them. The size of the samples that we take determines, ceteris paribus, how good our choices will be. Studies of decisions from experience have observed that people tend to rely on relatively small samples from payoff distributions, and small samples are at times rendered even smaller because of recency. We suggest one contributing and previously unnoticed reason for reliance on frugal search: Small samples amplify the difference between the expected earnings associated with the payoff distributions, thus making the options more distinct and choice easier. We describe the magnitude of this amplification effect, and the potential costs that it exacts, and we empirically test four of its implications.

1. Introduction

Being freed of tasks such as mating and territorial defense, worker bees are particularly well suited for evolutionary studies of choice (Real, 1992). Time and again, they need to decide which area of potential food resources to exploit—a choice rendered thorny due to changing ecological conditions. Using an artificial patch of flowers with varying colors signaling different amounts and probabilities of nectar rewards, Real (1991, 1992) investigated how a bee made choices in its uncertain habitat. He found that if a bee were maximizing expected utility, it chose as if it "misjudges the objective probabilities and underestimates the rare event"—a finding that is in "contrast with the conclusions of Kahneman and Tversky" (Real, 1992 p. S132). Kahneman and Tversky (1979) asserted that rare events—at least for humans—have more impact than they deserve according to their objective probabilities.

Investigations giving rise to the conclusion of opposite perceptions of rarity, however, have dealt their agents—bees versus humans—very different cards. In Real's studies (1991) bees were initially ignorant of the habitat's distribution of nectar and learned about their environment while foraging. What Real described as bees' misjudgments of rare events relates to the discrepancy between bees' implicit estimates of the likelihood of the rewards, inferred from their behavior, and actual probabilities. In contrast, during studies on how humans make decisions under risk, people receive perfect information about the payoff associated with each option and the probability of those payoffs (e.g., 3 with certainty versus 32 with probability of .1, 0 otherwise). They thus make what Hertwig, Barron, Weber, and Erev (2004) referred to as *decisions from descriptions*, rendering foraging for information and estimating probabilities superfluous.

2. Perception of rarity and small samples

What happens when people, like bees, sample information from uncertain environments, thus making *decisions from experience* (Hertwig et al., 2004)? Such decisions represent situations in which, being ignorant of the payoffs, agents may resort to experience: experience garnered

* Corresponding author. Tel.: +41 61 267 0611; fax: +41 61 267 0441.
E-mail address: ralph.hertwig@unibas.ch (R. Hertwig).

0010-0277/$ - see front matter
doi:10.1016/j.cognition.2009.12.009

through sampling from the payoff distributions and eventually arriving at what Knight (1921) called "statistical probabilities" based on the empirical classification of instances (see Hau, Pleskac, & Hertwig, 2010). Such sampling can occur in two fundamentally different ecologies, one with an inherent trade-off between exploiting and exploring options (see Barron & Erev, 2003; Berry & Fristedt, 1985; Erev & Barron, 2005), and one in which the agent's only objective, at least initially, is exploration for information (Hertwig et al., 2004; Weber, Shafir, & Blais, 2004). In the former ecology, the sampled outcomes simultaneously provide information and payoffs to the agent. In contrast, the sampled outcomes in the latter exploration-only environment merely provide information, much like attending a free wine-tasting fair, perusing the Gault-Millau or the Michelin Guide to select a restaurant, or taking a quick look at the online traffic cams. Exploitation of the options for payoffs—for example, dining at one of the acclaimed gastronomic temples—only comes after search for information was terminated.

We focus on this second ecology. Hertwig et al.'s (2004) study illustrates how it can be investigated. Participants were asked to choose between two gambles (payoff distributions). Lacking knowledge about them, they could explore each distribution by sampling from them. Specifically, people saw two boxes on a computer screen representing two possible gambles or outcome distributions. Clicking on a box triggered a random draw of an outcome from the associated distribution. They were encouraged to sample until they felt confident enough to decide which box was "better", in the sense that they would prefer to draw from it during a final trial involving real monetary payoffs. Hertwig et al. compared respondents' choices in this final trial with those of a group who made decisions from description from the same problems. Choices differed drastically and systematically. In description-based choices, rare outcomes were consistent with prospect theory's assumption of overweighting of rare events (Kahneman & Tversky, 1979). In experience-based choices, however, people behaved as if rare outcomes had less impact than they deserved according to their objective probabilities.[1] This description–experience gap has been replicated across a wide range of studies (e.g., Erev, Glozman, & Hertwig, 2008; Hau et al., 2010; Hau, Pleskac, Kiefer, & Hertwig, 2008; Rakow, Demes, & Newell, 2008; Ungemach, Chater, & Stewart, 2009; Weber et al., 2004; for a review of experience-based decision making see Hertwig & Erev, 2009, and for a special issue on experience-based decision making see Rakow & Newell, 2010).

For bees and humans, the source of the relative lack of appreciation of rare events appears to be the same: *the small samples on which they base their choice*. Based on investigations across different floral reward distributions, Real (1992) concluded that "bees frame their decisions on the basis of only a few visits" (p. S133). Focusing on humans, Hertwig et al. (2004) found that the typical number of draws that respondents made was approximately seven from each deck (and a median of 15 across both decks). Hau et al. (2010) reviewed the sample sizes in five subsequent decisions from experience studies in which people were instructed to sample until they felt confident enough to decide between the two payoff distributions, and in which monetarily only their final choice mattered (unlike in Barron & Erev, 2003). Consistent with Hertwig et al.'s original observation, sample sizes fell in a surprisingly small range from 9 to 19 draws, amounting to nearly 7 ± 2 draws from each deck. The only outlier was Hau et al.'s (2008) observation of 33 draws (Study 2) in a condition involving an order of magnitude larger payoff than in other studies.

Based on these observations, it seems fair to say that so far, ceteris paribus, the finding of relatively small sample sizes in decisions from experience studies is robust. Admittedly, referring to these sample sizes as "small" is somewhat arbitrary, as, of course, samples could have been even smaller. Yet, if one keeps in mind that each draw requires an investment of merely a few seconds and that people's small samples systematically and noticeably misrepresented the probabilities of rare events, the conclusion that people in these studies framed their decisions on relatively "small" samples appears justified. For example, in Hau et al.'s (2008) Study 1, respondents who sampled a median of 11 draws across both distributions did not even encounter the rare event in 50.3% of all trials, and the median difference between the experienced relative frequency and the objective probability of the nonzero outcomes was 10.9 percentage points.

Small samples interact with the statistical structure of the environment in systematic ways. To see this, consider a gambling environment with binomially distributed outcomes, and where n is the number of draws from a particular gamble or deck, and p is the probability of the maximum outcome in the gamble. When n is small (i.e., few draws) and/or p is small (i.e. the event is rare) the binomial distribution is skewed for the number of times this rare outcome will be observed in n independent trials. 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 a person samples 10 times from a distribution in which the critical event has a probability of .1 and estimates the event's probability to be the proportion in the sample that she observes. The probabilities that she will observe the critical event more than once, less than once, or exactly once are .26, .35, and .39, respectively. That is, the person is more likely to underestimate than to overestimate the frequency of the rare event (.35 versus .26). This asymmetry decreases the larger n becomes, for example, from 9 percentage points to 7, 6, and 5 percentage points for sample sizes of 10, 20, 30, and 40, respectively.

Next, we consider possible reasons why people appear content with small samples, and how the fact that the rare event is more likely to be under- than overrepresented in small samples changes the options that people actually experience.

[1] Fox and Hadar (2006) correctly pointed out that this weighting pattern in experience-based choice need not be in conflict with prospect theory, and that prospect theory can account for decisions from experience if its input is the sampled rather than the experienced probabilities (see also Hau, Pleskac, Kiefer, & Hertwig, 2008).