How short- and long-run aspirations impact search and choice in decisions from experience ☆

Dirk U. Wulff [a,b,*], Thomas T. Hills [c], Ralph Hertwig [a]

[a] Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany
[b] University of Basel, Switzerland
[c] Department of Psychology, University of Warwick, Coventry, UK

ARTICLE INFO

ABSTRACT

To what extent do people adapt their information search policies and subsequent decisions to the long- and short-run consequences of choice environments? To address this question, we investigated exploration and exploitation policies in choice environments that involved single or multiple plays. We further compared behavior in these environments with behavior in the standard sampling paradigm. Frequently used in research on decision from experience, this paradigm does not explicitly implement the choice in terms of the short or long run. Results showed that people searched more in the multi-play environment than in the single-play environment. Moreover, the substantial search effort in the multi-play environment was conducive to choices consistent with expected value maximization, whereas the lesser search effort in the single-play environment was compatible with the goal of maximizing the chance of winning something. Furthermore, choice and search behaviors in the sampling paradigm predominantly echoed those observed in the single-play environment. This suggests that, when not instructed otherwise, participants in the sampling paradigm appear to favor search and choice strategies that embody short-run aspirations. Finally, the present findings challenge the revealed preference approach in decisions from experience, while also suggesting that information search may be an important and potentially even better signal of preference or aspirations than choice.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Choices between uncertain options can be interpreted as representing either single-play or multi-play decisions. A lottery ticket, for instance, represents a single-play decision; its entry price entitles the player to exactly one play of the lottery. A choice to buy car insurance, on the other hand, guarantees against repeated plays of a gamble that is realized each time the car is driven. More generally, decisions to buy products that will be consumed either once (e.g., a dinner in a gourmet restaurant) or many times (e.g., a pair of running shoes) involve different time horizons. These may, in turn, prompt differences in the decisions made as well as in the information needed to render a decision. For illustration, consider the offer that Nobel-prize winning economist Paul Samuelson

(1963) once made his lunch partners: "to bet each $200 to $100 that the side of the coin *they* specified would not appear at the first toss" (p. 50). One colleague, whom Samuelson identified as a distinguished scholar but otherwise granted anonymity, responded to the offer by saying: "I won't bet because I would feel the $100 loss more than the $200 gain. But I'll take you on if you promise to let me make 100 such bets" (p. 2). Samuelson (1963) considered his colleague's preference to be inconsistent with expected utility theory and, by extension, to be irrational (a fallacy of large numbers): "... no sequence is acceptable if each of its single plays is not acceptable" (p. 3).

More recent analyses, however, have concluded that models of expected utility theory—by many considered *the* normative theory of individual decision making—can in fact capture the colleague's preference for safety in numbers, assuming that the 100 bets are aggregated to a single choice. Ex ante aggregation brings the final distribution of potential payoffs of a gamble much closer to its expected value and accordingly reduces in the above example the likelihood of a loss (Aloysius, 2007; Kahneman & Lovallo, 1993; Wedell, 2011; see also Peköz, 2002). Thus, in decisions under uncertainty, single-play and multi-play choice environments

effectively entail different payoff distributions. A single utility function can thus be consistent with both preferences of Samuelson's colleague, as well as the general observation that the higher expected value option is preferred in multi-play but not in single-play situations (Aloysius, 2007; Montgomery & Adelbratt, 1982).

What is less well understood—and the focus of this article—is how people respond to single- and multi-play environments in which they first have to search for information before making a choice. We address this question by implementing the two choice protocols described in Samuelson's anecdote within the sampling paradigm, a popular design used in research on decisions from experience (Hertwig, Barron, Weber, & Erev, 2004; Hertwig & Erev, 2009). In the sampling paradigm, people first explore the possible outcomes[1] of risky options in a self-directed and self-terminated sampling process before making a decision based on their sampled experience.

By investigating information search and choice, we can add search as a new dimension to the analysis of the effects of single-play and multi-play choice environments (DeKay & Kim, 2005; Montgomery & Adelbratt, 1982; Redelmeier & Tversky, 1992; see Wedell, 2011, for a brief review). Our investigation will also permit us to further analyze a recently discovered relationship between information search and choice that may originate from the pursuit of short-run versus long-run aspirations (Hills & Hertwig, 2010). Finally, systematic differences in information search between single- and multi-play environments will help us to further understand how preferences, as revealed by choices, are further impacted by the search that precedes them—a problem that generalizes to all tasks in which the actually experienced environment is a function of the organism's information search.

In the following, we first review pertinent empirical literature about expected utility in relation to single- and multi-play decisions. We then review findings on information search in decisions from experience, before describing how we link these lines of research.

1.1. Aspirations and the importance of expected utility in single- and multi-play decisions

References to expected utility often invite one to say, subtly and under one's breath, 'long-run' expected utility. Some may argue that the addition of 'long-run' is redundant. Given the broad class of single-play decisions where expected utility does not immediately apply (Lopes, 1981), however, we would not agree. An offer to pay $5 to play once a gamble that pays off $100 with probability .1 and $0 otherwise will leave the gambler poorer by $5 nine times out of ten (Fig. 1). This is true regardless of the amount of the non-zero payoff, be it $100, $1000, or even $100,000. However, the opportunity to play this gamble 100 times increases the probability of coming out ahead to above 50% (by 'coming out ahead,' we refer to the short-run aspiration of winning any non-zero amount). Anyone with a strict requirement of more than a non-zero return on their investment should avoid the single-play gamble, because in most realizations it will lead to losses.

The importance of achieving a minimal aspiration and its role in explaining many choice anomalies has been well explored (Koop & Johnson, 2012; Lopes, 1996; see also Lopes & Oden, 1999). The key argument is that many of the mathematical prosthetics added to expected value theory (e.g., polynomial utility functions and subjective probability curves) are unnecessary if one considers that
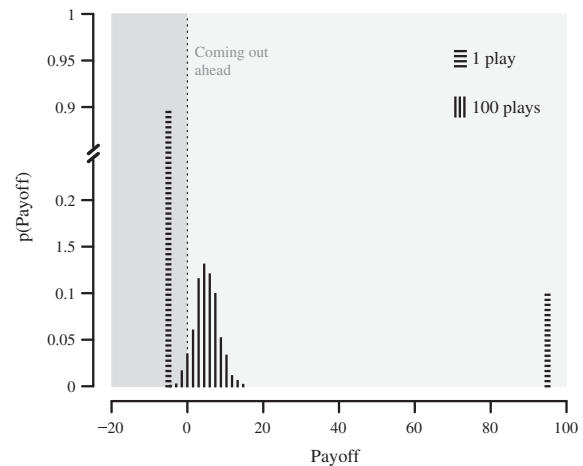

Fig. 1. The influence of 1 play versus 100 plays on the probability distribution of payoffs at the end of play. Results for a single-play gamble costing 5 to play and promising an outcome of 100 with a probability of .1 and otherwise 0. The results for the multi-play gamble reflect the expected payoff per single play (each costing 5) of that gamble.

in many situations it may not be rational to pursue the expected value or long-run expected utility, but rather "the probability of coming out ahead" (Lopes, 1981, p. 377)[2]. Indeed, studies investigating peoples' choices of single- and repeated-play gambles have found fewer violations of expected utility theory when people play repeatedly than when they play once (Camilleri & Newell, 2013; Keren, 1991; Keren & Wagenaar, 1987; Liu & Colman, 2009; Wedell & Böckenholt, 1990, see Caraco, 1980; Houston & McNamara, 1999; Stephens, 1981, 2001, for a similar discussion in behavioral ecology).

According to this line of theorizing, single-play and multi-play trigger short- and long-run aspirations, with short-run aspirations indicating an increased preference for the option that is most likely to come out ahead and long-run aspirations favoring the option offering the higher expected value. Yet, let us emphasize that aspirations are not the only way to conceptualize behavior across single-play and multi-play choice (Aloysius, 2007; Langer & Weber, 2001; Lopes, 1996; Tversky & Bar-Hillel, 1983; Wedell, 2011). For instance, any mechanism explaining risk aversion, that is, the preference for the option with the lower variance, is under most circumstances also capable of explaining differences in choice (but not in search), even when the expected values of single- and multi-play scenarios are the same (as in Fig. 1). Such explanations include non-linear transformations of outcomes and non-linear transformations of probabilities (Kahneman & Tversky, 1979; see Wedell, 2011).

1.2. The relation between aspirations and information search

Assuming that decision makers conceive of single- and multi-play environments differently, then one may ask whether and how the process of information acquisition differs across these environments. If decisions and decision rules in these different

---

[1] In what follows, we use the term *outcomes* to refer to the set of values that could result from choosing an option or that are experienced in the process of sampling from an option. The term *payoffs*, in contrast, refers to the monetary consequences of choosing an option and that are contingent on the payoff scheme employed.

[2] The short-run aspiration of maximizing the chance of coming out ahead has sometimes been used interchangeably with the aspiration of maximizing some percentile of the outcome distribution (e.g., the median). Although both criteria would essentially produce identical predictions in our study, we focus on the aspiration of coming out ahead for two reasons. First, for two-outcome gambles as used in our study, the median is not well-defined. Second, the aspiration of maximizing the chance of coming out ahead corresponds more closely to the short-run criteria implemented in the literature on risk-sensitive foraging (e.g., Stephens, 2001).

environments require more or less information, then adaptive search strategies that meet such differential demands could foster better decisions by increasing efficiency. Consistent with the idea of adaptive information search, Hills and Hertwig (2010) found that specific information search policies in decisions from experience are associated with specific decision policies. Individuals who took more samples and switched less often between options were more likely to choose options associated with maximizing expected utility, whereas individuals who took fewer samples and switched more frequently between options tended to choose options that offered a higher chance of coming out ahead. Specifically, individuals who showed frugal search and avid switching appeared to accomplish this by comparing the promised return on the options between switches and choosing the option that won most of the time. Individuals who showed avid search and frugal switching, on the other hand, appeared to choose the option with the higher mean return computed from all collected samples. Though important for understanding the interplay of information search and choice, these findings are correlative. Consequently, they cannot discern between two possibilities: Do search policies sway later decision strategies, or do preselected decision strategies shape subsequent search policies?

Hills and Hertwig (2010) speculated that the correspondence between search and decisions could be driven by different goals. However, it is also plausible that the cognitive control of attention drives search, irrespective of top-down aspirations. Specifically, the sampling paradigm in research on decisions from experience has participants commonly make a choice between two options after they have had the opportunity to explore (sample) them (Hertwig et al., 2004). A person may sample the outcomes $0, $0, $0, and $32 for one option and $3, $3, and $3 for the other. Following a choice, the person would receive the value of one randomly drawn outcome for the option he or she decided on. Using this sampling paradigm, Rakow, Demes, and Newell (2008, see also Ashby & Rakow, 2014; Wulff & Pachur, in press) observed that total sample size and subsample sizes (samples between switches) were positively correlated with working memory span, a measure proposed to be associated with attentional control (e.g., Conway, Cowan, & Bunting, 2001; Kane & Engle, 2000). Consequently, it is unclear to what extent cognitive control of attention, short-term versus long-term aspirations, or both drive search in this paradigm.

Apart from being a matter of theoretical interest, which mechanism—aspirations, cognitive control of attention, or both—drives information search has important methodological implications. If aspirations drive search, aspirations—that is, preference structures—determine not only what decision makers choose, but how they search for information prior to choice. As we will explain in the discussion, this means that their preference structure may not be uniquely identified on the basis of their choices, an issue with notable consequences.

1.3. Testing the impact of short- and long-run aspirations on decisions from experience

In order to determine the impact of short- and long-run aspirations on information search and choice in decisions from experience, we compared two conditions in which choices are followed by either short-run or long-run consequences (similarly to Camilleri & Newell, 2013; Wedell & Böckenholt, 1990). Specifically, we informed participants that their final payoff depended either on a single, randomly chosen outcome from *one* of their chosen options, multiplied by 100 (*single-play condition*), or on 100 random draws from one of their chosen options (*multi-play condition)*. Except for these payoff schemes, both conditions were identical. This means that—apart from the influence of sampling—both conditions see the same options. The distributional

consequences of a single play or multiple plays for the payoff (see Fig. 1) are thus left to participants to infer. This set-up allows us to directly evaluate and compare patterns of information search and choices across identical environments that, similar to real word environments, afford the organism the freedom to search and interpret them as they like.

Furthermore, we designed decision problems with a structure often employed in the decisions-from-experience literature, requiring a choice between a risky option (with two outcomes) and a safe option. These problems have the property that the number of samples needed to detect the option that promises the larger probability of coming out ahead is *less* than the number of samples needed to detect, with the same precision, the option with the larger mean. We expect that assigning short- and long-run consequences to otherwise identical choice environments will lead to individuals adapting their information search and choice policies to their aspirations. In other words, people may not only be adaptive decision makers (Payne, Bettman, & Johnson, 1993), but also adaptive information searchers. Furthermore, we predict that aspiration-induced differences in search will prompt systematically different experiences of identical options. Specifically, in the single-play condition, relative to the multi-play condition, fewer people will experience the rare event—because of smaller sample sizes. By extension, we predict that participants in the multi-play condition, relative to the single-play one, will more likely choose the option with the higher expected value.

In addition to the single- and multi-play condition we realize a third condition that implements the payoff scheme commonly employed in studies of decisions from experience (with the sampling paradigm; see Hertwig, 2015), that is, to reward a participants with a single draw from *every* chosen option. In this standard payoff scheme, a decision maker can choose to interpret the situation as concerning single-play or multiple-play environments. We therefore expect that behavior in the standard sampling paradigm will lie somewhere in between that of respondents in the single-play and multi-play conditions, both in search and choice, thus reflecting the interindividual heterogeneity of potential short- and long-run strategies previously observed (Hills & Hertwig, 2010).

Finally, in order to evaluate the potential role of attentional control on information search, we also measured each participant's operation span. Operation span is a complex working memory span measure that taps into a person's ability to store and retrieve a sequence of individual tokens over intermittent distractor tasks. Alternative to or in addition to short- and long-run aspirations, attentional control may determine search. Based on previous findings (Rakow et al., 2008), we expect higher operation spans to be associated with taking more samples and with fewer switches between the options. Moreover, one may speculate that attentional control is also linked to choice. The short- and long-run choice policies suggested by Hills and Hertwig (2010) are likely to require different levels of cognitive effort. Consequently, the coupling of search and choice may also be caused and explained by individual differences in attentional control capacities (vs. differences in short-run and long-run aspirations).

2. Method

2.1. Participants

We collected data from 124 students of the University of Basel, with a mean age of 24 years. Participants were rewarded with either course credit or a fixed payment of approximately $13. In addition, participants received a performance-based bonus as a result of their choices.

2.2. Materials

We designed 12 decision problems (Table 1). The two options within each problem offered the opportunity to maximize either the long-run expectation (higher mean or expected value) or the probability of coming out ahead (higher median). Each problem presented a choice between a high-outcome rare event ($p \leqslant .15$) in the higher expected value option (otherwise 0) and a small but relatively secure outcome ($p \geqslant .7$) in the lower expected value option (otherwise 0). All problems share the property that more samples are required to spot the option with the higher mean than the option with the higher probability of coming out ahead. To demonstrate this, we simulated 10,000 decisions for each problem and determined how many samples would be needed to identify the option with the higher mean versus the higher chance of coming out ahead, given some level of precision. Identifying the latter with a probability of, for instance, at least 80% requires a much smaller sample than identifying the higher mean option (on average, about 4 vs. 34 draws per option). Of course, we do not expect our participants to know in advance the options' possible outcomes and their probabilities; however, they may rapidly develop an understanding of the options once sampling begins. Sample sizes will reflect this grasp.

In order to prevent participants from inferring that the option with the rare and consequential event always promised the higher expected value, we intermixed four problems with a different structure (Appendix A, Table A1), resulting in a total of 16 decision problems. All were included in our analysis, because the predictions for the search and decision strategies are qualitatively independent of the structure of the decision problems. To measure individuals' working memory capacity, we used the automated version of the *operation span task* (Unsworth, Heitz, Schrock, & Engle, 2005). This computer-based task measures people's ability to remember sets of letters (e.g., E A D) that appear, one letter at a time, following simple math problems (e.g., $[1 \times 2] + 1 = ?$). The operation span score was determined as the sum of all correctly recalled letter sets (over a series of sets with lengths ranging from 3 to 7 letters).

2.3. Procedure

Problems were presented on a computer screen. Participants were randomly assigned to one of the three conditions. In the multi-play condition, they were instructed that their payoff would be determined by randomly selecting one of their final choices and then taking 100 random draws from the selected option (e.g., 100 draws from selected option *H* in Problem 1: 92 with probability .05; Table 1). In the single-play condition, participants were instructed that one of their chosen options would be randomly selected; a single random draw from this option would then be taken, and the resulting outcome would be multiplied by a factor of 100. This procedure renders the magnitudes of the expected payoff in the single- and multi-play conditions identical. Finally, the *standard* condition implemented the payoff modality used in past studies (e.g., Hertwig et al., 2004). Specifically, a random draw from each chosen option across the 16 problems determined the payoff. Because this meant that 16 draws were incentivized in the standard condition, relative to 100 draws in the other two conditions, we matched the expected returns across conditions by multiplying each randomly drawn outcome in the standard condition by a factor of 6.

Before participants turned to the 16 decision problems (presented in a random order), they worked on three practice trials. For each problem, they were able to sample from the two options as extensively and in whatever fashion they liked. Once search was terminated, they proceeded to their final choice by clicking a

Table 1
Decision problems employed in the three conditions.

Problem	H	L
1	92[a] with $p = .05$	3 with certainty
2	34 with $p = .05$	1 with certainty
3	120 with $p = .05$	5 with $p = .70$
4	44 with $p = .05$	2 with $p = .70$
5	70 with $p = .10$	4 with certainty
6	16 with $p = .10$	1 with certainty
7	54 with $p = .10$	4 with $p = .75$
8	23 with $p = .10$	2 with $p = .75$
9	35 with $p = .15$	3 with certainty
10	21 with $p = .15$	2 with certainty
11	48 with $p = .15$	5 with $p = .80$
12	9 with $p = .15$	1 with $p = .80$

Note: H = option with the higher expected value (as calculated by probability × monetary value); L = option with the lower expected value.
[a] In order to provide identical incentives across conditions, we matched the expected returns across conditions by multiplying each randomly drawn outcome in the standard condition by a factor of 6.

button. The operation span measure was administered once all choices were completed. At the end of the experiment, participants received a bonus as a result of their choices, paid out one-to-one in accordance with the conditions' payoff scheme.

For the purpose of data analysis, we set a threshold such that participants had to sample at least once from both options in at least half of the problems; five participants (of 124; 4%) failed to meet this criterion and were removed from the analysis. Additionally, all trials in which a person sampled only a single option were removed. The following analyses are thus based on 95% of all trials provided by 119 participants.

3. Results

3.1. The influence of single-play and repeated-play on search and switching

As Fig. 2 shows, amount of search in the repeated-play condition was very different from that in the single-play condition. On average, participants in the former condition took about 10.6 samples more than those in the latter condition (total sample size: $t[77] = 2.37$, $p = .02$, $d = 0.53$).[3] This difference was mainly driven by larger samples from the risky or riskier option (option *H* in Table 1) accounting for 7.1 of the additional 10.6 samples, $t(77) = 2.13$, $p = .04$. However, there also was more extensive search in the safe or safer option (option *L* in Table 1; $t[77] = 2.43$, $p = .02$). Because participants in all conditions took, on average, about two samples from option *H* for every one sample from option *L* ($n_H/n_L = 1.75$–2.05), it appears that participants in the multi-play condition increased their search effort about equally for both options. What about search in the standard sampling condition? Sample size was different from that achieved in the multi-play condition, $t(77.1) = 2.26$, $p = .03$, but almost identical to that in the single-play condition, $t(78) = .1$, $p = .95$.

Hills and Hertwig (2010) found that switch rate was inversely correlated with the choice of the option with the higher expected value. This observation invites the question as to whether decision

---

[3] Global tests of significance were omitted due to clear hypotheses for the pairwise group comparisons. All reported t values were derived from mixed effects analyses predicting the outcome variable on the problem level while controlling for the subject variable via the inclusion of a random intercept. Tests were performed using the statistical software R (R Development Core Team, 2008) and the packages lme4 and lmerTest. Specifically, Gaussian linear models were estimated using REML and Satterthwaite's approximation for degrees of freedom, the default method in lmerTest. The effect size d is a standardized measure, and $d = .2$, .5, and .8 denote small, medium, and large effects, respectively (Cohen, 1988).
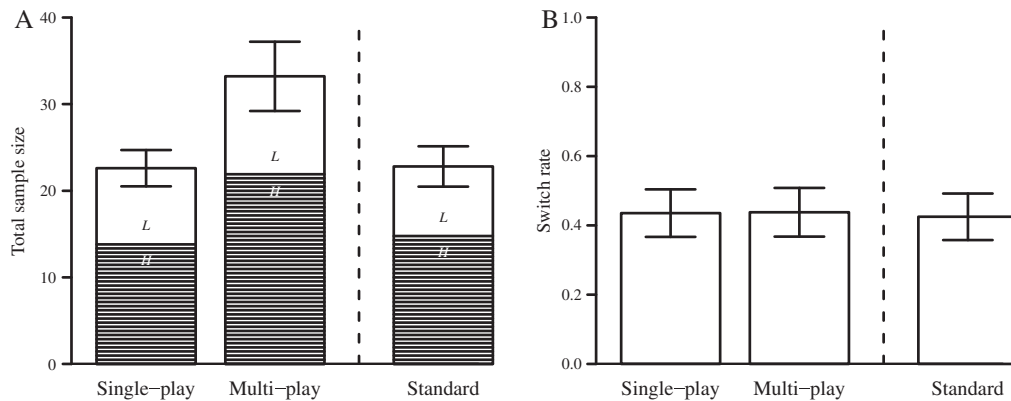
Fig. 2. Total sample size (A) and switch rate (B), separately for the three conditions: single-play, multi-play, and the standard sampling paradigm. In panel A, bars are further split into samples taken from the *H* and *L* options (see Table 1). Error bars represent the standard error of the mean.

strategies govern switch rate. Our results suggest they do not. The number of switches per sample was not significantly different between conditions, $t(77) = 0.33$, $p = .75$ (see Fig. 2B). This finding contradicts Hills and Hertwig's (2010) suggestion that frequent switching may be caused by the aspiration of short-run maximization. Moreover, switch rate in the standard paradigm was not statistically different from that in the multi-play (switch rate: $t[77] = 0.06$, $p = .95$) or single-play condition (switch rate: $t[78] = 0.23$, $p = .77$). To summarize, the single- versus multi-play instructions markedly affected sample size; the switch rate, in contrast, appeared to be less sensitive to the difference in aspiration.

3.2. How single- versus multi-play shape individuals' experience

One key characteristic of the sampling paradigm is that an individual's sample size inevitably shapes his or her experience of the events' probabilities. In particular, rare events are often not encountered when sample sizes are small; and even if they are observed, the number of people who experience them less frequently than expected exceeds the number who experience them more frequently than expected (as a consequence of the skewness of the binomial probability distribution for small *n*s and small *p*s; see Hertwig et al., 2004). The large difference in sample size observed for the multi-play and single-play conditions is thus likely to translate into different experiences: Participants in the multi-play condition are, ceteris paribus, more likely to be *cognizant* of the rare positive outcome in option *H* than are participants in the single-play condition; furthermore, the former can be expected to have *experienced* the rare event more often than the latter. Is this indeed the case?

In terms of awareness of rare events, our findings showed that participants in the single-play condition were 1.3 times more likely to miss the rare event than were participants in the multi-play condition (33% vs. 26%, see Fig. 3A). Although this difference was in the expected direction (in light of the different sample sizes; see Fig. 2A), it was not significant (logit link: $z = 1.42$, $p = .16$). But did the frequency with which people experienced a rare event (given that it was encountered once) differ between conditions? Fig. 3B shows that the rare event was encountered 1.5 times more often in the multi-play condition than in the single-play condition, $t(77) = 2$, $p = 0.049$, $d = .45$.

Why were people in the single-play condition not markedly more likely to miss the rare event than people in the multi-play condition? One possible explanation relates to optional stopping. To the extent that sample size (i.e., number of draws) is determined at the outset of the sampling process, the binomial probability distribution governs the sampling process; furthermore, it

implies a smaller chance to observe the rare event with smaller sample sizes (see Hertwig et al., 2004; Hertwig & Pleskac, 2010). Alternatively, however, stopping may be controlled by the actually experienced outcomes, thus rendering the binomial distribution an inappropriate model (see, e.g., Berger & Berry, 1988, for a discussion of optional stopping in statistical inference). Our data indicate that experience indeed matters for the decision to stop. Fig. 4 plots the differences between the observed probabilities of the rare events and their true probabilities in the multi-play, $t(37.2) = 4.66$, $p < 0.001$, and single-play condition, $t(39.2) = 2.4$, $p = 0.02$. In both conditions, people experienced the rare event more frequently than expected, consistent with outcome-dependent stopping—in other words, people appear to have stopped shortly after observing a rare event.

Fewer samples in total and fewer observations of the rare event in option *H* relative to the multi-play condition (Fig. 3A; multi-play vs. standard: $t[77.2] = 2.04$, $p = 0.04$) suggest that participants in the single-play and the standard sampling condition mustered experiences that were similar. This did not hold for every dimension, however. Participants in the standard condition experienced rare events more often than those in the single-play condition (see Fig. 3B; $t[77.2] = 2.04$, $p = 0.04$). Thus, although the single-play and the standard condition were alike in terms of sample size and switching, the samples on which they based their choices were not identical.

3.3. The influence of single-play and repeated-play on choice

We started out, among other hypotheses, by predicting that induction of long-run aspirations would lead to favoring the option with the higher expected value. Do our data support this hypothesis? Fig. 5A plots the proportion of choices of option *H*. Consistent with the above hypothesis, the proportion of choices of option *H* in the multi-play condition was 63%, substantially higher than the 49% observed in the single-play condition ($z = 3.16$, $p = 0.002$, $OR = 1.92$). Relatedly, the proportion of choices of option *H* in the standard condition was 51%, which was not statistically different from that observed in the single-play condition ($z = 0.67$, $p = 0.5$), but was different from that observed in the multi-play condition ($z = 2.66$, $p = 0.01$).[4] These similarities and differences thus suggest that long-run aspirations are conducive to expected-value maximization and that respondents appear to perceive the standard

---

[4] A mixed effects regression indicated that *H* choices were not influenced by the presence or absence of certainty in the *L* option ($z = 1.53$, $p = 0.13$). Because none of the following analyses were influenced by a comparison of safe and risky options, the following results are collapsed across both.
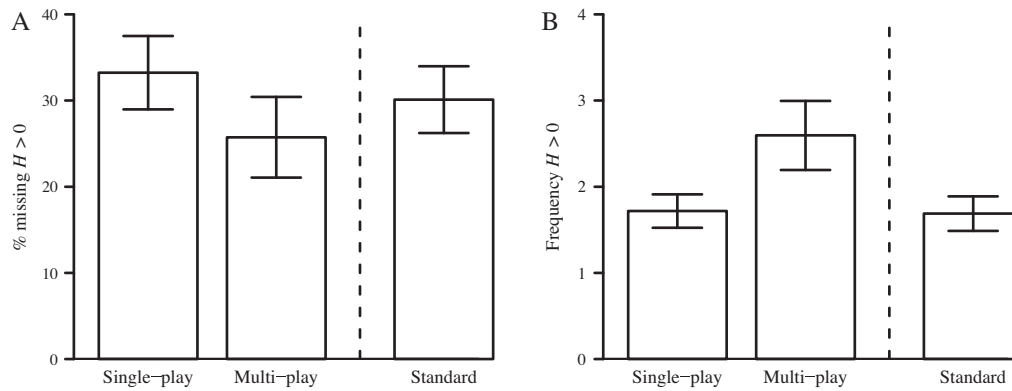
Fig. 3. Percentage of participants who missed (A) and frequency of encountering (B) the rare positive outcome in option *H*, separately for the three conditions: single-play, multi-play, and standard sampling paradigm. Analysis based on the 12 decision problems involving a rare positive event in option *H*. Error bars represent the standard error of the mean.

paradigm as a one-shot decision, notwithstanding the opportunity to aggregate choices across the sequence of decision problems in the experiment (Read, Loewenstein, & Rabin, 1999).

However, there is an important caveat to this interpretation. As spelled out before, differences in choices between single- and multi-play conditions could also be due to the different information people experienced. One and the same option, once filtered through experience, can take on many different 'phenotypes.' It is possible that respondents in both conditions maximize the same quantity, but that the quantity is the *experienced* mean reward, rather than the expected value. In other words, all the difference in choice might reside in the difference in the sampled information (see Figs. 2–4) and thus in the experienced phenotype rather than in different proclivities to maximize. In order to test this possibility, we calculated the proportion of choices of the option with the higher experienced mean (in those cases where participants experienced the rare event; 71% of cases) while accounting for the observed difference in means. Fig. 5B shows the results. Individuals in the multi-play condition continued to be much more likely to choose options with the higher experienced mean than were individuals in the single-play condition ($z = 3.3$, $p < .001$, $OR = 2.67$).[5] This effect was not reduced by the inclusion of the actual difference in means between the options ($z = 2.8$, $p = 0.01$, $OR = 2.68$). This suggests that differences in choices are not a mere function of sampling error (see Hertwig & Erev, 2009). Moreover, in further support of the notion that participants tend to view the standard paradigm as a short-run scenario, the standard condition did not differ from the single-play condition ($z = 0.76$, $p = 0.44$), but did differ from the multi-play condition ($z = 2.54$, $p = 0.01$).

Fig. 4. Observed probability minus true probability of the rare positive outcome in option *H*, separately for the three conditions: single-play, multi-play, and standard sampling paradigm. Analysis based on the 12 decision problems involving a rare positive event in option *H*. Relative frequencies are displayed in comparison to the expectation, i.e., the respective true probability. Error bars represent the standard error of the mean.

3.4. The role of working memory capacity

Based on a previously observed association between search and working memory (Rakow et al., 2008), we hypothesized that attentional control may serve as plausible explanation for the dependency between search and choice. To test this relationship, we measured participants' operation span and evaluated its association with sample size, switch rate, choices of the higher expected value option *H*, and choices of the option with the higher experienced mean. Table 2 shows the results of independent mixed effects analysis predicting these variables by operation span score. None of the effects were significant, suggesting that—if at all—attentional control plays a limited role in explaining search and choice in the sampling paradigm.

4. General discussion

Following up on Samuelson's anecdotal observation (1963), we investigated the suggestion that people making decisions from experience may choose differently when playing a gamble once versus multiple times (Lopes, 1996; Wedell, 2011). Using the sampling paradigm, we found differences in both choice and information search between single- and multi-play conditions. In the multi-play condition, individuals sampled more and were more

---

[5] Although choices in the single-play condition were much less likely to maximize expected value (5A) and experienced mean (5B) than were those in the multi-play condition, people still chose the higher mean option in 58% (versus 76%) of cases. This is of course not perfectly compatible with the notion that people in the single-play option tended to maximize the chances of coming out ahead (i.e., the median reward). However, people in the single-play condition also chose the option with the higher experienced chance of coming out ahead in 56% of cases (in 23% of cases, the options with the higher experienced mean and the higher chance of coming out ahead were identical). In contrast, people in the multi-play condition did so only in 44% of cases. One interpretation of this finding is that people in the single-play condition find themselves halfway between the two goals, with some betting on the rare but attractive gain, and others trying to come out ahead. In the multi-play environment, in contrast, the predominant course of action is maximization of the experienced mean.
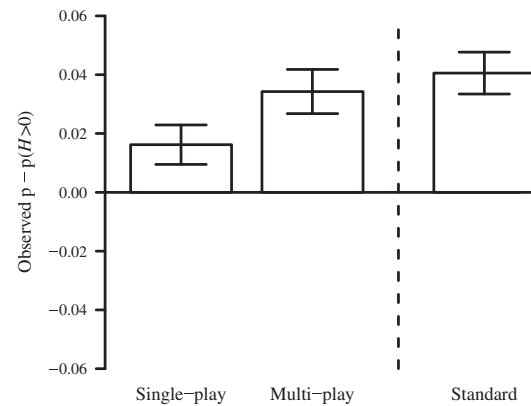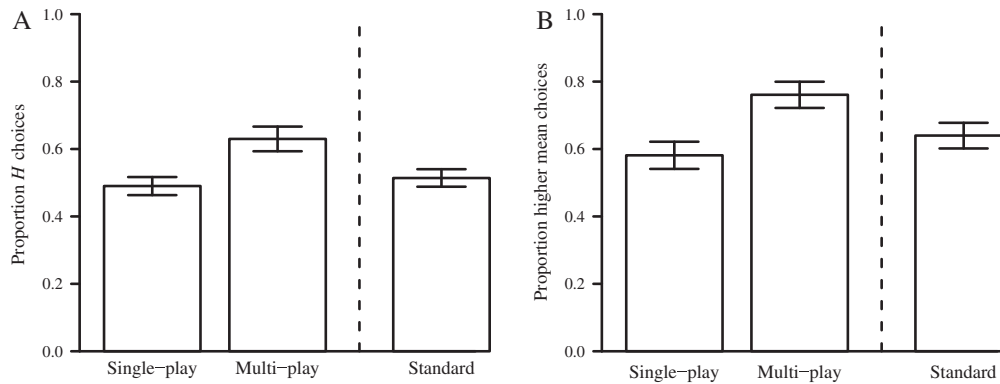
Fig. 5. Proportion of choices of the higher expected value option *H* (A) and choices consistent with the higher experienced sample mean (B), separately for the three conditions: single-play, multi-play, and standard sampling paradigm.

Table 2
Mixed-effects regression of search and choice on operation span.

	Sample size	Switch rate	*H* choices	Higher mean choices
Single-play	*b* = 3.14, *p* = 0.14	*b* = 0.03, *p* = 0.2	*b* = 0.01, *p* = 0.6	*b* = 0.02, *p* = 0.66
Multi-play	*b* = 2.65, *p* = 0.52	*b* = 0.01, *p* = 0.82	*b* = 0.05, *p* = 0.19	*b* = 0.04, *p* = 0.29
Standard	*b* = 0.25, *p* = 0.91	*b* = 0.03, *p* = 0.32	*b* = 0.02, *p* = 0.41	*b* = 0.02, *p* = 0.6

Note: Estimates (*b*) correspond to the change in the respective variable given a change of one standard deviation in the operation span score. Higher mean choices include only those choices where the rare event was observed.

likely to choose options with the higher expected values than did participants in the single-play condition. These differences were not mediated by experiencing different choice environments (and sampling error; see Fox & Hadar, 2006; Hertwig et al., 2004). Instead, they appear to stem from changes in decision strategy that were foreshadowed by changes in information search, a result consistent with the idea of adaptive information search as proposed by Hills and Hertwig (2010).

4.1. Implications for single- and multi-play choices

Our results add to the descriptive debate on single- and multi-play choices. The normative debate—i.e., whether people should have stable preferences across single- and multi-play situations—has cooled off. Single- and multi-play situations entail different payoff distributions. A change in preference can thus be consistent with 'classical' models of rationality (Aloysius, 2007). However, the discussion of how to best conceptualize the psychological processes involved in single- and multi-play situations is still ongoing. The first of two major positions is exemplified by Lopes' security potential and aspiration theory (SP/A; Lopes & Oden, 1999; see also Payne, 2005; Wang & Johnson, 2012) and proposes that (at least) two separate processes are executed in sequence: First, prospects are qualitatively compared against some aspiration level. When, and only when, the aspiration level is satisfied, the individual engages in a second, more systematic valuation of the prospect. If this is the case, the evaluation process of single- and multi-play situations could differ markedly, because multi-play prospects are more likely to surpass the aspiration level and trigger a systematic valuation than are single-play prospects (Wedell, 2011). The second position, exemplified by cumulative prospect theory (CPT; Tversky & Kahneman, 1992), proposes that just one process operates for single- and multi-play prospects

alike: Outcomes and probabilities of a prospect, once acquired and before being integrated into a single utility value, undergo non-linear transformations that allow single-play and multi-play versions of the same prospect to yield different utilities. Without making additional assumptions, this position implies identical evaluation processes for single-play and multi-play scenarios. Previous investigations using fully described single- and multi-play prospects have found evidence for differences in the behavioral patterns of choice and information acquisition (Joag, Mowen, & Gentry, 1990; Su et al., 2013), as well as in post-hoc verbal reports (Wedell & Böckenholt, 1990). Our investigation using decisions from experience adds to this debate by showing that the amount of information sampled prior to choice and the resulting experience varies in response to single- and multi-play instructions. Consistent with theoretical accounts assuming multiple processes, our findings suggests that the valuation process differs between single- and multi-play choices.

4.2. Implications for search and choice in decisions from experience

Our findings also provide new insights into the psychology underlying the standard sampling paradigm often used in recent research on decisions from experience. Behavior in this paradigm most resembled that observed in the single-play environment. This finding may be somewhat surprising, given that the standard condition, like the multi-play condition, offered multiple draws—one draw for each of the 16 choices. Participants in the standard condition could thus also aggregate the risk by bracketing the choices together (see Read et al., 1999). The results, however, suggest that participants evaluated each choice individually. This is consistent with previous research using decisions from description showing that people usually tend to segregate prospects when the cumulative nature of multiple prospects is not made apparent (Redelmeier & Tversky, 1992). Thus, one explanation may be that participants did not realize that they could aggregate the risk across their multiple choices (see DeKay & Kim, 2005, for the role of perceived fungibility in multi-play choices).

The finding that choice and search behavior in the standard condition resembled behavior in the single-play condition is of particular relevance to the discussion of as-if patterns of underweighting of rare events, and to what has been termed the description–experience gap (Hertwig & Erev, 2009). In contrast to the regularity of low-probability events tending to be overweighted in decisions from description (Hertwig & Erev, 2009; Tversky & Kahneman, 1992), it has been inferred from choices that rare events tend to receive *less* weight than they deserve in decisions from experience (Camilleri & Newell, 2011; Hertwig et al., 2004). Studies using yoked experiences and direct probability judgments

by participants have revealed that such underweighting also occurs when the subjective representation of the prospect accurately reflects its objective properties (Camilleri & Newell, 2009; Ungemach, Chater, & Stewart, 2009). The use of short-run strategies that maximize the chance of coming out ahead by inherently ignoring rare events may offer a new angle from which to address the persistent puzzle of why such underweighting occurs.

Further, our results suggest that different components of explorative behavior may be under different control processes. The lack of a difference in switch rate across all conditions (and across domains, see Lejarraga, Pachur, Frey, & Hertwig, 2015) appears to indicate that this property is not part of a participant's top-down aspiration level and associated decision strategy. Rather, switching may be under the control of more implicit processes, such as those associated with working memory (Hills & Pachur, 2012; Rakow et al., 2008; see also Hills, Todd, & Goldstone, 2010). Yet the complete independence between our measure of working memory and both search and choice variables throws into question working memory's potential role as a stopping rule as well as its previously suggested role (Rakow et al., 2008) as a common cause for search and choice policy (for similar findings, see Wulff, Hills, & Hertwig, 2014). One reason for the absence of a link may lie in the high complexity and difficulty of the operation span task relative to the simple digit span task used by Rakow et al. (2008). Yet, it could also mean that models for decision from experience that simplify information integration and choice, thereby taxing working memory less, may be good candidate models for experienced-based decision making. One particular class, often labeled as associative learning models (see Hertwig, 2015; Sutton & Barto, 1998), assumes a continuous updating of a single utility value per option and thus requires the storage of far less information. Similarly, the maximization of the experienced mean could be achieved by recruiting a simple strategy such as the natural mean heuristic (Hertwig & Pleskac, 2010). Future investigations should explore how short- and long-run aspirations may be captured in choice models based on an associative learning mechanism. In this context, attention should also be paid to the diversity of executive functions (Miyake et al., 2000), of which only some may be involved in experienced-based risky choice (see also Frey, Mata, & Hertwig, 2015).

Our results also shed light on an aspect of information search in decisions from experience that has been overlooked, namely, optional stopping. Deliberations into the statistical effects of small samples have often assumed that search is randomly terminated or is terminated once a preplanned size is reached (Fox & Hadar, 2006; Gonzalez & Dutt, 2011; Hertwig & Pleskac, 2010; Hertwig et al., 2004). Our investigation suggests that termination of search may also be subject to strategic concerns (for related findings, see Lejarraga, Hertwig, & Gonzalez, 2012; Phillips, Hertwig, Kareev, & Avrahami, 2014).

4.3. The problem of inferring risk preference in decisions from experience

Last but not least, our findings highlight a thorny inference problem concerning risk preferences in experienced-based choices. Following the revealed preference approach (Samuelson, 1938), researchers often infer an individual's preference directly from her choices in described and stable choice environments (decisions from description). Decision problems employing stated probabilities can easily be tailored to make different risk preferences discernible (e.g., Holt & Laury, 2002). In decisions from experience, though, inferring risk preference from choice is much more problematic. Due to variable sample sizes and the random composition of samples, it is often the case that no two individuals face identical decision problems. For this reason, it has been convincingly argued

that the individual-specific experienced choice environment—and not the objective choice parameters (outcomes and probabilities)—is the appropriate foundation for inferring the individual's risk preference (Fox & Hadar, 2006). Our findings, however, suggest that even inferring preferences contingent on experience is problematic. For illustration, consider a person with a *strong* preference for coming out ahead in the short term (in a single-play environment). She takes relatively few samples in each problem. Consequently, she may be faced with a decision between an apparently safe and modest positive outcome and an apparently safe zero outcome. This 'trivial' (dominated) choice reveals little about this person's risk preference. Now, consider a person with a *weak* preference for coming out ahead in the short term. She may sample a bit more than the first person, and even encounter the rare but attractive outcome. Because of her preference, she decides against the option offering this dicey but attractive outcome. This person will be 'revealed' to be risk averse, whereas the other appears, if anything, to be risk neutral. Of course, this is a constructed example (ignoring, among other factors, the role of optional stopping), but it illustrates a simple but consequential point: In environments that people experience and 'construct' through active sampling, inferences from choice to preference are problematic because the experienced environment can arise from the preference or aspiration level itself. Depending on which environment emerges, choices may or may not be informative about the underlying preferences or aspirations (see also Denrell, 2007). The good news, however, is that decisions from experience paradigms offer an *observable* psychological dimension that appears to afford researchers another window onto preferences or aspirations: the appetite for information.

Appendix A

See Table A1.

Table A1
Additional decision problems.

Problem	H	L
1′	1 with p = .75	0 with certainty
2′	1 with certainty	0 with certainty
3′	3 with p = .75	9 with p = .10
4′	2 with certainty	7 with p = .10

Note: H = option with the higher expected value (as calculated by probability $\times$ monetary value); *L* = option with the lower expected value.

References

Aloysius, J. A. (2007). Decision making in the short and long run: Repeated gambles and rationality. *British Journal of Mathematical and Statistical Psychology, 60,* 61–69.

Ashby, N. J., & Rakow, T. (2014). Forgetting the past: Individual differences in recency in subjective valuations from experience. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 40*(4), 1153–1162.

Berger, J. O., & Berry, D. A. (1988). Statistical analysis and the illusion of objectivity. *American Scientist, 76,* 159–165.

Camilleri, A. R., & Newell, B. R. (2009). The role of representation in experience-based choice. *Judgment and Decision Making, 4*(7), 518–529.

Camilleri, A. R., & Newell, B. R. (2011). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review, 18*(2), 377–384.

Camilleri, A. R., & Newell, B. R. (2013). The long and short of it: Closing the description–experience "gap" by taking the long-run view. *Cognition, 126,* 54–71.

Caraco, T. (1980). On foraging time allocation in a stochastic environment. *Ecology, 61,* 119–128.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.

Conway, A. R. A., Cowan, N., & Bunting, M. F. (2001). The cocktail party phenomenon revisited: The importance of working memory capacity. *Psychonomic Bulletin & Review, 8*, 331–335.

DeKay, M. L., & Kim, T. G. (2005). When things don't add up: The role of perceived fungibility in repeated-play decisions. *Psychological Science, 16*(9), 667–672.

Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review, 114*(1), 177–187.

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

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

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review, 118*(4), 523–551.

Hertwig, R. (2015). Decisions from experience. In G. Keren, G. Wu (Eds.), Blackwell handbook of decision making. Oxford, UK: Blackwell.

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.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences, 13*, 517–523.

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

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

Hills, T. T., & Pachur, T. (2012). Dynamic search and working memory in social recall. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 38*(1), 218–228.

Hills, T. T., Todd, P. M., & Goldstone, R. L. (2010). The central executive as a search process: Exploration and exploitation in generalized cognitive search processes. *Journal of Experimental Psychology: General, 139*, 590–609.

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *The American Economic Review, 92*(5), 1644–1655.

Houston, A. I., & McNamara, J. M. (1999). *Models of adaptive behavior: An approach based on state*. Cambridge, UK: Cambridge University Press.

Joag, S. G., Mowen, J. C., & Gentry, J. W. (1990). Risk perception in a simulated industrial purchasing task: The effects of single versus multi-play decisions. *Journal of Behavioral Decision Making, 3*(2), 91–108.

Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science, 39*(1), 17–31.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society, 47*(2), 263–291.

Kane, M. J., & Engle, R. W. (2000). Working memory capacity, proactive interference, and divided attention: Limits on long-term memory retrieval. *Journal of Experimental Psychology: Learning, Memory, Cognition, 26*, 333–358.

Keren, G. (1991). Additional tests of utility theory under unique and repeated conditions. *Journal of Behavioral Decision Making, 4*, 297–304.

Keren, G., & Wagenaar, W. A. (1987). Violation of utility theory in unique and repeated gambles. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 13*, 387–391.

Koop, G. J., & Johnson, J. G. (2012). The use of multiple reference points in risky decision making. *Journal of Behavioral Decision Making, 25*, 49–62.

Langer, T., & Weber, M. (2001). Prospect theory, mental accounting, and differences in aggregated and segregated evaluation of lottery portfolios. *Management Science, 47*(5), 716–733.

Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition, 124*(3), 334–342.

Lejarraga, T., Pachur, T., Frey, R., & Hertwig, R. (2015). Decisions from experience: From monetary to medical gambles. *Journal of Behavioral Decision Making*.

Liu, H. H., & Colman, A. M. (2009). Ambiguity aversion in the long run: Repeated decisions under risk and uncertainty. *Journal of Economic Psychology, 20*, 277–284.

Lopes, L. L. (1981). Decision making in the short run. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 377–385.

Lopes, L. L. (1996). When time is of the essence: Averaging, aspiration, and the short run. *Organizational Behavior and Human Decision Processes, 65*, 179–189.

Lopes, L. L., & Oden, G. C. (1999). The role of aspiration level in risky choice: A comparison of cumulative prospect theory and SP/A theory. *Journal of Mathematical Psychology, 43*, 286–313.

Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology, 41*(1), 49–100.

Montgomery, H., & Adelbratt, T. (1982). Gambling decisions and information about expected value. *Organizational Behavior and Human Performance, 29*, 39–57.

Payne, J. W. (2005). It is whether you win or lose: The importance of the overall probabilities of winning or losing in risky choice. *Journal of Risk and Uncertainty, 30*(1), 5–19.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge, UK: Cambridge University Press.

Peköz, E. A. (2002). Samuelson's fallacy of large numbers and optional stopping. *Journal of Risk and Uncertainty, 69*, 1–7.

Phillips, N. D., Hertwig, R., Kareev, Y., & Avrahami, J. (2014). Rivals in the dark: How competition affects information search and choices. *Cognition, 133*, 104–119.

R Development Core Team (2008). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria.

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

Read, D., Loewenstein, G., & Rabin, M. (1999). Choice bracketing. *Journal of Risk and Uncertainty, 19*, 171–197.

Redelmeier, D. A., & Tversky, A. (1992). On the framing of multiple prospects. *Psychological Science, 3*(3), 191–193.

Samuelson, P. A. (1938). A note on the pure theory of consumer behavior. *Economica, 5*(17), 61–71.

Samuelson, P. A. (1963). Risk and uncertainty: a fallacy of large numbers. *Scientia, 98*, 108–113.

Stephens, D. W. (1981). The logic of risk-sensitive foraging preferences. *Animal Behavior, 29*, 628–629.

Stephens, D. W. (2001). The adaptive value of preference for immediacy: When shortsighted rules have farsighted consequences. *Behavioral Ecology, 12*(3), 330–339.

Su, Y., Rao, L. L., Sun, H. Y., Du, X. L., Li, X., & Li, S. (2013). Is making a risky choice based on a weighting and adding process? An eye-tracking investigation. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 39*(6), 1765–1780.

Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning*. Cambridge, MA: MIT Press.

Tversky, A., & Bar-Hillel, M. (1983). Risk: The long and short. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*(4), 713–717.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5*(4), 297–323.

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science, 20*, 473–479.

Unsworth, N., Heitz, R. P., Schrock, J. C., & Engle, R. W. (2005). An automated version of the operation span task. *Behavior Research Methods, 37*(3), 498–505.

Wang, X. T., & Johnson, J. G. (2012). A tri-reference point theory of decision making under risk. *Journal of Experimental Psychology: General, 141*(4), 743–756.

Wedell, D. H. (2011). Evaluations of single and repeated-play gambles. In J. J. Cochran (Ed.), *Wiley encyclopedia of operations research and management science*. <http://onlinelibrary.wiley.com/doi/10.1002/9780470400531.eorms0670/abstract>.

Wedell, D. H., & Böckenholt, U. (1990). Moderation of preference reversals in the long run. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 429–438.

Wulff, D. U., Hills, T. T., & Hertwig, R. (2014). Online product reviews and the description–experience gap. *Journal of Behavioral Decision Making*. Advance online publication. doi: http://dx.doi.org/10.1002/bdm.1841.

Wulff, D. U., Pachur, T. (in press). Modeling valuations from experience: A commentary on Ashby and Rakow (2014). *Journal of Experimental Psychology: Learning, Memory and Cognition*.