Adaptive Information Search and Decision Making over Single and Repeated Plays

Dirk U. Wulff (Dirk.Wulff@unibas.ch)
Department of Psychology, Missionstrasse 64A
4055 Basel, Switzerland

Thomas T. Hills (T.T.Hills@warwick.co.uk)
Department of Psychology, University of Warwick
Coventry CV4 7AL, UK

Ralph Hertwig (Ralph.Hertwig@unibas.ch)
Department of Psychology, Missionstrasse 64A
4055 Basel, Switzerland

Abstract

For over 50 years expected value and expected utility theory has been challenged by behavioral findings in repeated and single plays of risky gambles. The inherent long-term nature of these models has been found to be at odds with preferences indicating short-term maximization in single play situations. With the present study we provide further evidence on the distinction between long-term and short-term oriented behavior. Evaluating experienced-based decisions over repeated and single play situations we demonstrate that both choice preferences and search behavior change in response to long and short-term framing. This suggests different cognitive approaches for single and repeated play situations, with single decisions often favoring risk-aversion and therefore the underweighting of rare events. These findings are in line with alternative models of risky choice as for example proposed by Lopes (1996) and also the literature on state-dependent foraging.

Keywords: Decisions from experience, information sampling; risky choice, single-play and repeated-play.

Introduction

Over lunch nobel-laureate Paul Samuelson offered his colleague the following bet: Win 200$ with a 50% chance or lose 100$. To his surprise the colleague rejected the offer but mentioned that he would agree to a series of a hundred such bets. Samuelson (1963) considered the choice of his colleague irrational. He later showed how this pattern of choice is incompatible with expected-utility maximization. A series of bets, which individually are all inacceptable, should not be accepted. But is the behavior of Samuelson's colleague irrational?

The prevailing answer after 50 years of scientific debate seems to be unanimously no. However, very different theoretical explanations have been invoked. On one hand, it has been demonstrated that the behavior of Samuelson's colleague can indeed be captured by models of expected utility theory when the decision is made over the aggregated outcome of repeated plays (Aloysius, 2007). On the other hand, "distinct process differences between the situations" (Wedell, 2011) have been emphasized. Particularly Lola Lopes argued in favor of a two-stage decision process in which the qualitative feature of "coming-out-ahead" plays an important role (Lopes, 1996; Lopes & Oden, 1999). Formally Lopes assumes next to decumulative weighting a second process that maximizes the probability of achieving an aspiration level.

Along the lines of Lola Lopes' explanation, Hills and Hertwig (2010) recently provided evidence for a decision strategy that focuses on winning most of the time. Hills and Hertwig used an alternative paradigm of risky choice – decisions from experience – in which the information about options is acquired in a prior inconsequential sampling phase. They found that the choice behavior of people who sampled less and evaluated options in shorter intervals – were best described by a "round-wise" decision strategy. Instead of aiming at the overall higher mean of both options, this strategy compares the options over all sampling transitions and tallies the wins. Consequently, this strategy favors the option with a higher probability of coming out ahead—i.e., the option that wins most of the time. The results of Hills and Hertwig (2010) indicate that distinct processes are indeed at work when people aim for either long-run expected value maximization or for short-run maximization that focuses on coming-out ahead.

The distinction between such strategies is prevalent also outside the field of human decision-making. For example, risk sensitivity and state-dependent foraging explain why a bird that needs food regularly to avoid starvation should concern itself with short-term outcomes and not seek for the better long-term average, which may not come in time (e.g., Caraco, 1980; Houston & McNamara, 1999; Stephens, 1981). In the present study we build on these ideas by examining information search and decision making in single and repeated play situations.

Behavioral Findings Under Repeated Plays

Behavioral studies have shown that in repeated-play situations—where people get to play the same gamble multiple times—people are more likely to act according to the principles of expected utility maximization (Wedell, 2011). Generally, under repeated plays, the preference for the option with the higher expected value is substantially increased (e.g. Montgomery & Adelbratt, 1982). In addition, repeated plays reduce a number of well-known decision anomalies including possibility and certainty effects (Keren & Wagenaar, 1987; Keren, 1991), violations of procedural invariance (Wedell & Böckenholt, 1990) and ambiguity aversion (Liu & Coleman, 2009).

The effect of repeated plays on choice preference has also been established in applied settings. Redelmeier and Tversky (1990, 1992) found that physicians make differential decisions for individual versus aggregated instances. Bernartzi and Thaler (1990) have also demonstrated the relevance of the single repeated play distinction for investment decisions.

Present Study

The goal of the present study is to evaluate the effect of single and repeated plays in the context of decisions from experience. By incorporating an active sampling phase decisions from experience account for the lack of full information about the available options in real life; a quality that is neglected in the commonly used paradigm of decisions from description. Based on the findings of Hills and Hertwig (2010) we expect the contrast between repeated and single play to impact pre-choice information search in addition to preferences. Specifically we expect people in single play as compared to repeated play situations to draw fewer samples and have shorter evaluation intervals for the available options. We also expect a higher preference in single play situations for options that have a higher probability of winning most of the time irrespective of expected. For repeated plays we expect the opposite pattern. Such findings would indicate an overall differential cognitive approach to single and repeated plays of the same gamble and speak in favor of choice models such as Lola Lopes' two-stage account.

Additionally, we are interested in the degree to which the pattern of results is dependent on the certainty of the available options. In much of the risky choice literature a risky high outcome option is paired with an option that offers a smaller but certain outcome. Thereby, the quality of "coming-out-ahead" is often confounded with certainty. However, given the prior assertions we would expect the pattern of results to be largely independent over cases where the secure option is certain (or only relatively secure).

Finally, we want to evaluate how the standard application of decisions from experience relates to situations where the single and repeated-play character is made salient. Decisions from experience have recently received much attention as they depart from decisions that are made based on full description of the options (see Hertwig & Erev, 2009). In this context it will be very telling to observe if the overall pattern of results in the paradigm used most often in the literature (following Hertwig et al., 2004) resembles more a single or repeated play instantiation.

Method

Participants We collected data from 124 participants. The mean age of the sample was about 24.2 years, 85% were students of the University of Basel. Participants were rewarded either by course credit or a fixed payment of about 13$. Every participant also received a performance-based bonus as a result of his or her choices.

Materials A set of 12 target problems was created (see table 1). Each problem required a choice between two gambles. Every gamble was comprised of two outcomes – one positive outcome and zero. One gamble posed a relatively secure ($p>.7$) or certain chance ($p=1$) to win a positive outcome. The other gamble was substantially more risky ($p<=.15$). This riskier option was always superior in expected value (1.5 to 1.8 times as high). We refer to these two options are referred to as the low (L) and high (H) expected value options, respectively. For control purposes, we included two problems where the high outcome options had the lower expected value (*C1*,

Table1: Study problems

Problem	H	L
1	92 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
C1	0 with certainty	1 with p=.75
C2	0 with certainty	1 with certainty
C3	9 with p=.10	3 with p=.75
C4	7 with p=10	2 with certainty

C2) and two problems where the supposedly high outcome option was actually a sure event of zero (*C3*, *C4*). Those problems were created to exclude participants showing unsystematic or extreme risk seeking behavior.

Procedure The experiment was entirely computer based. Participants were given verbal and visual instructions that explicitly explained the payoff modality according to their assigned experimental condition. Participants then made choices for three practice trials. On every problem the participants were able to sample from the two options as much and in whatever fashion they liked. They were instructed to proceed to the decision whenever they felt confident enough. Finally the participants made their decisions on all 16 problems (12 target and 4 control problems) in random order.

Payoff Manipulation The single and repeated play character was induced between-subjects through different payoff modalities. The payoff in the repeated play condition was determined by one hundred draws from one of the participant's chosen options. In the single play condition the payoff, one random sample was taken from one the participant's chosen options and then this value was multiplied by one hundred. A third condition (neutral) corresponded to the payoff modality usually applied in the literature (Hertwig et al., 2004), with the payoff equivalent to one random draw from each option chosen by the participant. As the payoff in the neutral condition is based on only 16 draws all presented outcomes were multiplied by a fixed factor to equal the expected values in the other two conditions.

Results

We applied two criteria to exclude participants showing extreme risk seeking or unsystematic behavior. First, we excluded participants that neither preferred the higher mean nor the higher median in the control problems *C1* and *C2* (see table 1). Second, we also excluded participants that did not sample at least once from every option. As a result the following analyses are based on 82 out of 124 participants. Inspection of control problems *C3* and *C4* reveals that this restriction very efficiently reduces zero EV choices, which can be regarded as an indicator for either extreme risk seeking or not very systematic behavior. The excluded participants chose the zero EV option in about 27% of the cases, whereas the included participants chose did not choose this option at all. Thus, the remaining participants showed neither extreme risk-seeking nor unsystematic behavior.

The main focus of this study is on the contrast between the single and repeated play conditions. Statistical tests are therefore mainly reported for this contrast. Additionally, comparisons for the neutral condition are reported to see where usual experimental instantiations of this paradigm fall in the context of salient single and repeated play situations.

Figure 2A illustrates choice patterns over the experimental conditions. In support of our predictions, H preferences increased for the repeated-play as compared to single-play situation ($t_{42.4}$=3.58, p<.001). H preference in the neutral condition resembled the pattern in the single-play condition ($t_{53.1}$=0.54, p=.594), but there was a difference for H choices between the neutral and repeated play conditions ($t_{43.8}$=-3.05, p=.004). This effect is not affected by the inclusion of the certainty of the L option in the prediction of H choices (X^2_2=.51, p=.776). Because none of the following analyses were influenced by a comparison of certain and uncertain gambles, the following results are collapsed over both problem types.

Figure 2B shows the information search behavior of the participants. Descriptively, it matches our predictions. The repeated-play condition elicited higher total sample sizes as well as longer evaluation intervals for the individual options (samples per transition: average number of samples taken per transition in the sampling sequence). However, only the contrast for total number of samples reaches statistical significance (total: $t_{32.9}$=2.28, p=.029; per transition: $t_{31.8}$=1.58, p=.124). Again the neutral condition closely resembled the single-play condition in both sampling variables (total: $t_{49.1}$=.52, p=.609; per transition: $t_{52.1}$=.85, p=.40), but was at least in terms of total sample size marginally different from the repeated-play situation (total: $t_{37.9}$=-1.85, p=.072; per transition: t_{34}=-1.07, p=.29).

So far we have shown that the payoff structure independently affected choice and information search in the predicted direction. However, we generally predicted that information search and choice behavior were likewise affected by the induction of a repeated play situation in that it was expected to elicit both a higher preference for the long run winner as well as higher total and per transition sample sizes. To demonstrate this, two additional analyses have to be carried out. First, it has to be shown that the contrast for H preference is not entirely caused by a sampling bias. Hertwig et al. (2004) showed how small sample sizes can systematically obscure rare outcomes in decisions from experience. When rare outcomes are obscured, the option that secures winning most of the time can also appear as the option with the higher long run expectation. Second, an association between the choice preferences and the sampling behavior needs to be established. Therefore we tested if this pattern of results holds when the qualities of winning most of the time – represented by the higher median – and higher long run expectation – represented by the higher mean – do not coincide in respect to actually observed outcomes. Then we evaluate if the preference for the higher mean option is associated with the sampling behavior.

Figure 3 shows the proportions of choices that were consistent with the higher experienced mean. The separate lines distinguish cases where the experienced median and mean predict the same (dotted lines) or different (solid lines) choices. Focusing on the cases where the mean and median do not fall together, it is apparent that participants in the repeated play situation opted more often for the option with the higher experienced mean as compared to the single play situation ($t_{49.1}$=2.72, p=.001). Thus, the differences in choices do appear to be associated with different decision strategies and not simply differences in observed outcomes. The neutral condition again matches the single ($t_{48.3}$=.12, p=.90) but not the repeated play condition ($t_{48.8}$=-2.26, p=.028). Overall, the observed differences are all the more convincing given the pattern of results for cases where the experienced mean and median fall together. Independent of the payoff condition we observed stong preferences for the option with the higher experienced mean pointing towards very systematic

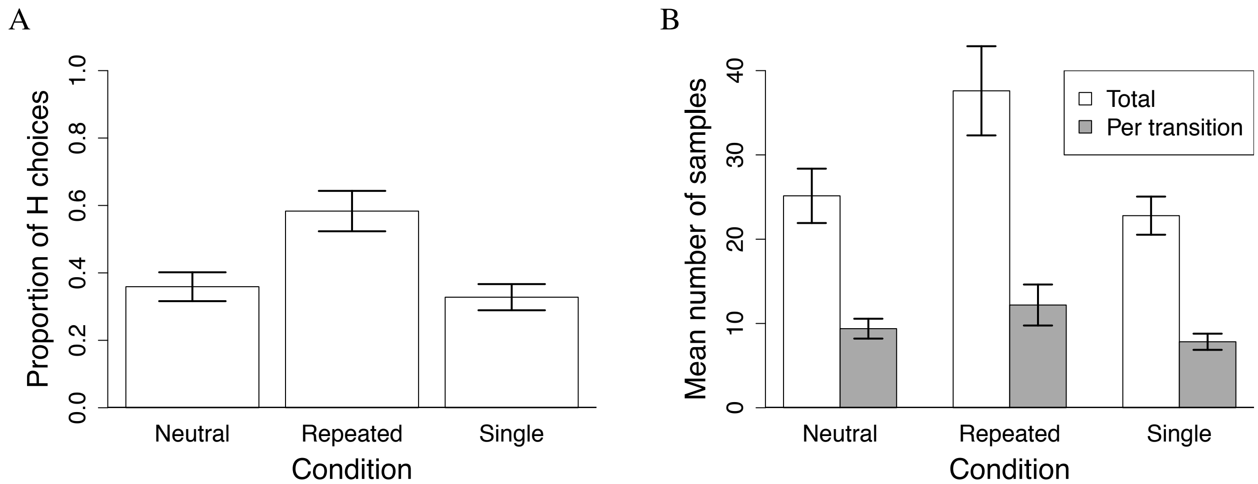

Figure 2: (A) Proportion of H choices as a function of the experimental condition. (B) Mean number samples in total and per transition in the sampling sequence as a function of the condition. Error bars represent standard error of the mean.

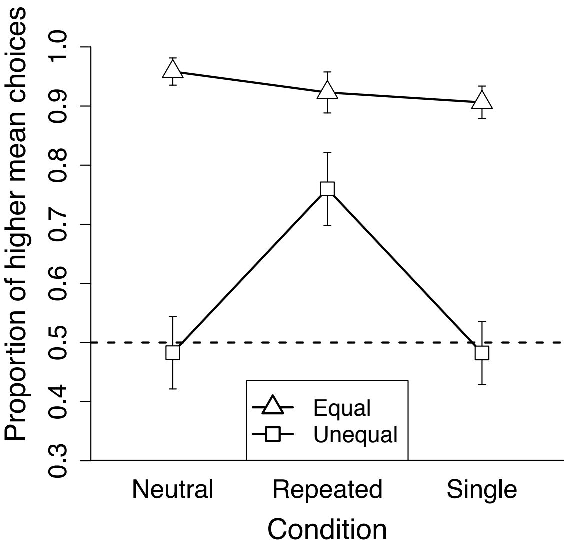

Figure 3: Higher mean choices as a function oft he experimental condition.

behavior.

Now, is the shift in strategy use for the repeated play condition associated with changes in the sampling behavior? In line with our predictions we found this to be the case. Regressing higher mean choices on sample size and the group contrast of single and repeated play situations yields a significant interaction (t_{52}=-2.73, p=.009), while the main effect of the contrast vanishes (t_{52}=.54, p=.59). The same pattern was found for samples per transition (interaction: t_{52}=-2.36, p=.022; contrast: t_{52}=-0.33, p=.746). It was not possible to determine which of these sampling variables were more strongly associated with the strategy use because they were highly correlated (r = .83). In sum, we thus were able to show, that the repeated play condition influenced both choice preference and sampling behavior.

Discussion

The primary aim of our study was to use experience-based decisions to investigate information search and decision making over repeated and single plays. Based on prior findings of Hills & Hertwig (2010) we predicted that the situations of repeated and single plays would elicit information search behavior that has been found to be associated with short and long-term maximization, respectively. Our results confirm this prediction. Participants in the situation of repeated plays – particularly those that sampled more avidly and switched less frequently between the options – exhibited higher preferences for the superior option in the long run. These results did not appear to be mediated by a simple sampling bias (Hertwig et al., 2004; Fox & Hadar, 2006). Rather, they were associated with different decision strategies.

We believe that our findings provide new evidence for the theoretical debate on single and repeated-play decisions. So far behavioral studies that separated these situations have focused entirely on decision outcomes. We extend this literature by showing that differences can also be demonstrated on more process-oriented measures such as information search. Our results imply that possibly more is changing over the situations of single and repeated play than the weighting of probabilities and outcomes (Aloysius, 2007), as suggested by unitary theoretical accounts, e.g. prospect theory (Kahneman & Tversky, 1979). We suggest that an overall different cognitive approach towards repeated plays is applied. These results appear to be consistent with theoretical accounts such as Lola Lopes' (1996, 1999) two-stage model.

Our findings also hold important insights for the literature on experienced-based decision making. On all of our measures, the standard payoff procedure resulted in behavior that resembled the condition where the single-play character was strongly emphasized. This might not appear surprising as in both cases only single draws were taken out of the chosen option. However, in the standard realization all decisions contribute to the overall payoff. Thus, the whole set of choices could be mentally combined (see Read, Loewenstein, & Rabin, 1999) and thereby the set of problems could more resemble a repeated play situation. In this context our findings suggest that under the standard paradigm of decisions from experience, all choices are evaluated independently. This finding might contribute in the explanation of the description-experience gap.

In conclusion, we find that single and repeated play situations impact more than just choice preferences. The way people search for information and specific decision strategies change as well. Our findings emphasize the distinction between long-term and short-term maximizing behavior in human decision-making - a distinction that is long established in the literature on animal foraging.

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.

Benartzi S., & Thaler, R. H. (1999). Risk aversion or myopia? Choices in repeated gambles and retirement investments. *Management Science, 45*, 364–381.

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

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.

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.

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.

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

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. (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.

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

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–292.

Keren G. (1991). Additional tests of utility theory under unique and repeated condition. *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.

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*, 191–193.

Redelmeier D. A., & Tversky, A. (1990). Discrepancy between medical decisions for individual patients and for groups. *New England Journal of Medicine, 322*, 1162–1164.

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.

Wedell, D. H. (2011). Evaluations of single and repeated-play gambles. *Wiley Encyclopedia of Operations research and Management Science.*

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.