Journal of Behavioral Decision Making
J. Behav. Dec. Making, **19**: 17–42 (2006)
Published online in Wiley InterScience
(www.interscience.wiley.com). **DOI**: 10.1002/bdm.512

When Mixed Options are Preferred in Multiple-Trial Decisions†

YUH-JIA CHEN[1] and JAMES E. CORTER[2]*
[1]*Middle Tennessee State University, USA*
[2]*Teachers College, Columbia University, USA*

ABSTRACT

We report three studies showing that in prospective multiple-trial decisions people often select a mix of sure and risky options over pure bundles of either option. Such a preference is not 'rational' because a mixed option cannot be the EV-maximizing choice. Experiment 1 confirmed a mixed-option preference for gains but not for losses. Showing a graph of the multiple-trial outcome distribution reduced but did not eliminate this effect, suggesting that it is not due purely to a failure to aggregate correctly over the multiple trials. Experiment 2 replicated the mixed option preference using a wider range of problems. Experiment 3 compared choices in the trinary choice conditions used in Experiments 1 and 2 with binary choices between *pairs* of the multiple-trial sure, mixed, and risky options. In the binary choice condition the mixed option was no longer the modal choice, suggesting that the strong mixed option preference found in the trinary choice conditions is mainly due to a compromise effect. However, the binary choice probabilities did show violations of strong stochastic transitivity in a pattern that suggested a slight bias toward the mixed option. Copyright © 2006 John Wiley & Sons, Ltd.

KEY WORDS decision-making under risk; multiple trials; rationality; compromise effect

INTRODUCTION

Some decisions in life are faced repeatedly. For example, each morning we might choose whether to get to work by bus or by car. Every year we might have the opportunity to direct a financial contribution into any of several different investments in an individual retirement account. Leading theories of decision-making under risk, including Expected Utility Theory (Luce & Raiffa, 1957; von Neuman & Morgenstern, 1947) and Prospect Theory (Kahneman & Tversky, 1979), were originally developed to explain how people behave when faced with a single trial of a decision problem. But evidence from decision-making research has accumulated

* Correspondence to: James E. Corter, 525 W. 120th St., rm. 453 Grace Dodge Hall, Teachers College, Columbia University, New York, NY 10027, USA. E-mail: jec34@columbia.edu
† Experiments 1 and 2 were conducted as part of the first author's dissertation research (Chen, 2001).

showing that people may behave differently when faced with multiple or repeated trials of a decision problem.

In this paper we investigate risky choice in multiple-trials situations. First, we draw a distinction between *multiple-trial* decision tasks and *repeated independent* decisions. In *multiple-trial* decisions, the decision maker (DM) is faced with a single prospective choice between strategies to be applied uniformly to some number of identical decision trials. For example, participants in a study might be offered a choice between ten plays of a risky option (e.g., a 75% chance of receiving \$20) or ten 'sure-thing' payments of a constant amount with lower EV (\$10). In contrast, the *repeated-trials* paradigm is one in which a decision maker makes repeated decisions or choices, one for each trial, usually experiencing the outcome or consequence of each decision after the trial. Note that it is also important to distinguish decision tasks in which the decision maker is provided with the relevant probabilities and payoffs ('description-based' tasks) from decision tasks in which these are initially unknown, but discovered through feedback experience over repeated trials (Barron & Erev, 2003; Hertwig, Barron, Weber, & Erev, 2004).

Analogs to these multiple-trials and repeated-trials situations arise in real life. For example, a day-trading investor might make repeated decisions every day as to whether or not to purchase or sell options in a particular stock. An investor who makes investment allocation decisions only once a quarter or once a year can be seen as making prospective multiple-trial decisions (if each day of market trading is viewed as a separate 'trial,' or opportunity to change one's allocation).

The present paper seeks to extend previous research on prospective multiple-trial decisions. Previous studies of prospective multiple-trial decision making (e.g., Keren & Wagenaar, 1987; Wedell & Böckenholt, 1990; Redelmeier & Tversky, 1992; Benartzi & Thaler, 1999) have examined the consistency of choices between single- and multiple-trial situations by offering participants prospective choices between bundles of risky or risk-free options. For example, in the single-play condition participants might be offered a choice between a sure gain of \$5 and a 50% chance to win \$15, while in the multiple-play condition they might be offered a choice between ten sure gains of \$5 and ten plays of the gamble that consists of a 50% chance to win \$15.

But we hypothesize that in this situation most participants would prefer *neither* the purely sure-thing nor the purely risky options. For a number of reasons (described below), we suspect that people will often prefer a *mix* of the risky and risk-free options in these prospective multiple-trial choices. Such a finding would be highly interesting, because preference for the mixed option is non-normative, at least from the point of view of expected value (since the expected value of any mixed option must lie between the expected values of the other two options).

One reason to suspect that decision makers might prefer mixed options in prospective multiple-trial decisions under risk is that when participants experience *repeated independent* trials of a risky decision problem, in which they are given 'feedback' after each trial in the form of the obtained outcome, they rarely select one alternative consistently (Barron & Erev, 2003; Busemeyer & Myung, 1992; Camerer & Ho 1998; Chen, 2001). Rather, they often alternate or mix their selections. One simple explanation for response variability is that people might have natural tendencies to exhibit a certain level of exploratory behavior, in the form of variability in their choices. In fact, the observation that repeated choices are often inconsistent motivated the development of stochastic models of choice in psychology (see e.g., Luce, 1959; Tversky, 1972).

Inconsistency of choices has also been documented in consumer decision research, where it is termed *variety-seeking* behavior. Interestingly, the results of several studies (Read & Loewenstein, 1995; Simonson, 1990) of prospective multiple-trials consumer purchase decisions suggest that preferences for mixed bundles of products (options) might even be *exaggerated* in multiple-trial prospective choice compared to single-trial choices made at the time of consumption, a phenomenon that Read and Loewenstein termed 'diversification bias.' It is possible that an analogous bias might exist for prospective multiple-trial decisions under risk.

Another reason to suspect that decision makers might exhibit a mixed-option preference in prospective multiple-trial financial decisions comes from evidence of social norms promoting a financial 'diversification heuristic' in investment decisions. It is well established in portfolio theory that there is (normative) value in diversification, i.e., in selecting a variety of financial instruments rather than a single one. Empirical data show that patterns of behavior echoing this normative theory are apparent in the behavior of even relatively naïve individual investors. Individual investors often choose to allocate their investments evenly among the available investment instruments, a strategy that Benartzi and Thaler (2001) termed 'naïve diversification.' Selecting a 50–50 mix of sure and risky gains in prospective multiple-trial choices can be viewed as analogous to this investment heuristic.

Thus, the main goal of the present work is to test for a possible mixed-option preference in prospective multiple-trial decisions under risk, and to investigate the reasons for any such finding. In the following sections, we first review previous studies of decision behavior in repeated-trials decision tasks and in prospective multiple-trial decision tasks, before turning to the present experiments.

Previous research on multiple-trial decisions

Samuelson (1963) was one of the first to point out that risky decisions may differ for single and multiple plays of a gamble. He told of a colleague who refused to accept a gamble involving a 50% chance of winning \$200 and a 50% chance of losing \$100, but offered to take the bet if he could play it 100 times in a row. Samuelson argued that his colleague's pattern of choices was inconsistent and therefore illogical. In fact, Samuelson proved that this pattern of choices is inconsistent with expected utility theory (EUT) in the following sense: if his colleague would reject the single play at every wealth position reachable in the multiple-plays gamble, then he should not accept the multiple gamble, according to EUT.

Samuelson's (1963) article was intriguing and influential partly because the preferences of Samuelson's colleague (SC) seem intuitively reasonable on first glance, but upon reflection can easily be criticized. SC rejected the single play of the mixed gamble, but offered to take it if he could play it 100 times (i.e., in prospective multiple trial choice). But what would happen if we offered SC the single gamble each day at lunch for 100 days (i.e., a repeated independent trials situation)? Based on his choice in the single-play condition, it seems he might refuse the gamble on each individual occasion. However, if he would reject the single gamble 100 times yet accept 100 'bundled' plays when offered, his behavior in these situations could be described as inconsistent, therefore irrational. Note that this argument (and Samuelson's proof) depend on assumptions about SC's hypothetical behavior in repeated independent trials, even though SC's choices are 'observed' only in the single-trial and prospective multiple-play situations.

Samuelson's colleague seemed to adopt an EV- (or EU-) maximizing strategy when faced with a prospective multiple-trial decision. Since Samuelson's (1963) article, a number of papers have investigated multiple-trial decision making, usually focusing on how choices in multiple-trial decision making differ from those made in single decision trials. For example, Coombs and Bowen (1971) found that participants ranked the risk of gambles differently depending on whether the gamble was to be played once or 25 times. In single-play, gambles with low odds of winning were considered riskier, but this was not true for repeated gambles. Keren and Wagenaar (1987) found that two standard violations of EUT described by Kahneman and Tversky (1979), the *certainty effect* and the *possibility effect*, were not replicated in multiple-trial situations. In two experiments, they also found a consistent preference for the higher-EV options in multiple-play problems with ten or 100 trials, for both gains and losses. Keren (1991) replicated the finding that the certainty effect does not obtain for multiple trials (where the number of plays was 5), and found that the multiple-play preference for the higher-EV gambles also obtained when the explicit distribution of outcomes for the five-play gamble was presented to participants and described as a one-stage gamble.

Wedell and Böckenholt (1990) demonstrated that *preference reversals* between choice and pricing are reduced when participants play gambles multiple times. Although Wedell and Böckenholt interpreted their results as supporting the general idea that violations of EUT are reduced for multiple-play gambles, their full set of results do not support the hypothesis that participants consistently behave in an EV-maximizing manner in multiple-play situations. To be sure, in their Experiment 1 most participants chose the p-bet (with lower EV) in single play and the $-bet (with higher EV) in multiple-play. However, in their Experiment 2 the set of gambles was constructed so as to unconfound EV and type of gamble: for half the gambles the p-bet had higher EV, and for half the $-bet had higher EV. Results showed that participants did not consistently choose the bet with higher EV in the multiple-play (gain) condition. Rather, a majority of participants chose the p-bet, the bet with the higher probability of winning, in all conditions: one play, ten plays, or 100 plays.

In the area of behavioral finance, several studies (Benartzi & Thaler, 1999; Gneezy & Potters, 1997; Thaler, Tversky, Kahneman, & Schwartz, 1997) have examined how the frequency with which investment payoffs are monitored ('evaluation periods') can affect risk propensity in investment decisions. By and large these studies have found that shorter evaluation periods lead to more risk-averse allocation decisions, and that longer periods induce more 'rational' (i.e., EV-consistent) behavior. These findings are consistent with the findings of studies of single- vs. multiple-trial decision making, because longer evaluation periods entail more aggregation of individual gain and loss events, just as do 'bundled' multiple decision trials.

To summarize, previous research contrasting single- and multiple-trial decision making has led some researchers to conclude that decision makers behave more in accord with the predictions of expected value theory (EVT) or expected utility theory (EUT) in the multiple-trials situation. We will refer to this as the 'long-run rationality' hypothesis. However, the empirical evidence for this idea is mixed at best. Importantly, this hypothesis runs counter to our prediction that in multiple-trial decision making most participants will prefer a mix of the sure-thing and risky options, because the EV of a mixed option must lie between the EVs of the purely sure-thing and purely risky strategies. Thus, testing the 'long-run rationality' hypothesis is another goal of our paper.

Bounded rationality in prospective multiple-trial decisions

One reason why people may not always maximize EV or EU in multiple-trial decisions is that it can be very difficult to evaluate the complex outcomes of prospective multiple-trial gambles. For example, suppose a decision maker is offered a choice between ten sure gains of $5, or ten plays of a simple gamble consisting of a 60% chance to win $10. The sure option here can be aggregated to a single gain of $50, but the distribution of outcomes for the risky option is complex, consisting of 11 different possible outcomes from $0 to $100, each with an associated probability. Some researchers (e.g., Benartzi & Thaler, 1999; Redelmeier & Tversky, 1992) have attempted to separate people's *preferences* in the multiple-trial situation from their cognitive capacity to mentally generate this distribution, by presenting such multiple-trial choice problems with or without a graph of the aggregate outcome distribution.

Of course, a simplifying decision heuristic that people could adopt in these multiple-trial choices is to simply evaluate the multiple-trial options by evaluating representative single plays of the sure and risky options. Redelmeier and Tversky (1992) referred to these two alternate choice strategies as involving either *aggregating* trials or *segregating* trials. Redelmeier and Tversky presented data on how risky choices change when a gamble is played multiple times, and when the aggregated multiple-trial distribution is calculated and presented explicitly as a probability distribution of outcomes. They found that the proportion of participants electing to play a simple mixed gamble (with positive EV) increased when the gamble was to be played five times rather than once (from 43% to 63%), and increased even more when the five-fold gamble's outcomes were presented explicitly in a chart (to 83%). The pattern of results suggests that the five-play repeated gambles are not always properly aggregated unless the explicit outcome distributions are shown, and that decision makers are more likely to pick the EV-maximizing option if the problem is made more 'cognitively transparent.'

This issue was also investigated by Benartzi and Thaler (1999), who described several experiments examining how people's choices in the multiple-play situation are affected by presenting a chart of the overall distribution of outcomes. Consistent with Redelmeier and Tversky's results, Benartzi and Thaler found that many people who decline multiple plays of a simple mixed gamble with positive EV will accept it if shown a chart of the distribution. For example, for a sample of evening-program MBA students from the University of Chicago, 64% accepted a single-play version of the Samuelson bet, while 65% accepted the multiple-play version. However, when shown the appropriate distribution chart of the multiple-play version, 86% accepted the gamble. Results differed somewhat for two samples of relatively naïve participants. For these participants, the multiple-play version was somewhat *less* desirable than the single-play version unless the explicit distribution of outcomes was shown in chart form, in which case 87% of them accepted the gamble.

These studies examining how charts of the outcome distributions affect multiple-trial decision making seem to call for caution in accepting the idea of long-run rationality. People may behave more 'rationally' (i.e., consistently with EV or EU) in multiple-trial decision making only if the multiple-trial problem they are facing is sufficiently *transparent* to allow them to find the optimal choice, or if decision aids (e.g., distribution graphs) are used to make the problem more transparent. This idea is consistent with Simon's (1955) notion of 'bounded rationality,' that people are rational only to the extent allowed by the limits of their information-processing capabilities. This hypothesis too was investigated in the present experiments. In particular, because our conjectured mixed-option preference is a form of non-rational behavior (in the sense of being inconsistent with EV maximization), the mixed-option preference may not arise when graphs of the multiple-trial outcome distributions are provided to the decision maker.

Overview of the experiments

Experiment 1 examined participants' preferences in multiple-trial situations when they are offered prospective choices between a purely sure-thing option, a purely risky option, and a mixed option consisting of several plays of the sure thing plus several plays of the risky option. In particular, we expected participants to show a 'mixed option preference,' i.e., to prefer the mixed option over the purely risky and purely riskless option. The results can also be used to examine the general claim that decision makers' choices are more consistent with EV maximization in multiple-trial decision making. Experiment 1 also tested if presenting participants with the graph of the outcome distribution affects preferences in prospective multiple-trial choices, as found by Redelmeier and Tversky (1992) and Benartzi and Thaler (1999).

Experiment 2 replicated the basic findings of Experiment 1 with a wider range of problems, and checked if a possible confound (the order in which problem version are given to participants) could have accounted for the mixed option preference found in Experiment 1.

Experiment 3 explored possible explanations for the observed mixed option preference. In particular, Experiment 3 tested whether the mixed option preference found in Experiments 1–3 might be due to a compromise effect (Simonson, 1989; Simonson & Tversky, 1992), in which the mixed option is viewed by decision makers as a compromise between the extremes of the purely sure-thing and purely risky options. The overall pattern of results from the trinary and binary choice conditions used in the study supported this explanation.

EXPERIMENT 1

This experiment examined participants' preferences in multiple-trial situations when they are offered prospective choices between a purely sure-thing option, a purely risky option, and a mixed option consisting of several plays of the sure-thing plus several plays of the risky option. Each risky gain has a higher EV than the sure gain it is paired with, and the risky losses have a larger expected loss than the sure losses. For the reasons listed in the introduction, we expected most participants to show a 'mixed option preference,' i.e., to prefer the mixed option over the purely risky and purely riskless options. The results can also be used to evaluate the 'long-run rationality' hypothesis, that participants' choices should show a higher proportion of EV-maximizing behavior in the multiple-trials situation. Finally, Experiment 1 also tested whether showing a chart of the outcome distribution increases the proportion of EV-maximizing behavior, as suggested by the results of Redelmeier and Tversky (1992) and Benartzi and Thaler (1999). Specifically, does showing a chart of the distribution decrease the proportion of choices of the mixed option and increase the proportion of choices of the EV-maximizing option?

Materials

Three different simple decision problem types were used, each offering a choice of a risky option (receive $\$ \times$ with probability p) versus a sure thing (a certain gain of \$y): G1 = (\$5; \$100, 0.1), G2 = (\$9; \$100, 0.1), and G3 = (\$9; \$20, 0.5). Three problems in the domain of losses were created from these three gain problems by reflecting the payoffs around 0: the resulting loss problems are labeled L1–L3. Each problem was presented in three versions: single-play versus multiple plays versus multiple plays with a mixed-strategy option. Another manipulation was whether or not a graph of the explicit distribution of outcomes for each option was shown. A single participant saw only one of the six basic problems listed above (Problem 1, 2, or 3 in either the Gain or the Loss condition), but saw it presented in six different versions: (a) single play, (b) multiple play with choice of purely risky and purely risk-free strategies, (c) multiple plays with risky, risk-free, and mixed strategy options, (d) single play with graphs showing the explicit distribution of outcomes for each option, (e) multiple play with choice of purely risky and purely risk-free strategies (with graphs), and (f) multiple plays with risky, risk-free, and mixed strategy options (with graphs). Use of such a design runs the risk that memory or stimulus comparison processes may affect how participants react to the more complex versions of the decision problem they face, but has the advantage that comparisons of the Version factor are more powerful, by virtue of the fact that they are within-subject comparisons.

For example, the following version of Problem 1 was presented in the condition that offered multiple trials with risky, risk-free, and mixed strategy options:

> You are offered a choice between the following two options ***ten times in a row***:
>
> Option: you have a 10% chance to receive \$100; otherwise you receive nothing
> Option: you receive \$5
>
> BUT, before we begin you must make your choices for each play. That is, you will actually be choosing between taking the first option on every play, the second option on every play or selecting a mix of choices. So you are really choosing between the following options:
>
> Option **K**: ten payments of \$5
> Option **D**: five payments of \$5, plus five 'plays,' each of which is a 10% chance to receive \$100.
> Option **H**: ten 'plays' of the game in which every play is a 10% chance to receive \$100.
>
> Which option do you prefer: **K**, **D**, or **H**? ________ (fill in)

An example version of Problem 1 that includes graphs of the distributions of outcomes for the three options is shown in the Appendix.

Procedure

The participants were graduate students, mostly in psychology and education, enrolled in large statistics or research-related courses at Teachers College, Columbia University. Students were contacted at the end of classes, and asked to volunteer to stay to fill out a 5-minute questionnaire. No payment was offered. There were six different forms of the survey, corresponding to each of the six basic decision problems. Each form contained six versions of one basic problem, as described above. A mix of the six forms of the instrument was distributed in each class. Participants were asked to answer one question at a time, not going back to look at or change their previous answers.

Results

The results are shown in Tables 1 and 2. Table 1 shows the frequency and percentage of choices for the three base problems in the domain of gains, for the single-, multiple-, and multiple-play-with-mixed-option versions. Results are shown separately according to whether or not a graph of the outcome distribution for each option was presented. Table 2 shows the corresponding data in the domain of losses.

Table 1. Experiment 1: frequency (and %) of choices for six problems in the domain of gains, when a chart of the outcome distribution is presented, and when it is not (for G1 $n=51$; for G2 $n=53$; for G3 $n=54$)

	No chart			With chart		
Gain problems:	Sure (%)	Mixed (%)	Risky (%)	Sure (%)	Mixed (%)	Risky (%)
G1 = ($5) versus ($100, 0.10)						
Single play	22 (43)	—	29 (57)	32 (63)	—	19 (37)
Multiple plays	26 (51)	—	25 (49)	17 (33)	—	34 (67)
Multiple/mixed	9 (18)	33 (65)	9 (18)	9 (18)	24 (47)	18 (35)
G2 = ($9) versus ($100, 0.10)						
Single play	29 (55)	—	24 (45)	25 (47)	—	28 (53)
Multiple plays	36 (68)	—	17 (32)	28 (53)	—	23 (47)
Multiple/mixed	11 (21)	34 (64)	8 (15)	14 (26)	22 (42)	17 (32)
G3 = ($9) versus ($20, 0.50)						
Single play	34 (63)	—	20 (37)	34 (63)	—	20 (37)
Multiple plays	29 (54)	—	25 (46)	30 (56)	—	24 (44)
Multiple/mixed	9 (17)	39 (72)	6 (11)	13 (24)	28 (52)	13 (24)

Table 2. Experiment 1: frequency (and %) of choices for six problems in the domain of losses, when a chart of the outcome distribution is presented, and when it is not (for L1 and L2, $n=51$; for L3 $n=50$)

	No chart			With chart		
Loss problems:	Sure (%)	Mixed (%)	Risky (%)	Sure (%)	Mixed (%)	Risky (%)
L1 = (−$5) versus (−$100, 0.10)						
Single play	33 (65)	—	18 (35)	30 (59)	—	21 (41)
Multiple plays	25 (49)	—	26 (51)	34 (67)	—	17 (33)
Multiple/mixed	19 (37)	13 (26)	19 (37)	29 (57)	10 (20)	12 (23)
L2 = (−$9) versus (−$100, 0.10)						
Single play	22 (43)	—	29 (57)	23 (45)	—	28 (55)
Multiple plays	15 (29)	—	36 (71)	21 (41)	—	30 (59)
Multiple/mixed	14 (28)	19 (37)	18 (35)	12 (24)	26 (51)	13 (26)
L3 = (−$9) versus (−$20, 0.50)						
Single play	15 (30)	—	35 (70)	12 (24)	—	38 (76)
Multiple plays	22 (44)	—	28 (56)	11 (22)	—	39 (78)
Multiple/mixed	8 (16)	23 (46)	19 (38)	7 (14)	16 (32)	27 (54)

Our first hypothesis was that most participants would prefer a mixed-strategy option over purely risky- and purely risk-free strategies in the multiple-trial with mixed-option situation. For gains (Table 1) the predicted pattern clearly held. When no chart was presented, the mixed-strategy option was chosen by a majority of participants for all multiple-trial problems with a mixed option (65% for Problem G1, 64% for Problem G2, and 72% for Problem G3). These choice proportions were higher than chance (i.e., 33%), Pearson χ^2 (2, $n = 158$) = 81.35, $p < 0.05$. For the Loss domain (Table 2), with no chart, results varied across problems, though the mixed option was still chosen frequently. For Problems L2 and L3 the mixed option was the modal choice (37% and 46%, respectively), but for Problem L1 participants were evenly split between the risky and sure-loss strategies (37% each). A log-linear analysis indicated that for losses there was a strong Problem by Strategy association: the likelihood ratio (LR) χ^2 (2) = 30.64, $p < 0.05$, indicating that Strategy choice did differ across Problems. Thus, a strong mixed option preference was observed for gain problems, but not for all loss problems.

The results did not support the 'long-run rationality' hypothesis that participants should show a higher proportion of EV-consistent choices when confronted with multiple plays of a decision problem or gamble. In the multiple-play gain problems (Table 1), a majority of participants (51% in Problem G1; 68% in Problem G2; 54% in Problem G3) selected the sure thing for gains rather than the EV-maximizing risky option. For multiple-trial loss problems (Table 2) a majority of participants (51% in Problem L1; 71% in Problem L2; 56% in Problem L3) selected the risky option rather than the EV-maximizing sure loss. Comparing these percentages to those obtained in the single-play condition, it can be seen that for Problems 1 and 2, the proportion of EV-consistent choices actually *declined* when problems are offered as multiple plays, both for gains and for losses. Only for Problem 3, which offered a 50% chance of winning (or losing), were the changes in the direction predicted by the long-run rationality hypothesis, for both gain and loss problems.

Our third hypothesis proposed that increasing the transparency of the multiple-trial problems by showing a chart of the explicit distribution of outcomes would increase the proportion of EV-consistent behavior. For gains (Table 1), when a chart was shown (Table 1), the mixed option was no longer preferred by a majority of participants in all three problems, although it was still the modal choice for each, with 47%, 42%, and 52%, respectively, selecting the mixed option. A test of whether these choice proportions were higher than chance level (i.e., 33%) was significant, Pearson χ^2 (2; $n = 147$) = 14.329, $p < 0.05$. Thus, a mixed option preference is observed for gain problems, even when a chart of the outcome distribution is presented. However, these proportions are lower than the proportions when a chart was *not* presented (65%, 64%, and 72%, respectively). A direct test of this effect was accomplished by a log-linear analysis, using a model with factors Problem (= G1, G2, G3), Chart (= no, yes), and Strategy (= mixed, other). In this model, the Chart $\times$ Strategy association was significant, LR χ^2 (1) = 13.404, $p < 0.05$, demonstrating that the proportion of participants selecting the mixed option differed in the chart and no-chart conditions. Overall, 67% of participants chose the mixed option with no chart versus only 47% when a chart was shown. Thus, presenting a chart of the outcome distribution decreases the proportion of choices of the mixed option.

For gains, participants not only selected the mixed option less often when shown a chart, but also chose the EV-maximizing risky option more often. For mixed-option gain problems, when no chart was shown 15% of participants chose the (EV-maximizing) purely-risky strategy, as compared to 30% with a chart. This difference was significant, as shown by the significant Chart $\times$ Strategy association, LR χ^2 (1) = 11.643, $p < 0.05$, in a log-linear analysis with Problem, Chart, and Strategy as factors.

However, for the loss domain no such effects were observed. Showing a chart did slightly lower the proportion of participants selecting the mixed option, but only from 37% to 34% overall (ignoring problem). This difference was not significant. Inspection of the data reveals that for losses the effect of showing a chart varied across problems: the modal choices with a chart were the sure loss for Problem L1 (57%), the mixed option for Problem L2 (51%), and the purely risky strategy for Problem L3 (54%). For losses, the percentages of participants selecting the EV-maximizing pure sure-loss option did increase slightly with a chart (from 27% to 32%), but this difference was not significant in the appropriate log-linear analysis, as shown by the non-significant Chart × Strategy association, LR χ^2 (1) = 0.852, $p > 0.05$. Thus, for the mixed-option problems in the loss domain no significant effect of showing the chart was found.

The effects of showing a chart of the outcome distributions can also be assessed for the multiple-trial problems with *no* mixed option. An increase in the proportion of participants selecting the EV-maximizing option would replicate and extend the findings of Redelmeier and Tversky (1992) and Benartzi and Thaler (1999). The results in Tables 1 and 2 show that for the multiple-play gain problems (with no mixed option) the proportion of participants choosing the EV-maximizing risky option increased with a chart, from 42% to 53%. However, in a log-linear analysis of the multiple-trials Gain condition data with Chart, Problem and Strategy as factors, the three-way association of all three factors was significant. This means that the effect of Chart on Strategy differs by Problem for these data. Specifically, the percentage of participants selecting the risky options increased from 49% to 67% for Problem G1, and from 32% to 47% for Problem G2, but was essentially flat (changing only from 46% for 44%) for Problem G3, which offered a sure gain of \$9 versus the (EV-maximizing) choice of a 50% chance to win \$20. Problem G3 is unique among the problems in that it offers a 50% probability of the risky gain.

For multiple-trial loss problems (with no mixed option), without a chart 41% of participants chose the EV-maximizing sure-loss options overall, compared to 43% with a chart. This difference was not significant, and the apparent effect differed across problems. Just as for the gain-condition data, for Problems L1 and L2 there was an increase in the proportion of EV-consistent choices with a chart (from 49% to 67% for Problem L1, and from 29% to 41% for Problem L2), but this did not hold for Problem L3, which offered the EV-maximizing option of a sure loss of \$9 versus the option of a 50% chance to lose \$20. Here, the chart actually reduced the proportion of EV-maximizing sure-thing choices, from 44% to 22%.

Discussion

Our first hypothesis was that most participants would prefer a mixed-strategy option, when available, over purely risky- and purely risk-free strategies. This hypothesis was strongly supported: the mixed option was the modal choice for participants for all three problems in the domain of gains, and for two out of three problems in the domain of losses. This result casts the findings of previous studies of multiple-trial decision making in a new light. Some previous studies examined differences in single- and multiple-trial decision making by offering participants a prospective choice between purely risky and purely risk-free strategies. The present results show that this limited range of options may not include participants' true preference (the mixed option) in multiple-trial situations. Note that a multiple-trial mixed option cannot be optimal in terms of expected value (Corter & Chen, 2005): the ordering of options by expected value must be either: EV(risky) > EV(mixed) > EV(sure), or EV(sure) > EV(mixed) > EV(risky). Thus the current findings offer a new type of evidence that decision behavior is not always 'rational' in an EV-sense, even in multiple-trial decision making.

A possible limitation is that this strong mixed-option preference for gains was obtained in a design in which participants viewed multiple versions of the same decision problem: first the single trial version, then the version with multiple trials, first without and then with a mixed option. It is possible that the mixed option was favored in the multiple-trials with mixed options because of a novelty effect. That is, after participants had encountered the sure-thing and risky options in the single-trial and multiple-trial versions, the mixed option may have been especially attractive in part because of its novelty. This possibility is addressed by Experiment 2, which is designed so that participants encounter only one version of each decision problem.

The second hypothesis, long-run rationality, holds that decision makers will show a higher proportion of EV-consistent choices when faced with multiple trials of a decision problem or gamble. This hypothesis was not supported. For Problems 1 and 2, where the risky option had a probability of 0.10, participants chose the EV-maximizing options *less* often in the multiple-trials situation, both for gains and for losses. Only for Problem 3, where the risky prospect had a probability of 0.5, did the claimed pattern hold.

It is not immediately clear why participants might behave differently for Problem 3. According to Prospect Theory (Tversky & Kahneman, 1992), people systematically underweight moderate to high probabilities (such as 0.5), and systematically overweight low probabilities (such as 0.1). But this assumption does not explain why these two types of problems should differ in the single- and multiple-trial versions. Our explanation is as follows. People find it difficult to compute EV for the multiple-trial problems, thus they are less likely to discover that the (multiple-trial) risky option is optimal by this criterion. However, problems involving risky prospects with $p = 0.5$ may be especially transparent for participants, and for that reason may be especially attractive when the risky option has higher EV. Support for this interpretation can be found in the literature: Edwards (1954) found that problems with risky options involving $p = 0.5$ were chosen especially often in binary choice.

Our third hypothesis was that providing a chart of the explicit distribution of outcomes will reduce the percentage of participants preferring the mixed-play option, and increase the percentage selecting the EV-maximizing option. Indeed, this seemed to be the case. For gains the presence of the chart significantly reduced the proportion of participants selecting the mixed option and increased the proportion of participants selecting the purely risky (EV-maximizing) option. This suggests that the mixed-option problems, like the multiple-trial problems studied by Redelmeier and Tversky (1992) and Benartzi and Thaler (1999), are complex enough so that they are not typically aggregated by decision makers into a single distribution of outcomes, unless this framing process is aided by providing a chart. When the chart is available to assist such framing, increased levels of EV-consistent behavior are observed for the multiple-trial gain problems that include a mixed option. However, for the multiple trial problems without a mixed option, charts increased the proportion of participants selecting the EV-maximizing option only for gain Problems G1 and G2, but not for Problem G3 (with a 0.5 probability for the risky gain). For losses the effects of providing charts were neither consistent across problems, nor significant.

Why does presenting a chart of the outcome distribution have these effects? In the multiple-trial decision problems, the mixed-play option (like the purely risky option) has a complicated outcome distribution. Presenting a chart of this distribution makes the decision problem more transparent. When the decision problem is more transparent, participants' decisions often seem to be more EV-consistent (Redelmeier & Tversky, 1992), a trend that seems based on a general (though not universal) desire by naïve participants to be 'rational' in their decisions, where the optimal choice is clear to them. In the present Experiment 1, the proportion of choices of the maximal-EV option increased for gain problems when a chart was shown, and the proportion of choices of the mixed option declined. This pattern suggests that people's preferences for the mixed option may reflect operation of some simple decision heuristic invoked by participants in the face of decision problem complexity, rather than reflecting deep-rooted and stable preferences. Accordingly, Experiments 2 and 3 were designed to assess the robustness and stability of the mixed option preference.

EXPERIMENT 2

Experiment 2 was designed to try to replicate the preference for mixed options in multiple-trial decisions, and to extend these results by using a wider range of problems. Demonstrating the mixed-option preference with a wider range of problems is important, because the results of Experiment 1 suggest that how people respond in single- versus multiple-trial decision situations may depend on specific aspects of problems, in particular the probability of winning or losing for the risky option and on the relative size of the EVs of the risky and riskless options. Experiment 2 was also designed to remove a possible confound in the design of Experiment 1. In Experiment 1, the different versions of a given problem were given in a fixed order, so that the multiple-trials-with-mixed-option versions were always given after the single- and multiple-trials versions. This confound could be producing the mixed-option preferences, since the newly introduced mixed option has the advantage of novelty over the purely-sure and purely-risky options. In Experiment 2, order of these three versions of each problem was counterbalanced across forms to remove the confounding of version with order of presentation.

The main hypothesis investigated in Experiment 2 was that in the multiple-play condition with a mixed option available, decision makers will prefer the mixed option over the purely risky and purely risk-free strategies ('mixed-option preference'). However, we also evaluated the 'long-run rationality' hypothesis, namely that in multiple-play situations participants will be more prone to pick the highest-EV option (the risky option for gains, and the sure thing for losses). Also, in this study a variety of risky prospects were used, including options with payoff probabilities ranging from 0.01 to 0.67. Systematically investigating the effect of payoff probability seems desirable, because it has been established (Tversky & Kahneman, 1992) that people are often risk averse for high-probability gains but risk-seeking for low-probability gains (and conversely for losses). Finally, based on the results of Experiment 1, we surmised that decision problems involving risky prospects with $p = 0.5$ may be especially transparent for participants, leading to more EV-maximizing behavior. Thus, decision problems with payoff probabilities of 0.5 were specifically included.

Materials

Six basic decision problems were used. These problems are shown in Table 3. Each problem had two variants, one in the domain of gains and one in the domain of losses, leading to 12 different basic decision problems.

Problems 1–3 are the same as those used in Study 1. With the addition of Problem 4 here, these four problems formed a 2×2 factorial design with factors Probability Level (of the risky option), set at either 0.1 or 0.5, and EV ratio, set at either 5:10 or 9:10. The effects of Probability Level can also be gauged across a wider range of values, because for Problems 2, 3, 5, and 6 the probability of the risky gain (or loss) is varied from 0.01 to 0.67, while holding EV ratio constant at 9:10. In this experiment, a single participant saw all six experimental problems, two problems in each of three Versions: (a) single play, (b) multiple play (10 trials) with choice of purely risky and purely risk-free strategies, and (c) multiple plays with risky, risk-free, and mixed strategy options. Thus, problem Version was varied across problems rather than within problems.

Procedure

The participants were undergraduate students at Columbia University, drawn from multiple sections of an introductory statistics course. They were contacted at the end of a class session and asked to stay to fill out a 5-minute questionnaire. No payment was offered. Participants were instructed to answer one question at a time, not going back to look at or change their previous answers.

Each participant encountered all six basic decision problems, in either the gain or the loss domains. Thus, Domain was the only between-participants factor. For each participant, two problems were encountered in each of the three Versions: single play, multiple plays, and multiple plays with mixed option. Problem and Version were counterbalanced across three test forms. So, for example, one group of participants received Problems 1 and 4 in the single-play version, Problems 3 and 5 in the multiple-play, and Problems 2 and 6 in the multiple-play with mixed option version. Within each form, several sub-forms were used, in which order

Table 3. Decision problems used in Experiment 2

Gain domain	Loss domain
G1 = (\$5) versus (\$100, 0.1)	L1 = (−\$5) versus (−\$100, 0.1)
G2 = (\$9) versus (\$100, 0.1)	L2 = (−\$9) versus (−\$100, 0.1)
G3 = (\$9) versus (\$20, 0.5)	L3 = (−\$9) versus (−\$20, 0.5)
G4 = (\$5) versus (\$20, 0.5)	L4 = (−\$5) versus (−\$20, 0.5)
G5 = (\$9) versus (\$1000, 0.01)	L5 = (−\$9) versus (−\$1000, 0.01)
G6 = (\$9) versus (\$15, 0.67)	L6 = (−\$9) versus (−\$15, 0.67)

of presentation of the questions was randomized. All participants received one single-trial decision problem as a warm-up or practice problem.

Results

Results of the experiment are shown in Table 4. In this table the problems are ordered by the probability of the risky gain.

The overall patterns in the table are as follows. For the single play problems in the gain domain participants tended to select the risky option (with higher EV) for all problems except Problem 5. This tendency was most pronounced for the two problems (3 and 4) that have $p = 0.5$. In the multiple plays gain condition (with no mixed option) participants preferred the sure thing for all three of the low-probability problems (Problems 5, 1, and 2), but tended to select the risky (EV-maximizing) option for the high-probability problems (Problems 3, 4, and 6). For losses, in single play participants tended to select the higher EV sure loss option for all problems except Problem 3, with a 9:10 EV ratio and a $p = 0.5$ chance of losing for the risky option. In the multiple plays Loss condition, participants tended to be risk seeking, except for Problem 4, with an extreme 5:10 EV ratio and $p = 0.5$, in which a majority of participants prefer the (EV-maximizing) sure loss of $5.

The results for the multiple-play problems with a mixed option enable a test of the hypothesis that the mixed option will be the most preferred. For gains, this hypothesis was strongly confirmed: for all six

Table 4. Experiment 2 results: frequency (and %) of choices for six problems in three versions, shown separately for gains and losses

	Gains			Losses		
Problem version:	Sure (%)	Mixed (%)	Risky (%)	Sure (%)	Mixed (%)	Risky (%)
Problem 5	G5 = ($9) versus ($1000, 0.01)			L5 = (−$9) versus (−$1000, 0.01)		
Single play	20 (54)	—	17 (46)	21 (53)	—	19 (47)
Multiple plays	24 (71)	—	10 (29)	9 (26)	—	26 (74)
Multiple/mixed	9 (26)	21 (60)	5 (14)	7 (19)	11 (30)	18 (50)
Problem 1	G1 = ($5) versus ($100, 0.10)			L1 = (−$5) versus (−$100, 0.10)		
Single play	16 (46)	—	19 (54)	25 (70)	—	11 (30)
Multiple plays	24 (65)	—	13 (35)	18 (45)	—	22 (55)
Multiple/mixed	10 (29)	16 (47)	8 (24)	11 (31)	14 (40)	10 (29)
Problem 2	G2 = ($9) versus ($100, 0.10)			L2 = (−$9) versus (−$100, 0.10)		
Single play	16 (47)	—	18 (53)	19 (54)	—	16 (46)
Multiple plays	19 (54)	—	16 (46)	16 (44)	—	20 (56)
Multiple/mixed	9 (24)	24 (65)	4 (11)	7 (17)	19 (48)	14 (35)
Problem 4	G4 = ($5) versus ($20, 0.50)			L4 = (−$5) versus (−$20, 0.50)		
Single play	6 (18)	—	28 (82)	34 (97)	—	1 (3)
Multiple plays	5 (14)	—	30 (86)	26 (72)	—	10 (28)
Multiple/mixed	6 (16)	19 (52)	12 (32)	18 (45)	13 (33)	9 (22)
Problem 3	G3 = ($9) versus ($20, 0.50)			L3 = (−$9) versus (−$20, 0.50)		
Single play	11 (30)	—	26 (70)	16 (40)	—	24 (60)
Multiple plays	16 (47)	—	18 (53)	17 (49)	—	18 (51)
Multiple/mixed	12 (34)	15 (43)	8 (23)	7 (19)	17 (47)	12 (33)
Problem 6	G6 = ($9) versus ($15, 0.67)			L6 = (−$9) versus (−$15, 0.67)		
Single play	16 (46)	—	19 (54)	20 (56)	—	16 (44)
Multiple plays	17 (46)	—	20 (54)	17 (42)	—	23 (58)
Multiple/mixed	10 (29)	17 (50)	7 (21)	13 (37)	13 (37)	9 (26)

problems the mixed option was the modal choice, chosen 58.3% of the time overall (vs. 26.3% for the purely-sure and 20.8% for the purely risky options). This proportion of choices of the mixed option is significantly higher than chance (33.3%), Pearson χ^2 (1; $n = 212$) = 34.29, $p < 0.05$. For losses, results were more variable: the mixed option was the modal choice for only three out of the six problems. Overall, the mixed option was selected most often for loss problems: 38% of the time versus 29% of the time for the purely-sure options and 33% for the purely-risky options, but the overall proportion of choices of the mixed option was not significantly higher than chance, Pearson χ^2 (1; $n = 222$) = 2.204, $p > 0.05$. These results show that participants often prefer mixed options for risky gain problems, but do not do so consistently for loss problems.

The second hypothesis, long-run rationality, was not supported. In fact, for four out of six gain problems the results were in the opposite direction, with *fewer* participants choosing the higher-EV risky option in the multiple-play problem versions, compared to the single play. In a log-linear analysis with Problem (= 1–6), Version (= single, multiple), and Strategy (= sure, risky) as factors, this difference in the direction opposite to that supposed by the hypothesis was significant (as measured by the Version × Strategy association): LR χ^2 (1) = 4.327, $p < 0.05$. The sole exception to this reverse trend was Gain Problem 4, with an EV ratio of 5:10 and a win probability of 0.5. Only for this problem did more participants choose the risky option under multiple play. For losses, the second hypothesis was again not supported, and the observed pattern was again the reverse of that predicted: for five out of six loss problems, *fewer* participants chose the higher-EV sure loss in the multiple-play problem versions. This difference was also significant by a test of the Version × Strategy association in the appropriate log-linear model, LR χ^2 (1) = 11.095, $p < 0.05$. Only for loss Problem 3 (EV ratio = 9:10, $p = 0.5$) did participants choose the sure loss more for the multiple-play condition. To summarize, participants in Experiment 2 made *more* EV-consistent decisions for single-trial problems than in the corresponding multiple-trials versions, contrary to the predictions of the long-range rationality hypothesis. Note that the only two exceptions to this otherwise consistent pattern in the *reverse* direction from that expected under long-run rationality are two problems that may be especially transparent to participants, because they involve payoff probabilities of 0.5 for the risky option.

It is instructive to examine more closely the pattern of risky choice and its relationship to the probability of the risky gain (or loss). Figure 1 shows the pattern of choice probabilities for those problems (Problems 2, 3, 5, and 6) having an EV ratio of 9:10, with win probabilities ranging from 0.01 to 0.67. For gain problems (both in single- and multiple-trial forms), participants were more prone to select the risky (max-EV) option for the high-probability problems ($p = 0.5$ and $p = 0.67$) than for the low-probability problems ($p = 0.01$ and $p = 0.1$), consistent with the findings of Wedell and Böckenholt (1990). For loss problems, the reverse trend was found, with more participants selecting the *sure-thing* (but max-EV) option when the probability of the loss was high ($p = 0.5$ or $p = 0.67$).

Discussion

The results of Experiment 2 replicate the most important findings of Experiment 1. First, the strong preference for mixed options over purely risky and purely risk-free strategies was confirmed for problems in the domain of gains: the mixed option was the modal choice for all six gain problems, and was selected 58% of the time overall. For losses, the mixed option was the modal choice for three out of the six problems, and was selected 39% of the time overall. The results of Experiment 2 also eliminate the alternative explanation for the mixed-option preference in Experiment 1 that the preference could have been due to a novelty effect arising because those participants experienced multiple versions of a single decision problem, because in Experiment 2 participants saw only one version of each decision problem. This finding from Experiments 1 and 2, that mixed options are typically preferred over purely risky and purely risk-free strategies for multiple-trial decisions involving gains, is striking and important. As noted, this pattern of preferences violates EV theory. It may also violate EU theory, at least with certain classes of utility functions (Corter &

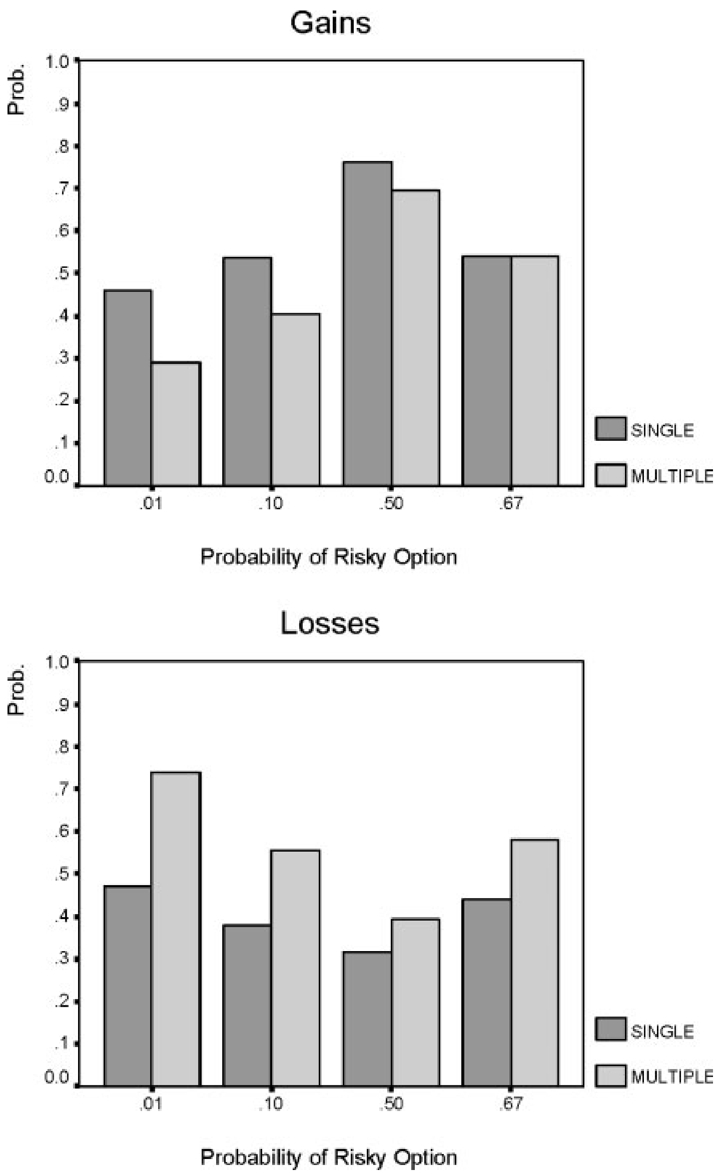

Figure 1. Probability of selecting the risky option, by Domain (gains or losses) and probability of winning (or losing) for four problems used in Experiment 2 (Problems 5, 2, 3, and 6, respectively)

Chen, 2005). Thus, the present results extend the range of findings showing that people's choices in simple monetary gambles do not always follow 'rational' principles.

Single- versus multiple-play

The results of Experiment 2 do not support 'long run rationality,' namely the idea that people behave more in accordance with expected value and expected utility theories when faced with multiple plays of a gamble. In fact, the pattern of choices here shows the opposite pattern, with *fewer* EV-consistent choices in the multiple-play condition. It may be that participants are better at estimating EV and basing choices on those estimates for single-trial decisions, but find it harder to estimate EV for multiple-trial problems unless graphs of the distributions of outcomes are provided. Therefore, for the more complicated multiple-play problems participants may fall back on simple heuristics to make their choices. For example, in Experiment 2 participants tended to increase their proportion of risky choice as the *probability* of a risky gain increased (both for single- and multiple-trial problems), and showed the converse pattern for risky losses. This finding is consistent with previous research on single-trial decision-making: Lopes (1981) noted that the probability of a risky gain was an important factor for many participants in single play, and Kahneman and Tversky (1979; Tversky & Kahneman, 1992) found that participants tended to choose those risky options with the higher probability of winning for moderate-probability problems, like those used here. For multiple plays, previous results in the literature are inconsistent: Keren and Wagenaar (1987; Keren, 1991) found an increased preference for higher-EV gambles in multiple play, but in Wedell and Böckenholt's (1990) Experiment 2 the higher-probability gain was generally preferred, even when it was not the higher-EV option.

Why are mixed-strategy options preferred?

In the introduction we mentioned several possible reasons to expect a mixed option preference, reasons that can now be evaluated as potential explanations for the mixed option preference found in Experiments 1 and 2. The first reason (though perhaps unsatisfying as an explanation), is that decision behavior is inherently inconsistent. For example, consumers typically buy a variety of products, a phenomenon known as 'variety seeking.' Furthermore, in *prospective multiple-trial* decisions variety seeking may be enhanced, a tendency that has been dubbed 'diversification bias' (Read & Loewenstein, 1995; Simonson, 1990). A skeptic might argue that there is little reason to expect a type of 'variety seeking' in the present type of experiments involving simple financial gambles rather than consumer choices, because money is fungible. However, in *prospective* decisions the prospect of receiving a sure gain in the future and the prospect of a future probabilistic gain are not the same type of good.

Similarly, in the introduction we mentioned the fact that in repeated decisions with feedback, decision makers often select a mix of options across many trials. However, the present tasks involve 'description-based' prospective decisions in which all relevant outcomes and probabilities are provided to participants for each problem, and no feedback is experienced. In such description-based tasks there may be little or no reason to expect response variability. For example, EUT predicts response variability only as a result of changes in the decision maker's wealth position, while Prospect Theory (Kahneman & Tversky, 1979) predicts such variability only via differences in framing of the problem.

The literature in behavioral finance offers a reason to expect a mixed option preference in prospective description-based decision making, namely the fact that lay investors often exhibit an apparent 'naïve diversification' heuristic (Benartzi & Thaler, 2001) for personal investments. This heuristic refers to the observed tendency of non-professional investors to allocate resources to a mix of different portfolio components (e.g., a 50–50% allocation of money to stocks and bonds). However, Benartzi and Thaler do not provide evidence (nor claim) that this heuristic is applied outside the context of investment portfolio decisions, so it is not obvious that these investing norms or heuristics would be invoked in the present context.

A more plausible explanation for the observed mixed option preference invokes the notions of framing and *mental accounting* (Kahneman & Tversky, 1979). The descriptions of the mixed options presented to participants in the present studies may promote a framing of the problem in which the five sure gains (or losses) are treated separately from the five risky prospects. The offered sure things and risky prospects are *qualitatively* different types of options, and this distinction may be emphasized by the wording of the problems, in which the sure things are described as payments (in the gain domain) and the risky options are described as 'plays' of a game. For these reasons, we suspect that participants might frame the mixed option by separately aggregating the five sure-gain trials and the five risky-gain trials, a possibility we will term 'quasi-aggregation.' Thus, a participant might aggregate five sure gains of \$5 into a single sure gain of \$25, and aggregate five chances to win \$20 (each with 50% probability) into the corresponding distribution of possible gains ranging from \$0 to \$100.

Why should decision makers use this 'quasi-aggregation' framing, aggregating the five risky prospects separately from the five sure things, rather than completely aggregating all ten risky and sure prospects into one overall expected outcome distribution? One possible reason is that participants may not be sure how to mathematically combine sure gains with probabilistic ones. Another possible explanation stems from the work of Thaler (1985), who advanced the *hedonic editing* hypothesis that decision makers will tend to mentally aggregate or segregate events so as to maximize subjective values. Empirical support for a modified version of the hedonic editing hypothesis was reported by Thaler and Johnson (1990), and by Linville and Fisher (1991). Quasi-aggregation framing should increase the attractiveness of the mixed options studied here, because the usual concave utility function for gains makes the anticipation of receiving two moderate gains more attractive than the anticipation of receiving their sum. Thus, a tendency toward hedonic editing would tend to elicit the quasi-aggregation framing for the mixed options.

This line of argument suggests that quasi-aggregation framing might account for the mixed option advantage, if the mixed option is 'quasi-aggregated,' while the ten prospects comprising the purely sure-thing and the purely risky options are completely aggregated. This is because receiving two small gains (a sure and a risky one) should be preferred to receiving one large gain (comprised of either all sure or all risky gains), assuming that the sure and risky options are roughly equally attractive. The results of Experiment 1 offer only mixed support for this account, however, because when the graphs of the outcome distributions of the various options were shown to participants, presumably promoting total rather than quasi-aggregation for the mixed option, their preference for the mixed option diminished but did not vanish.

Why then might a mixed option preference remain even when graphs of the outcome distributions are shown? Under quasi-aggregation, the mixed option can be seen as a multidimensional alternative that offers a moderate amount of two kinds of gains: sure and risky. This analysis suggests another possible reason for the mixed-option advantage observed in Experiments 1 and 2: that it might be due to a 'compromise effect' (Simonson, 1989; Simonson & Tversky, 1992). A compromise effect occurs if an option is chosen because it represents a compromise between two extreme options.

If participants frame the mixed options by 'quasi-aggregation,' then they may view the three choice options as essentially two-dimensional alternatives, each comprised of some number of sure gains (either 0, 5, or 10) combined with some number of risky gains (10, 5, or 0). In this framing, the sure thing option is viewed as having ten sure gains and no risky gains, the purely risky option as offering no sure gains but ten risky gains, and the mixed option as offering five of each. The three choice options framed in this way resemble the two-dimensional choice stimuli that Simonson and Tversky (1992) used to demonstrate compromise effects. More formally, suppose choices are being made among multidimensional options az, by, and cx, where the utility values of the three options on Dimension 1 are $a > b > c$, and on Dimension 2 are $z < y < x$. A compromise effect is demonstrated if the ratio of choice probability of by over az is increased when cx is added to the choice set, making option by the compromise choice between az and cx. In Experiment 3 we sought to test if the mixed-option preference documented in Experiments 1 and 2 could arise through a compromise effect, by offering the three types of multiple-trial options in both binary and trinary choice sets.

EXPERIMENT 3

Experiment 3 was designed to investigate possible reasons for the mixed option advantage, in particular if the mixed option preference is due to a compromise effect. Thus, one condition of the experiment was a *trinary choice condition* like that used in Experiments 1 and 2, offering a choice between ten sure gains, ten risky gains, or a mixture of five of each. In a second condition, however, participants faced *binary choices* between *pairs* of the same three multiple-trial options: risky versus mixed, mixed versus sure, and sure versus risky. If the mixed option preference were found in the trinary choice condition but not in the binary, that would be strong evidence that the mixed option preference is due to a compromise effect, in which it is selected because it a compromise between two extreme options. In addition, in the trinary-choice condition participants were asked to give justifications for their choices in the mixed-option multiple-play problems. These justifications were coded and analyzed, to see if they lent support to any of the possible explanations for the mixed-option advantage discussed above.

Materials

Four base decision problems in the domain of gains (or losses) were used, each offering a choice of a sure thing (a certain gain of \$y) versus a risky option (win \$ × with probability p). These four problems were: G1 = \$32 versus (\$160, 0.25), G2 = \$20 versus (\$75, 0.33), G3 = \$60 versus (\$150, 0.50), and G4 = \$40 versus (\$75, 0.67). The corresponding Loss versions of the four problems were identical (except that all payoffs were reversed in sign), and are referred to as L1–L4. Note that in each problem the ratio of EVs for the certain versus the risky option is 8:10.

Procedure

The participants were graduate students, mostly in psychology and education, enrolled in large research-related lecture courses at Teachers College, Columbia University. Students were contacted at the end of classes, and asked to stay to fill out a 5-minute questionnaire. No payment was offered.

In the Trinary Choice condition, each participant received all four decision problems, each presented in two versions: single play, and multiple plays with a mixed-strategy option (resulting in a total of eight decision problems presented to each participant). Each participant received the four simple decision problems in the single play condition (randomly ordered), then in the multiple-play-with-mixed-option version (randomly ordered). Some participants ($n = 65$) received the Gain version of the questionnaire while others ($n = 47$) received the Loss version. Participants were asked to answer one question at a time, without going back to look at or change their previous answers. On the final page of each survey, participants were asked for justifications; specifically, for the reason why they chose the option they did on the last problem (always a multiple-trial-with-mixed-option version).

In the Binary Choice condition, each participant received all four decision problems (one single-play plus three binary-choice multiple-play problems) generated from a single base problem (G1–G4, or L1–L4). For example, in the multiple-play versions of decision problem G1, a participant would face three binary choices between each possible pair of options G1(sure), G1(risky) and G1(mixed), where

G1(sure) = ten payments of \$32
G1(risky) = ten 'plays' of the prospect (\$160, 0.25)
G1(mixed) = five payments of \$32, plus five 'plays' of the prospect (\$160, 0.25).

The four resulting decision problems (one single-play plus three multiple-play binary choices) were ordered randomly in a single participant's booklet, with one unrelated single-play problem given as an initial warm-up. A total of $n = 50$ participants faced Gain problems, while $n = 42$ encountered Loss problems.

Results

Trinary choice condition

In the single-play Gain condition participants were generally risk averse (Table 5). Specifically, they chose the sure-gain option on 67% of trials (averaged across problems) even though the risky option had a higher

Table 5. Experiment 3, Trinary Choice condition: frequencies (and %) of sure, risky, and mixed choices for four decision problems (G1-G4) in the domain of gains ($n = 65$) and four problems (L1–L4) in the domain of losses ($n = 47$)

	Single trial		Multiple trials with mixed option		
Problem	Sure (%)	Risky (%)	Sure (%)	Mixed (%)	Risky (%)
G1 = ($32; $160, 0.25)	42 (65)	23 (35)	35 (54)	26 (40)	4 (6)
G2 = ($20; $75, 0.33)	45* (74)	16* (26)	24 (37)	33 (51)	8 (12)
G3 = ($60; $150, 0.50)	39 (60)	26 (40)	24 (37)	30 (46)	11 (17)
G4 = ($40; $75, 0.67)	44 (68)	21 (32)	19 (29)	33 (51)	13 (20)
Total (Gain)	170 (66)	86 (34)	102 (39)	122 (47)	36 (14)
L1 = (−$32; −$160, 0.25)	24 (51)	21 (45)	10 (21)	17 (36)	20 (43)
L2 = (−$20; −$75, 0.33)	26 (55)	23 (49)	12 (26)	11 (23)	24 (51)
L3 = (−$60; −$150, 0.50)	24 (51)	23 (49)	14 (30)	17 (36)	16 (34)
L4 = (−$40; −$75, 0.67)	32 (68)	15 (32)	13 (28)	15 (32)	19 (40)
Total (Loss)	106 (56)	82 (43)	49 (26)	60 (32)	79 (42)

*$n = 61$ due to missing data.

expected value for all four problems. For the multiple-play gain problems with a mixed option, the mixed option was the modal choice overall (also for 3 of 4 individual problems). Overall, participants selected the mixed option 47% of the time, compared to 39% for the sure thing and 14% for the risky. Thus, the mixed option preference was confirmed for gain problems. For loss problems, in the single-play versions participants tended to choose the sure-thing option (on 56% of trials, averaged across problems). In the multiple-trial versions, the mixed option was the modal choice for loss Problem L3, but the risky option was the modal choice for the remaining three loss problems. Overall, 26% of loss-problem choices were of the sure loss, 32% were of the mixed option, and 42% were of the risky option. These results confirm the existence of a robust mixed-option preference for gains (but not for losses) in the Trinary Choice condition. Note that the loss condition results are also inconsistent with the long-run rationality hypothesis.

Analysis of coded justifications

Participants' decision justifications for their final choice problem in the Trinary Choice condition were analyzed as follows. A coding scheme was developed to code for presence of various types of justifications used for the final-problem choice. This coding scheme included codes for justifications mentioning person or trait attributions by the decision maker ('I am not lucky'), justifications mentioning certain specific goals (security, maximizing gains, or both), expressions of confusion or unwillingness to expend mental effort to decide on the best option, mentions of a general desire or heuristic for diversifying choices, mentions of an explicit motive to segregate or aggregate gains, and justifications involving potential regret or rejoicing (cf. Bell, 1982; Loomes & Sugden, 1982; Zeelenberg, van den Bos, van Dijk, & Pieters, 2002). Finally, uninformative responses ('it seemed like the best option') and blank or omitted justifications were also noted. Two coders independently coded the text of the justification offered for each decision problem for which the mixed option was chosen. A single response by a participant could contain more than one type of justification. Only a few discrepancies between raters were found in the initial coding, and they were easily resolved.

For choices of the purely sure-thing option ($n = 25$), the most common justification was a security motive, followed by various attributions mentioning aspects of the participant's character or personality traits (for example, 'I'm just not a risk-seeking person'). Only $n = 6$ participants chose the purely-risky option on their final problems, and of these, two indicated that they were motivated to maximize gains or expected value, while the other four chose to omit a justification or gave uninformative responses. For the $n = 25$ participants who selected the mixed option on their final problem, the most common justification given (by 10 respondents) was a joint mention of both a security motive and a maximizing-gain motive. The second most common response was an omitted justification. No justifications for mixed option choices explicitly mentioned a diversification heuristic, while only a single participant mentioned potential regret. For the Loss problems, no mixed-option preference was observed. Only 11 of 39 participants chose the mixed option in the multiple-trial loss problems, and 8 of these 11 omitted any justification for this choice. Thus, these data were not analyzed further.

Binary choice condition

The preferences of all participants ($n = 50$ for gains, $n = 42$ for losses) in the Binary Choice condition are shown in Panel A of Table 6 ('Sample 1'). In the binary-choice condition, a majority of participants were risk-averse in the domain of gains. This was true both for the single play gain problems, where 38 of 50 participants (76%) chose the sure-thing option, and for the multiple-play problems offering a choice between the purely sure-thing and the purely-risky options, where 33 of 50 (66%) chose the sure thing. In the Loss domain, most participants were risk-seeking. For the single-play loss problems, 23 of 42 participants (55%) chose the risky option, while in the multiple-play problems, 28 of 42 (67%) chose the purely-risky option over the purely sure-thing option. Thus, Table 6 shows mixed evidence concerning the long-run rationality hypothesis: for gains participants showed *more* EV-maximizing risky choices in the multiple trials condition (34% risky choices versus 24% in single-play), consistent with long-run rationality, but for losses they showed *fewer* EV-maximizing sure thing choices in multiple trials (33% versus 45% for single-play).

The most important thing to note about the results in Panel A of Table 6 is that there is little or no evidence for a mixed option preference in these binary-choice multiple-trial problems. Instead, for gain problems a general pattern of risk aversion was observed, with a majority (62%) of participants preferring the purely sure-thing option to the mixed and a majority (76%) preferring the mixed to the risky. For multiple-trial loss problems, the general pattern was of risk-seeking, with the mixed option being preferred to the purely sure thing option (by 74% of participants), but the risky being preferred to the mixed (by 64%). Examination of the individual-level data revealed that the mixed option was preferred to *both* the purely sure-thing and the purely risky options for only 11 of 50 participants (22%) in the gain domain, and for only 11 of 42 participants (26%) in the loss domain. Note that in the gain domain 10 of 50 participants showed intransitive patterns of preferences for the binary-choice problems. For losses, 4 of 42 participants showed intransitive patterns. However, reanalysis of the data after eliminating participants who showed intransitivities did not change the pattern of results concerning preferences for the mixed option.

Although the mixed option was not the modal choice, the binary choice probabilities in Panel A do offer some evidence for a slight bias toward the mixed option. Note that for gains the majority of participants (66%) preferred the multiple-trial purely sure-thing option to the purely risky one. The choice probability for the sure-thing over the mixed option is *less* extreme (only 62%), but the choice probability for the mixed option over the purely risky one is *more* extreme (76%). This pattern of choice probabilities violates strong

Table 6. Experiment 3, binary choice condition: choices among sure, risky and mixed options for single- and multiple-trial binary choice decision. Results for the replication sample are also shown

	N	Single-trial: sure versus risky		Multiple-trial: sure versus risky		sure versus mixed		mixed versus risky	
A. Main sample									
Gains (%)	50	38 (76)	12 (24)	33 (66)	17 (34)	31 (62)	19 (38)	38 (76)	12 (24)
Losses ()	42	19 (45)	23 (55)	14 (33)	28 (67)	11 (26)	31 (74)	15 (36)	27 (64)
B. Replication sample									
Gains ()	26	19 (73)	7 (27)	19 (73)	7 (27)	17 (65)	9 (35)	21 (81)	5 (19)
Losses ()	20	13 (65)	7 (35)	6 (30)	14 (70)	6 (30)	14 (70)	2 (10)	18 (90)

stochastic transitivity (SST), which states that if $P(a, b) > 0.5$ and $P(b, c) > 0.5$, then $P(a, c)$ must be greater than or equal to the maximum of $P(a, b)$ and $P(b, c)$. Here, P(sure, risky) is *not* greater than or equal to the maximum of P(sure, mixed) and P(mixed, risky).

The violation of SST means that the binary choice probabilities among the three options cannot be explained by a simple scalability model (Tversky & Russo, 1969), which implies that we cannot explain the pattern of choices by assuming that 1) participants are generally risk-averse for gains, and that 2) choices are made among the three alternatives based solely on ordering them along a continuum of increasing risk, *sure* < *mixed* < *risky*. Rather, the observed pattern of choice probabilities: P(sure, risky) = 0.66, P(sure, mixed) = 0.62, P(mixed, risky) = 0.76, suggests a tendency for participants to be risk-averse coupled with a slight bias in favor of the mixed option. The analogous pattern is also observed for losses, demonstrating a slight bias in favor of the mixed option in addition to the basic pattern of risk seeking in this domain.

Replication sample

These results, showing that the mixed option preference disappears (or nearly so) when decision makers face binary choices among mixed and pure options, are important enough to merit replication. Accordingly, we replicated the binary choice condition with a sample of $n = 44$ undergraduate students enrolled in a social science research methods course at a large university in the southeastern United States. Their pattern of binary choices is shown in Panel B of Table 6 ('Replication Sample'). The results for gain problems ($n = 24$) were virtually identical for this replication sample: participants tended to choose the sure gain over the risky for both single and multiple trial decision problems (73% for both versions). In the multiple-trial problems involving a mixed option, the mixed option was preferred over the risky (81–19%), but was not preferred over the sure gain (65% select the sure gain versus only 35% for the mixed). This pattern of binary choice probabilities again violates SST in a way that suggests a combination of risk aversion coupled with a slight bias in favor of the mixed option, since the probability of choosing the mixed option over the risky is *more* extreme than the probability of choosing the sure gain over the risky option. For the loss condition ($n = 20$), results for the replication sample differed somewhat from the main sample: participants tended to select the (EV-maximizing) sure loss in the single trial problems, and although their pattern of choices for the multiple-trial problems violated SST, there was no evidence for a mixed-option preference.

Discussion

The Trinary Choice condition of Experiment 3 confirmed the existence of a mixed option preference in multiple-trial decisions, but in the Binary Choice condition the mixed option was not found to be the modal choice. This pattern of results supports the interpretation of the mixed option preference found in the trinary choice conditions as a compromise effect (Simonson, 1989; Simonson & Tversky, 1992), in which the mixed option is preferred mainly because it is perceived as the 'middle option,' i.e., as a compromise between the purely risky and the purely sure-thing options. As in the demonstrations of the compromise effect by Simonson and Tversky, the present results involve choice tasks in which options with intermediate values on two dimensions are generally preferred to options with a high value on one dimension and a low value on the other dimension. Note that this characterization of our choice problems as involving two-dimensional alternatives is based on an assumption that decision makers use 'quasi-aggregation' to frame the mixed options, that is, they tend to separate the sure and risky gains (or losses), and aggregate the two types of prospect separately.

Simonson (1989) argued that the compromise effect occurs because compromise options are more justifiable. Our data on how participants justify their decisions provides some support for this explanation. In the

Trinary Choice condition, participants' justifications for choice of the mixed option in gain problems most often cited the twin motives of obtaining security and maximizing potential gains, implying that they believed the mixed option simultaneously satisfied both these motives to some degree. Thus, the mixed option tends to be justified by citing two justifications versus only a single justification for the sure and risky options.

These post-decision justifications could be also seen as supporting Lopes's (1987) multiple-goals account of decision making under risk. In this account, decision makers have drives for both *security* and *potential*; i.e., they desire both to have security (by locking in a certain gain) and to maximize potential gain (achieved in these experiments by choosing the higher-EV risky prospect). The type of risky problem used in the present studies creates conflict between these goals because one option offers security and the other offers higher potential. The mixed option, when present, offers the decision maker a chance to address both of these needs to some degree.

GENERAL DISCUSSION

The present experiments document another way in which human decisions involving financial prospects can be characterized as non-normative, at least from the standpoint of maximizing expected value. In multiple-trial decision tasks using trinary choice sets that include purely sure-thing, purely risky, and mixed options, people often select the mixed option consisting of a mix of sure and risky gains over the pure-strategy options, a strategy that cannot be optimal from the standpoint of expected value. For decisions involving gains, this mixed option preference is robust and stable across a wide variety of problems presented as trinary choice sets. For loss problems no stable preference for mixed options was seen.

The mixed option preference as a compromise effect

The pattern of results observed in Experiment 3 suggests that the mixed-option preference is an example of a compromise effect, in which the mixed option is seen as adequate on two 'dimensions,' and therefore as a compromise choice between the risky and sure-thing options. Simonson (1989) observed that one possible reason that compromise options may be preferred is that they can be viewed as more justifiable. The mixed option can be seen as more prudent (perhaps due to diversification arguments), and offers two different types of benefits: sure gains and EV-maximizing risky options (though only five 'units' of each type of benefit). Thus, it is possible that reason- or justification-based choice processes (e.g., Slovic, 1975) could be playing a role in the mixed-option advantage.

However, this explanation for the compromise effect could be criticized. If the mixed option is more justifiable in the trinary choice condition, why is it not also more justifiable in the binary choice problems? Simonson and Tversky (1992) offered another explanation for the compromise effect. For two-dimensional choice options that involved tradeoffs on two salient dimensions (e.g., options *ax, by, cz*), the extreme options may be compared to the other alternatives on Dimension 1 or Dimension 2. For example, option *ax* (compared to *by* and *cz*) involves a gain on Dimension 1 but a loss on Dimension 2. Because losses loom larger than gains, the extreme options fare badly in this comparison.

This explanation too can be evaluated in light of our data. As we have pointed out, viewing our choice problems as involving two-dimensional alternatives is based on an assumption that decision makers use 'quasi-aggregation' to frame the mixed options, i.e., they tend to separate the sure and risky gains (or losses), and aggregate the two types of prospect separately. Note that in Experiment 1 the mixed option preference still obtained (but was reduced in magnitude) when charts of the outcome distribution were presented. Because presenting these charts should have promoted *complete* aggregation of the sure-thing and risky options comprising the mixed options, this reduction in the size of the mixed option preference would be expected under Simonson and Tversky's explanation for the compromise effect.

If we accept the conclusion that the mixed option preference is a mainly due to a compromise effect, we might be motivated to ask why the effect was so much weaker or nonexistent in the loss condition. Recall that in the loss conditions of Experiments 1 and 2 the mixed option was not favored by a majority of participants (though it was still the modal choice overall), and in Experiment 3 it was not even the modal choice. It is not immediately clear why this should be so, because even in the loss domain the mixed option involves balanced tradeoffs of risky versus sure losses.

One possible explanation is that quasi-aggregation of the sure-thing and risky options comprising the mixed option may not be the dominant framing for the loss domain for hedonic editing reasons. Specifically, aggregating the mixed option as two moderate losses (one comprised of sure losses and the other of risky losses) should be *more* aversive to decision makers than total-aggregation framing. Thus, in the loss domain participants may be less prone to frame the mixed option as a compromise option involving tradeoffs on two dimensions, and less attracted to the mixed-option if they do frame it in that manner. Alternatively, Simonson's (1989) reason-based explanation for compromise effects may be relevant: in the loss domain the mixed option has two arguments *against* it (sure losses and large potential risky losses), rather than for it. This reason-based explanation could also account for why there are no demonstrations of compromise effects involving aversive stimuli elsewhere in the literature, to our knowledge.

Long-run rationality and bounded rationality

Expected Utility Theory (von Neuman & Morgenstern, 1947; Luce & Raiffa, 1957) has long been the dominant normative theory of decision making, though its descriptive validity has often been questioned (e.g., Kahneman & Tversky, 1979). Previous researchers (e.g., Lopes, 1981; Keren & Wagenaar, 1987) have suggested that decision makers are or ought to be more consistent with EU theory when making multiple-trial decisions. This idea of 'long-run rationality' also seems consistent with frequentist notions of probability that underlie many applications of the concept of expected value (EV). But as Tversky and Bar-Hillel (1983) point out, modern EU theory does not depend in any way on long-run considerations, thus long-run rationality is crucial neither to EU nor to Prospect Theory.

The present results present a strong challenge to the idea of long-run rationality. In the present experiments participants actually showed the *reverse* pattern to that predicted by long-run rationality, tending to make more EV-consistent choices for single-play than for multiple-play decisions. This finding seems explainable, at least post hoc. Many decision makers, particularly those who have received some training in calculating expected value (like those who participated here), may find it easy to calculate EV for the single-play risky options, and thus may tend to use that criterion to select an answer for single-play problems, but may not be able or willing to do the same calculations for the multiple-trial risky and mixed options. Thus for multiple-trial problems they may choose in some other way, perhaps by falling back on some simple decision heuristic, such as a naïve diversification principle.

The pattern of results does support what we might call the 'bounded long-run rationality' hypothesis, namely that participants' choices tend to more often be EV-consistent in the multiple-trial problems when the problem is made more transparent, for example by including a chart of the outcome distribution or by using a familiar and computationally easy probability such as $p = 0.5$. This pattern of findings is consistent with results previously reported in the literature (Benartzi & Thaler, 1999; Redelmeier & Tversky, 1992). In fact, the present results parallel the findings of Benartzi and Thaler (1999), who found that relatively naïve participants become somewhat more reluctant to select a risky (but higher EV) alternative in the multiple-trials case (contrary to long-run rationality), but are quite willing to select it when a chart of the outcome distribution is shown. The present results extend their findings as well, by showing that the presence of the chart also draws participants away from the compelling (but lower-EV) mixed option.

Mixed-option preferences in multiple trials and repeated trials

The present studies establishing a mixed-option preference in prospective multiple-trial choices have been restricted to decision problems in the domain of financial decision making under risk. However, as noted in the introduction, the mixed-option preference seems analogous to the phenomenon of *variety-seeking* in consumer choice (e.g., Kahn & Lehmann, 1991). In particular, consumer research (Read & Lowenstein, 1995; Simonson, 1990) has revealed a 'diversification bias' in prospective consumer decisions. Specifically, there is evidence that consumers tend to choose more diverse bundles when they select many products simultaneously (and prospectively), compared to when they choose them sequentially, at the time of use. Participants in the present experiments are also making multiple-unit simultaneous prospective choices, thus they too can be expected to select bundles with more variety (here corresponding to the mixed option), compared to the mix of choices they might show when facing repeated independent trials. This is an interesting avenue for future investigations.

ACKNOWLEDGMENTS

We are indebted to a number of colleagues, including Elke Weber, Ido Erev, and Eric Johnson, for useful discussions. We also wish to acknowledge the contribution of an anonymous reviewer, who suggested the interpretation of the mixed option preference as a compromise effect.

REFERENCES

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16*, 215–232.

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research, 30*, 961–981.

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

Benartzi, S., & Thaler, R. (2001). Naïve diversification strategies in defined contribution saving plans. *The American Economic Review, 91*, 79–98.

Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General, 121*, 177–194.

Camerer, C. F., & Ho, T.-H. (1998). EWA learning in normal-form games: probability rules, heterogeneity and time variation. *Journal of Mathematical Psychology, 42*, 305–326.

Chen, Y.-J. (2001). The role of learning in risk preferences for single- and multiple-play gambles. Unpublished Ph.D. dissertation, Teacher College, Columbia University.

Coombs, C. H., & Bowen, J. N. (1971). A test of *Ve*-theories of risk and the effect of the central limit theorem. *Acta Psychologica, 35*, 15–28.

Corter, J. E., & Chen, Y.-J. (2005). Mixed risky/risk-free strategies are (sometimes) non-optimal. Unpublished manuscript, Teachers College, Columbia University.

Edwards, W. (1954). Probability preferences among bets with differing expected values. *Psychology Journal, 66*, 349–364.

Gneezy, U., & Potters, J. (1997). An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics, 112*, 631–645.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events. *Psychological Science, 15*, 534–539.

Kahn, B., & Lehmann, D. R. (1991). Modeling choice among assortments. *Journal of Retailing, 3*, 274–298.

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

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.

Linville, P. W., & Fisher, G. W. (1991). Preferences for separating or combining events. *Journal of Personality and Social Psychology, 60*, 5–23.

Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty, *Economic Journal, 92*, 805–824.

Luce, R. D. (1959). *Individual choice behavior.* New York: John Wiley & Sons.

Luce, R. D., & Raiffa, H. (1957). *Games and decisions: Introduction and critical survey.* Oxford, England: Wiley.

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

Lopes, L. (1987). Between hope and fear: the psychology of risk. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 20, pp. 225–295). New York: Academic Press.

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

Read, D., & Loewenstein, G. (1995). Diversification bias: explaining the discrepancy in variety seeking between combined and separated choices. *Journal of Experimental Psychology: Applied, 1*, 34–49.

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

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics, 69*, 99–118.

Simonson, I. (1989). Choice based on reasons: the case of attraction and compromise effects. *Journal of Consumer Research, 13*, 411–420.

Simonson, I. (1990). The effect of purchase quantity and timing on variety seeking behavior. *Journal of Marketing Research, 27*, 150–162.

Simonson, I., & Tversky, A. (1992). Choice in context: tradeoff contrast and extremeness aversion. *Journal of Marketing Research, 29*, 281–95.

Slovic, P. (1975). Choice between equally-valued alternatives. *Journal of Experimental Psychology: Human Perception and Performance, 1*, 280–287.

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing Science, 4*, 199–214.

Thaler, R., & Johnson, E. (1990). Gambling with the house money and trying to break even: the effects of prior outcomes on risky choice. *Management Science, 36*, 643–660.

Thaler, R., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: an experimental test. *Quarterly Journal of Economics, 112*, 647–661.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review, 79*, 281–299.

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

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 1*, 297–323.

Tversky, A., & Russo, J. E. (1969). Substitutability and similarity in binary choice. *Journal of Mathematical Psychology, 6*, 1–12.

von Neuman, J., & Morgenstern, O. (1947). *The theory of games and economic behavior.* Princeton, NJ: Princeton University Press.

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.

Zeelenberg, M., van den Bos, K., van Dijk, E., & Pieters, R. (2002). The inaction effect in the psychology of regret. *Journal of Personality and Social Psychology, 82*, 314–327.

APPENDIX

Experiment 1 sample materials: Problem 1 in multiple-trial with mixed option format (with charts):
You are offered a choice between the following two options ***ten times in a row***:

Option: you have a 10% chance to receive $100; otherwise you receive nothing
Option: you receive $5

BUT, before we begin you must make your choices for each play. That is, you will actually be choosing between taking the first option on every play, the second option on every play or selecting a mix of choices. So you are really choosing between the following options:

Option **D**: five payments of $5, plus five ‘plays,’ each of which is a 10% chance to receive $100.
Option **K**: ten payments of $5
Option **H**: ten ‘plays’ of the game in which every play is a 10% chance to receive $100.

These choices will result in the following distribution of outcomes:

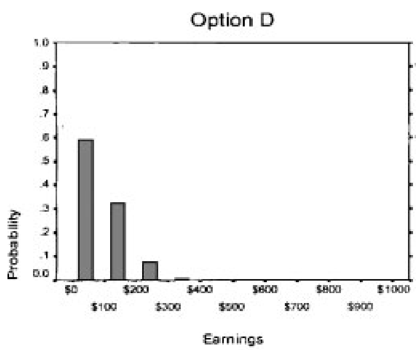

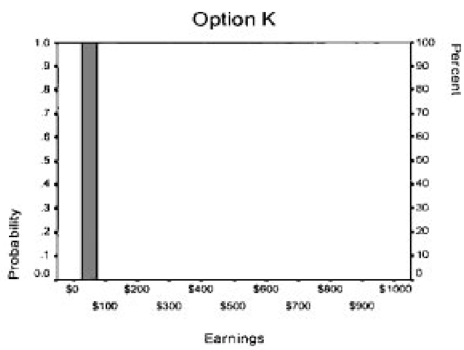

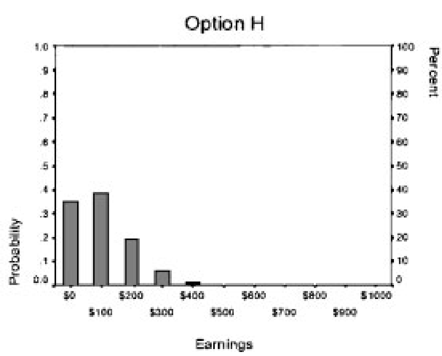

Which option do you prefer: D, K or H? ________ (fill in)

Authors’ biographies:

Yuh-Jia Chen received his PhD in the Measurement, Evaluation, and Statistics Program at Teachers College, Columbia University (with an M.A. in Organizational Psychology and an M.S. in Applied Statistics). His research interests lie in human problem solving, choice and decision making under risk, money attitudes, and resource allocation behavior. Currently, he is Assistant Professor of Psychology at Middle Tennessee State University.

James E. Corter is Associate Professor of Statistics and Education at Teachers College, Columbia University, where he is Chair of the Department of Human Development. He received his graduate training in psychology and quantitative methods at the L. L. Thurstone Psychometric Laboratory of the University of North Carolina—Chapel Hill and at Stanford University, where he obtained his PhD in 1983 under the direction of Professor Amos Tversky. His research interests include clustering and scaling methods, judgment and decision making, categorization, learning, and problem solving.

Authors' addresses:

Yuh-Jia Chen, Department of Psychology, Middle Tennessee State University, Murfreesboro, TN 37132, USA.

James E. Corter, 525 W. 120th St., rm. 453 Grace Dodge Hall, Teachers College, Columbia University, New York, NY 10027, USA.

