Southern Economic Journal 2009, 75(4), 1045–1066

The Neglect of Correlation in Allocation Decisions

Ido Kallir* and Doron Sonsino†

We study the effect of variation in correlation on investment decision in an experimental two-asset application. Comparison of allocations across problems suggests that subjects neglect probabilistic information on the joint distribution of returns and base their allocations on the observed return levels for the two assets. When asked to predict future returns, subjects try to replicate the historical distribution, thereby falling into the probability-matching bias. Predictions drastically vary when correlations become negative, while allocations are not significantly affected by changes in sign of correlation. The observed allocation patterns contradict the predictions of standard models of choice; the inconsistency is attributed to common behavioral bias in financial decision. Field implications of the results are discussed.

JEL Classification: G11, C91, D81

1. Introduction

The observation that nonperfect correlation in asset returns may increase diversification possibilities is an essential component of efficient portfolio theory (Markowitz 1952). Finance textbooks commonly demonstrate the expansion of efficient frontier as the correlation in returns of underlying assets decreases.[1] This paper deals with an experimental investigation of the effect of changes in correlation on allocation decisions in practice. We design a simple two-asset investment problem where changes in correlation must affect optimal allocations in a predetermined, intuitive direction. More than 140 subjects with advanced backgrounds in finance are asked to predict future returns and allocate funds, in several distinct problems, for different levels of correlation. The experiments reveal that subjects consistently neglect the distributional data in their allocation decisions; although, predictions are adapted to changes in correlation. While the results for the prediction tasks demonstrate that subjects recognize differences in experimental correlation, the observed allocations suggest they fail to incorporate

* Graduate School of Business, The College of Management, 7 Rabin Boulevard, P.O. Box 9017, Rishon LeZion 75190 Israel; E-mail ikallir@colman.ac.il.

† Graduate School of Business, The College of Management, 7 Rabin Boulevard, P.O. Box 9017, Rishon LeZion 75190 Israel; E-mail sonsinod@colman.ac.il; corresponding author.

We thank the research authority at the College of Management for financial support. We thank participants at the Instituto de Estudios Sociales Avanzados (IESA) meeting on Experimental and Behavioral Economics in Cordba, Spain; the Economic Science Association (ESA) meeting in Alessandria, Italy; and the decision-group seminar in Ben-Gurion University for comments on preliminary versions of this study. Particular thanks to Ofer Azar, Rachel Barkan, Uri Benzion, Werner Güth, David Leiser, Joachim Meyer, Lau Morten, Marcus Nöth, Bradely Raffle, and two anonymous referees.

Received December 2006; accepted February 2008.

[1] See, for example, William Goetzmann's e-book "An Introduction to Investment Theory," available at http://viking.som.yale.edu/will/finman540/classnotes/notes.html.

such variation into their investment decisions. Our simple two-asset design moreover allows for direct derivation of formal results regarding the optimal allocation patterns of rational investors. The experimental allocations consistently contradict formal theory; the inconsistency can be attributed to documented bias in allocation of funds (Benartzi and Thaler 2001) and perception of probability (Kahneman et al. 1982).

In each experimental allocation problem, subjects are asked to distribute funds between two "virtual" assets, A and B, which admit only two levels of return: high or low. The marginal distribution of return on A always dominates the corresponding distribution on B, and the only rationale for allocating funds to the second asset (B) follows from the possibility that realized return on A would be low when the return on B is high; that is, from nonperfect correlation in returns. Different versions of this problem are presented to more than 140 master of business administration (MBA) and business undergraduate students in three distinct versions of the experiment. In the first two versions, subjects are asked to allocate funds in five different problems where the correlation in asset returns takes the values +2/3, +1/3, 0, $-1/3$, and $-2/3$. The high/low return levels on each asset are randomly drawn for every subject and each problem to examine the effect of changes in correlation in general for different combinations of underlying returns. The random drawing of returns for different problems is also intended to preclude schematic response to the distributional stimuli. Information on the joint distribution of returns is present in the form of empirical frequencies for the 12 preceding periods. Subjects are requested to predict returns for four additional periods under the assumption that future returns are sampled from the empirical distribution.

The experiment was designed to investigate three specific issues/hypotheses: (i) test the rationality of predictions; (ii) examine if subjects increase the allocations to the dominated asset B as the correlation in returns decreases; and (iii) check the compatibility of allocations with rational models of choice, such as expected utility (EU) and rank-dependent utility (RDU).[2] The results reveal that subjects with advanced backgrounds in finance violate the rational hypotheses in all three domains: subjects try to replicate the empirical distributions when filling in their predictions, while they should fill in the most probable return combination in order to maximize the probability of correct forecasting. In this sense, subjects exhibit the well-known probability-matching, or "law of small numbers," bias (Tversky and Kahneman 1971).

The allocation to the dominated asset B does not increase significantly as the level of correlation with A decreases; although, the results for the prediction tasks suggest that subjects "recognize" shifts in correlation. Actual allocation decisions are mostly affected by the magnitude of returns on A and B, rather than being affected by correlation levels. In this sense, subjects neglect the joint distribution of returns in their allocation decisions.

Finally, we prove that more than 50% of the subjects violate conditions required for consistency with EU or RDU maximization by allocating too much of the funds to the dominated asset B. The inconsistency is attributed to either a common bias in allocation

[2] Quiggin's (1982) rank-dependent utility model is considered one of the prominent alternatives to von Neuman Morggenstern's (1944) expected utility. The model has recently been applied to various financial applications (see, for example, Aizenman 1997; Haliassos and Hassapis 2001; Bleichrodt and Eeckhoudt 2005; Polkovnichenko 2005) and shown to provide insights that cannot be captured with EU maximization.

patterns (naïve diversification) or perception of probability (the overweighting of salient, representative events).

To further examine the robustness of results we run a third version of the experiment where the joint distribution of returns was presented in a multicolored pie-chart (in addition to the numerical table), and the number of problems in each questionnaire was decreased from five to two. The experiment was run in advanced MBA seminars in finance. The allocations still did not vary with level of correlation. The neglect of distributional data therefore reappears in a focused "shorter" experiment with modified instructions.

Kroll et al. (1988) ran an experimental examination of the Capital Asset Pricing Model (CAPM), where subjects are asked to allocate funds between three risky assets with normally distributed returns. Between-subject comparisons suggest that allocations are not affected by the correlation in returns (of two assets), even after 20 successive investment rounds. Anderson and Settle (1996) more recently demonstrate that subjects are insensitive to distributional characteristics when creating portfolios from field 1-year and 10-year stock data. Benzion et al. (2004) let subjects diversify between a bond, a stock, and an option and found that inefficient naïve allocations persist even after 40 investment rounds where subjects receive feedback on their realized payoff at the end of each period. Our current experiments test the effect of correlation on investment, within-subject, in a simpler two-asset binary-return design where the effect of decrease in correlation on optimal allocation is intuitively transparent. Moreover, we run a multi-task experiment where subjects are asked to predict returns and allocate funds concurrently. The results demonstrate that even when subjects comprehend the difference between positive and negative correlations (differences are reflected in the prediction tasks), they fail to incorporate such variation into their allocation decisions.

The nonprofessional financial media is frequently engaged with analyzing mean returns, trends, and cycles in the stock market. Data on correlation, volatility, and higher moments of historical returns are rarely examined. Our experiments may suggest that the media's emphasis on average data is complemented by low demand of potential investors for higher-level distributional information. Higher moments and joint-distribution statistics are harder to process or digest. The investors therefore naturally focus on "salient" mean return statistics while ignoring finer properties of underling distributions (Reyna 2004; discussed further in section 7).

The paper is organized as follows: Section 2 presents the main experiment (Versions 1 and 2), the results for the prediction tasks are discussed in section 3, and the allocations are examined in section 4. Version 3 of the experiment is discussed in section 5, section 6 analyzes the allocation problem of an expected utility maximizer and reexamines the data in light of the model (a generalization to RDU is provided in the supplementary Appendix), and section 7 concludes.

2. The Main Experiments

Typical Problem Format

The experiment consisted of five prediction/allocation problems. The format of a typical problem is illustrated in Appendix A. Subjects receive tabulated information on the empirical

Table 1. Frequency of Realization Types for Each Problem

	hh	*hl*	*lh*	*ll*	correlation
P1	5/12	1/12	1/12	5/12	+2/3
P2	4/12	2/12	2/12	4/12	+1/3
P3	3/12	3/12	3/12	3/12	0
P4	2/12	4/12	4/12	2/12	−1/3
P5	1/12	5/12	5/12	1/12	−2/3

distribution of returns on virtual assets A and B in a sample of 12 observations. The returns on each virtual asset take only two possible values: high or low. The joint distribution table thus refers to four possible scenarios: the case where both assets earn their high return level, the case where both assets pay their low return level, and the two cases where one asset gains high and the second asset pays low. Subjects are told (see instructions in Appendix B) that the tabulated data describes the joint distribution of returns on A and B and are asked to assume that a sample of five additional observations (observations 13–17) was drawn from the tabulated distribution to compose the experimental tasks. The first experimental task dealt with predicting the return for one asset from the realized return on the second asset for observations 13–16 (see the prediction table in Appendix A). The asset with concealed return was randomly selected for each prediction task (and each questionnaire) and could vary across the four prediction periods. The prediction task was followed by an allocation task, where subjects allocated a fixed endowment between the two assets under the assumption that payoff will be determined by the realized returns in observation 17.[3] The instructions explained that at the end of the experiment one of the five problems and one prediction period would be randomly selected for each subject. Correct prediction in the selected prediction task paid the subject 20 New Israeli Shekel (N.I.S.).[4] In addition, subjects received a payoff that depended on the realized return of their portfolio in the selected problem (see details in subsection below).

The Allocation/Prediction Problems

Let *hh* denote the realization where both assets score the high return level; *ll* will similarly denote the case where the level of return on both assets is low; *hl* (*lh*) will then describe the cases where asset A (B) scores the high level of return while the other asset scores low. The five decision problems used in the experiment differ in the empirical frequency of the four return combinations. All five problems are symmetric in the sense that the frequency of *hh* is equal to the frequency of *ll*, and the frequency of *hl* equals the frequency of *lh*. The specific frequencies for each problem (P1–P5) are described in Table 1.

Let $h(i)$ denote the high return level on asset i; $l(i)$ will similarly denote the low return level on i, for i = A or B. Note that in all five problems asset i pays $h(i)$ or $l(i)$ with equal (marginal) probability 0.5. A simple argument (see Appendix C, Claim C.1) shows that the Pearson coefficient of correlation between the return on A and the return on B for problem Pi is

[3] No restrictions were imposed on the allocations. Subjects could thus allocate all their funds to one of the assets; they could also short one asset and accordingly increase the allocation to the second asset. Information on the frequency of cases where subjects allocated all their endowment to one asset is provided in footnote 16. Only one participant has short asset B to buy A in one of the problems. The questionnaire of this subject was otherwise incomplete, and he was removed from the sample.

[4] The rate of the New Israeli Shekel at the time of the experiment was about 4.5 N.I.S. for one U.S. dollar.

independent of the specific high/low return levels on the two assets. The correlation levels for problems P1 and P2 are +2/3 and +1/3, respectively, while the correlations for problems P4 and P5 are −1/3 and −2/3, respectively. In P3 the returns on the two assets are statistically independent and the correlation is 0 (see the right column on Table 1).

Return Levels

Each decision problem is characterized by four return levels: $h(\mathrm{A})$, $h(\mathrm{B})$, $l(\mathrm{A})$, and $l(\mathrm{B})$. To avoid schematic anchoring of allocations across the five problems, we randomly select the four return levels for each decision problem and each subject.[5] The four return levels thus vary across problems (for each subject) and are diverse across subjects (for each problem). The specific numbers are selected so that A always dominates B in marginal distribution. In particular, we choose the return levels so that $h(\mathrm{A}) \geq h(\mathrm{B}) > l(\mathrm{A}) > l(\mathrm{B})$ in each decision problem. This scheme is chosen to control subjects' incentives to allocate funds to asset B. The only rationale to invest in B in this setting follows from the possibility that its realized return may be high when the return on A is low (the *lh* combination). In all other return mixtures (*hh*, *ll*, *hl*) the return on A is higher than the return on B. A potential problem with such a design is that subjects might not allocate any funds to the marginally dominated asset B. We therefore keep the difference $h(\mathrm{B}) - l(\mathrm{A})$ "large" to increase incentives to invest in B for risk diversification.[6] Specifically, the four return levels are chosen by the following scheme: $h(\mathrm{A})$ is randomly selected from the set {11%, 12%, 13%, 14%, 15%}; to determine $h(\mathrm{B})$ we randomly subtract one of the numbers in the set {0%, 1%, 2%, 3%} from $h(\mathrm{A})$; $l(\mathrm{A})$ is independently selected from the set {2%, 3%, 4%}; $l(\mathrm{B})$ is finally determined by subtracting one of the numbers {1%, 2%} from $l(\mathrm{A})$. Note that this scheme generates 120 combinations of return levels; the difference $h(\mathrm{B}) - l(\mathrm{A})$ varies from 4–13%.[7]

Subjects and Method

We run two different versions of the experiment (Version 1 and Version 2). Version 1 is run on 49 MBA students in Tel-Aviv University and at the College of Management.[8] All the participants are advanced MBA students attending core courses in finance and investments. The endowment used with this group is 1000 N.I.S. Subjects receive the actual return on their portfolio in the randomly selected allocation problem, plus the 20 N.I.S. payoff for correct prediction (see subsection 2.1). The translated instructions for the MBA group are provided in Appendix B; we henceforth refer to this group as "MBA." The experiment on MBA is run in class (at the last 30 minutes of a 90-minute class). The printed instructions are first presented on

---

[5] An alternative design where the returns are fixed across the five problems for each subject could produce the anticipated "monotonic" response; however, this design is not interesting since it invites schematic response to changes in frequency of *lh* realizations.

[6] To obtain a large difference between $h(\mathrm{B})$ and $l(\mathrm{A})$, we increase the distance between $h(\mathrm{A})$ and $l(\mathrm{A})$ while decreasing the distance between $h(\mathrm{A})$ and $h(\mathrm{B})$.

[7] We confine the study to positive returns, including negative returns would complicate the experiment (e.g., since subjects can lose some of their endowment, the direct payoff method must be changed) and some parts of the analysis (cumulative prospect theory, the generalization of rank-dependent utility for cases where payoffs can be negative, should be examined).

[8] Thirty-five of the MBA students are from Tel-Aviv University and 14 students from the College of Management; we did not find significant differences between the two subgroups.

the blackboard. Subjects are then asked to complete the five prediction/allocation assignments at their own pace.

The second version of the experiment is run on 67 business undergraduate students at the College of Management. Students attending an advanced course in finance are invited to prescheduled experimental sessions. The endowment in version 2 is 100,000 N.I.S., but the instructions explain that subjects would only receive 1% of the payoff on a randomly selected portfolio (in addition to the correct prediction bonus). The expected payoff in Version 2 is thus similar to the expected payoff in Version 1. The instructions for Version 2 include a paragraph (that did not appear in Version 1) as follows: "*The purpose of the experiment is to compare your investment patterns in five different cases. The data on the empirical joint distribution of returns varies across the different problems. We ask you to examine the data for the five problems before you complete the assignments for each problem*." The instructions are otherwise similar to those used with the MBA. We henceforth refer to this experiment as "BUS" (for business undergraduates).

The order in which the five problems appear in each questionnaire is randomly selected for each subject in both versions of the experiment.[9] No time constraints are imposed; most subjects take 20–30 minutes to complete the five prediction/allocation tasks. The average actual payment per subject is 75 N.I.S. (about 16.7 U.S. dollars); the maximal payment is 144 N.I.S. ($32).

3. Predictions

Recall that the data for the prediction tasks are generated by sampling four observations from the joint distribution table. The procedure is disclosed to the subjects in the written instructions (see Appendix B) and emphasized in the oral presentation. The random sampling of the additional observations implies that the probability of correct predictions is maximized by filling in the most likely return combination for each prediction task. In P1, for example, the *hh*/*ll* return combinations are five times more likely than the *hl*/*lh* mixtures. Rational subjects should thus fill in the "high" return level for asset B (A) when the prediction table reveals that the return on A (B) was "high," and similarly fill in the "low" level of return when the table reveals that the return on the other asset was "low." The same prediction rule should be applied in each of the four prediction tasks (13–16) in order to maximize the overall probability of correct prediction. We accordingly say that a given subject is rational in P1 if the prediction table of the subject for P1 only consists of *hh* and *ll* combinations. The same reasoning applies to P2: Since the *hh*/*ll* realizations are twice as likely as the *hl*/*lh* cases in this problem, predictions are classified as rational if the prediction table does not include any *hl*/*lh* mixtures. In problems P4 and P5, however, the *hl*/*lh* realizations become more frequent than the *hh*/*ll* cases. Rational prediction thus implies that subjects should fill in *l* for *h* and *h* for *l*, and the prediction tables for these problems should only consist of *hl*/*lh* realizations. Note that since P3 assigns equal weight to all four realization types, the distinction between rational and irrational prediction is irrelevant for this problem.

[9] The order in which the *hh*, *ll*, *hl*, *lh* realizations appear in the empirical distribution table is also randomly assigned; asset A appears at the left column for some subjects while appearing at the right column for others (see Appendix B).

Table 2. Distribution of $\#RP$

	$\#RP = 0$	$\#RP = 1$	$\#RP = 2$	$\#RP = 3$	$\#RP = 4$
MBA ($N = 49$)	30.61%	8.16%	22.45%	10.20%	28.57%
BUS ($N = 67$)	55.22%	22.39%	7.46%	7.46%	7.46%

Table 3. Frequency of *hh*/*ll* Combinations in Prediction Tables

	Frequency of *hh* and *ll* Predictions in MBA	Frequency of *hh* and *ll* Predictions in BUS	Frequency of *hh* and *ll* in Joint-Distribution Data
P1	77.55% (21/28)	**65.67% (47/20)**	83.33%
P2	73.98% (20/29)	65.30% (34/33)	66.67%
P3	47.96% (8/6)	51.87% (12/13)	50.00%
P4	26.53% (30/19)	41.79% (26/41)	33.33%
P5	26.02% (24/25)	**39.55% (14/53)**	16.67%

The smaller brackets provide the number of subjects with strictly lower/strictly higher frequency of *hh*/*ll* predictions (compared with the corresponding joint-distribution frequency). The bolded boxes represent the two cases where the matching hypothesis is rejected at $p = 0.05$.

To summarize the data, let $\#RP$ denote the number of problems for which the subject confirmed with the rational prediction hypothesis (in all four tasks); $\#RP = 4$ then describes an all-rational subject who consistently follows the rational prediction rule in all four problems (P1–P2, P4–P5); $\#RP = 0$, on the other extreme, describes subjects who make at least one irrational prediction in each of the four problems. Table 2 presents the distribution of $\#RP$ in the two experiments.

Less than 30% (14 subjects) of the MBA students always follow the rational prediction rule in all four problems. The proportion of all-rational prediction in the BUS pool is significantly lower: only 5 of the 67 subjects. More than 30% of the MBA students and 55% of the BUS subjects do not follow the rational prediction rule in all four prediction tasks. A closer look at the individual prediction tables suggests that many subjects, in both versions of the experiment, act as though they are trying to "match" or reproduce the "empirical" joint distribution frequencies in the prediction tables.[10] Probability matching is most apparent in P3, where 61.22% of the MBA subjects and 59.7% of BUS students fill in exactly one return combination of each type (*hh*, *ll*, *hl*, *lh*) in their prediction tables. "Exact matching" of this type is not possible in the other problems because (i) the empirical frequencies cannot be exactly reproduced in a four-period prediction table,[11] and (ii) the hidden return (A or B) for each prediction period is randomly selected.[12] Therefore, in Table 3 we compare the frequency of *hh*/*ll* combinations in the individual prediction tables to the underlying joint-distribution frequency.

The first row of Table 3, for example, reveals that the frequency of *hh*/*ll* combinations in the individual prediction tables for P1 is 77.55% for the MBA and 65.67% for the BUS. This is

[10] See Vulkan (2000) for a comprehensive survey on probability matching in economics.

[11] When the empirical frequency of the *hl*-realizations is 2/12, for example, subjects can approximate the empirical frequency from below (by constructing a prediction table with no *hl* cases) or from above (by constructing a prediction table with one *hl* realization), but they cannot exactly reproduce the 2/12 frequency.

[12] Some subjects, for example, are asked to fill in a prediction table with four successive *h* realizations for one asset and thus cannot produce any predictions where the return on that asset is low.

Table 4. Mean Allocation to Dominant Asset A

	P1	P2	P3	P4	P5
MBA ($N = 49$)	738 (200)	710 (200)	721 (185)	715 (196)	715 (187)
BUS ($N = 67$)	62,082 (21,673)	63,017 (21,296)	59,878 (21,720)	60,317 (19,367)	57,783 (23,631)

Numbers in parentheses are standard deviations.

compared to the empirical frequency of *hh/ll* cases in the P1 data, 10/12 = 83.33%. A sign test is used to test the hypothesis that "approximations of the empirical frequency from above" are as likely as "approximations of the empirical frequency from below" (*the matching hypothesis*).[13] Similar tests are applied to problems P2–P5. The matching hypothesis is not rejected for the MBA, but it is rejected twice (at $p = 0.05$), in problems P1 and P5, for the BUS. A closer look at the BUS data suggests that the undergraduates are especially reluctant to build "nondiversified" (Rubinstein 2002) prediction tables with four or zero *hh/ll* combinations; this generates significant "overdiversification" in these two problems.[14]

4. Allocations

Recall that asset B could produce a higher return than A only when the return on A was low and the return on B was high. Since the probability of such *lh* combinations increases with the index number of the problem, we conjecture that the amount allocated to B would increase with the problem number. We term this hypothesis "*monotonicity*."[15]

Note that because of the random drawing of returns for each problem, monotonicity is not expected to hold at the individual level. Subjects, for instance, may choose to allocate a larger amount to asset A in P5 compared to P1 when the differences $h(\mathrm{A}) - h(\mathrm{B})$ and $l(\mathrm{A}) - l(\mathrm{B})$ are larger in P5. We still anticipated that the effect of differences in returns on individual allocations would cancel out on aggregate and the allocations to asset A would decrease with the correlation in returns. The experimental data (see Table 4 and the histograms in Figure 1), however, do not reflect the anticipated trends.

While the mean allocation to asset A in P1 is larger than the mean allocation to A in P5 (in both groups), statistical tests suggest that the differences are far from being significant (signed-rank Wilcoxon tests; $p = 0.4206$ for MBA; $p = 0.3762$ for BUS). Subjects' allocations when the correlation is +2/3 are thus similar to the allocations when the correlation is $-2/3$. Also, the trends across other problems often contradict the anticipated monotonicity. The mean

[13] Consider P1 for the MBA: 21 subjects produce prediction tables with *hh/ll* frequency lower than 10/12 (the empirical benchmark); 28 subjects, on the other hand, produce prediction tables with *hh/ll* frequency higher than 10/12. A sign test cannot reject the hypothesis that approximations from below are as likely as approximations from above (p = 0.3916, two-tails).

[14] Detailed comparison with Rubinstein's results is relegated to the concluding discussion.

[15] In section 6 we show that when the returns levels $h(\mathrm{A})$, $l(\mathrm{A})$, $h(\mathrm{B})$, $l(\mathrm{B})$ are fixed across the five problems, subjects who maximize EU or RDU should satisfy monotonicity. In this sense, the hypothesis is implied by standard models of choice.

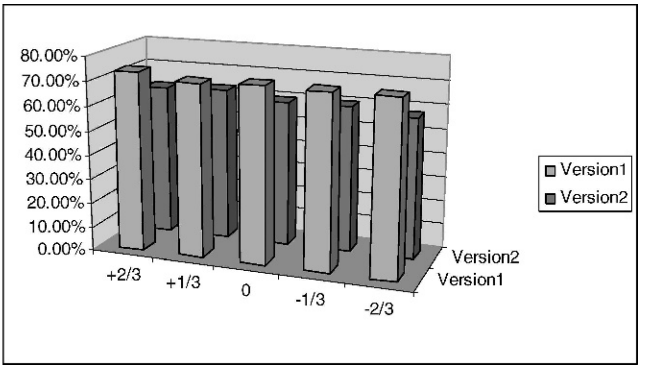

Figure 1. Proportion Allocated to Asset A

Table 5. Linear Regressions

	h(A)	l(A)	h(B)	l(B)	$p(lh)$
MBA (N = 245)	6.27	9.37	−2.46	−8.34	0.1015
	(t = 5.7)	(t = 3.6)	(t = −2.2)	(t = −3.4)	(t = 0.32)
BUS (N = 335)	7.62	7.15	−4.36	−5.71	−0.045
	(t = 7.22)	(t = 2.83)	(t = −4.23)	(t = −2.5)	(t = −0.44)

The R^2 levels are 0.93 for the MBA group and 0.89 for the BUS group; adding a constant to the model does not give a significant coefficient.

allocation of the MBA to asset A in P2, for example, is lower than the mean allocation to A in P3, P4, and P5; the differences, however, are slight and far from being statistically significant.[16]

A possible explanation to the surprisingly weak response to changes in the level of correlation across problems might be that subjects concentrate on the return levels on the two assets, thereby almost ignoring the distributional data when contemplating their allocations. To check this explanation, we run regressions where the proportion of endowment allocated to asset A depends on five variables: (i) the returns on A and B: h(A), l(A), h(B), l(B); and (ii) the frequency of lh combinations in the joint distribution data: $p(lh)$. The results of linear regressions for the two versions of the experiment are summarized in Table 5.

The estimations suggest that an increase of 1% in h(A) increases the allocation to A by 6.27% of endowment in MBA and 7.62% of endowment in BUS. Both coefficients are statistically significant at $p < 0.01$. An increase in l(A) has a similar positive, significant effect on the allocation to the dominant asset. The returns on B, h(B), and l(B) have weaker, yet significant, negative effects on the allocation to A (see the coefficients in Table 5). The coefficient of $p(lh)$, however, is not statistically significant in either pool (t = 0.327 for the MBA; t = −0.445 for the BUS). In this sense, subjects neglect the probability distribution of returns in their allocations. Recall (from Table 3) that subjects' predictions for P1 (P2) are drastically different from their predictions for P5 (P4).[17] The multi-task experiment thus

[16] About 9% of the MBA allocations and 13% of the BUS allocations assigned all the funds to the dominant asset A. The proportion of extreme allocations was similar in all five problems. Removal of the extreme allocations does not change the data significantly; the differences between the allocations to A in P1 and P5 remain insignificant for both groups.

[17] Sign tests for comparing the frequency of hh/ll combinations in the prediction tables for P1 to the corresponding frequencies in P5 confirm that the hh/ll cases are significantly more frequent in P1 ($p < 0.001$ for MBA and for BUS). Similar results are obtained in comparison of individual predictions in P2 and P4.

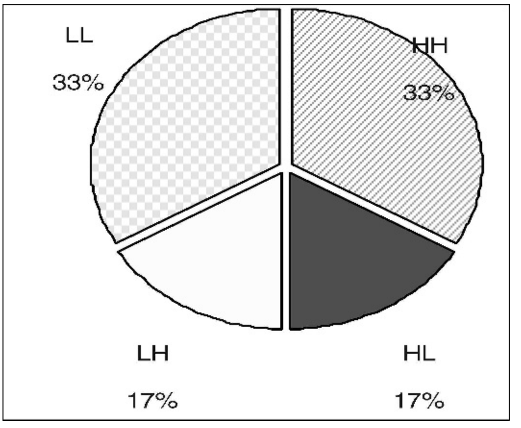

Figure 2. Pie-Chart for Version 3

demonstrates that subjects fail to incorporate changes from positive to negative correlation in their allocation decisions; although, they recognize the differences in concurrent prediction tasks.

Finally, it is interesting to note that the proportion allocated to asset A by the MBA students is more than 10% larger than the corresponding proportion for the BUS. Our preferred explanation is that subjects act more risk-aversively in the BUS version where the endowment is one thousand times larger. The inclination to act more risk-aversively when larger amounts of money are involved has been documented in Bosch-Domènech and Silvestre (1999), Holt and Laury (2002), and most recently in Harrison et al. (2005).[18] Note, however, that the two versions of our experiment did not control for the difference in endowments. The difference in allocation levels could therefore arise from other factors (e.g., the difference in subject pool or the modification of instructions).

5. Experiment III

To further examine the robustness of results, we run a third version of the experiment where each subject fills in only two prediction/allocation problems. The design and instructions are similar to those used in the preceding versions with slight modifications. The joint distribution of returns is illustrated in a multicolored pie-chart using the *hh*, *hl*, *lh*, *ll* notation to describe each realization (see Figure 2 for an illustration). The pie-chart is appended to the return-table used in preceding versions in order to highlight the difference in distributions across the two problems in each questionnaire. Also, the following paragraph is added to the instructions (in bolded font): "*The goal of the experiment is to compare your allocation in two different cases. The levels of return and joint distribution of returns vary across problems. We ask*

[18] The difference in endowments between the two versions of our experiment is "hypothetical" since the BUS subjects only received 1% of their payoff. In Holt and Laury (2002), subjects did not respond to changes in scale of payoffs when choices were hypothetical. Harrison et al. (2005) remark that this may be explained by subjects' attempts to economize on cognitive effort and anchor on previous results when consequences are hypothetical. The current results indeed suggest that a significant hypothetical payoff scale effect on risk preferences may arise in between-subject comparisons.

Table 6. Predictions and Allocations in Version 3 vs. Version 1

	P1 (+2/3)	P2 (+1/3)	P4 (−1/3)	P5 (−2/3)
Proportion of *hh/ll* predictions	78.3% (77.5%)	70.3% (74.0%)	23.3% (26.5%)	20.3% (26.0%)
Mean allocation to asset A	816.7 (738)	628.1 (710)	696.7 (715)	659.7 (715)

Numbers in parentheses represent the predictions and allocations from Version 1.

you to inspect the data (and pie-charts) for both problems before filling-in the prediction/allocation assignments for each problem."

The experiment is run on two groups of MBA students attending advanced seminars in finance. The first group ($N = 15$) receives the +2/3 (P1) and −1/3 (P4) problems while the second group ($N = 16$) fills in the −2/3 (P5) and +1/3 (P2) problems. To increase control, we fix the levels of returns in each problem (for all subjects). The returns in the problems with correlation +2/3 or −2/3 are fixed at $h(\mathrm{A}) = 15\%$, $l(\mathrm{A}) = 2\%$, $h(\mathrm{B}) = 12\%$, and $l(\mathrm{B}) = 0\%$, while the return levels in the +1/3 and −1/3 problems are $h(\mathrm{A}) = 11\%$, $l(\mathrm{A}) = 3\%$, $h(\mathrm{B}) = 9\%$, and $l(\mathrm{B}) = 1\%$.[19] The main results of the additional experiment (henceforth Version 3) are summarized in Table 6. For comparison, we represent the corresponding results for the MBA experiment (Version 1) in brackets.[20]

First, we briefly examine the results for the prediction tasks. The predictions of 11 of the 31 subjects (35.5%) are all-rational in both problems ($\#RP = 2$); 15 subjects (48.4%), on the other hand, mix the *hh/ll* and *hl/lh* combinations in both problems ($\#RP = 0$). The first row of Table 6 presents the frequency of *hh/ll* return combinations in the prediction tables for each problem. The proportions are similar to those observed in Version 1 of the experiment (see the bracketed numbers). The results for the prediction tasks, in conclusion, are not significantly affected by the changes in experimental design. Subjects still exhibit the probability-matching bias in the modified two-problem design, just as in Version 1 of the experiment. The frequency of *hh/ll* combinations in P1 (P2) is therefore significantly higher than the frequency of these combinations in P5 (P4) (sign tests, $p < 0.01$ in both comparisons).

The mean allocations to the dominant asset A in Version 3 and Version 1 are contrasted in the second row of Table 6. The results for Version 3 appear stronger in comparison of the +2/3 and −2/3 allocations. Between-subject comparison confirms that the allocation to A when the correlation is +2/3 (mean: 816.7) is significantly larger than the allocation to A in the −2/3 problem (mean: 659.7) (Mann-Whitney test for comparison of independent samples, $p = 0.03$). The mean allocation to asset A in the +1/3 problem, however, is slightly smaller than the mean allocation to A in the −1/3 problem (Mann-Whitney test, $p = 0.24$; nonsignificant). Within-sample comparisons reveal that only 12 subjects (of 31) decrease the allocation to A in the problem with negative correlation, 11 subjects increase the allocation to A in the problem with negative correlation, and eight subjects choose the same allocation in both problems. The hypothesis that subjects are as likely to increase or decrease the allocation to A as the level of

[19] Note that $h(\mathrm{B}) - l(\mathrm{A}) = 10\%$ in the 2/3-correlation problems, while $h(\mathrm{B}) - l(\mathrm{A}) = 6\%$ in the 1/3-correlation problems; by this we attempted to make asset B relatively more attractive in the 2/3 problems. We therefore anticipated that the increase in allocation to asset B (with decrease in correlation) would be stronger in the P2/P5 group. This conjecture is contradicted by the data.

[20] The corresponding data for the BUS can be traced in Tables 3 and 4; the comparison with MBA/Version 1 is more relevant because the type of subjects (advanced MBA students in Versions 1 and 3 vs. undergraduates in BUS) and allocation budgets (1000 in Versions 1 and 3 vs. 100,000 in BUS) are similar in Versions 1 and 3 but different in Version 2.

correlation decreases is not rejected for either group (sign test significance: $p = 0.38$ for the P1/P4 group and $p = 0.22$ for the P2/P5 group). The changes in experimental design, in conclusion, induce larger significant differences between the +2/3 and −2/3 allocations; but individual-level comparisons still suggest that most of the subjects (in both groups) do not increase the allocation to B as the frequency of *lh* observations increases. Subject-level neglect of probabilistic information (in the allocation tasks) persists.

6. Violation of Formal Choice Theory

We briefly examine the optimal allocation problem of a representative risk-averse expected-utility maximizer Decision Maker (DM) and argue that subjects contradict theory in all three versions of the experiment.

Let $U[x]$ denote the von Neumann Morgenstern utility function of DM and assume it is twice continuously differentiable, strictly monotone ($U' > 0$), and strictly concave ($U'' < 0$). Use $P(xy)$ to denote the probability of the return combination: x on asset A and y on asset B; $\pi_q(xy)$ is similarly used to denote the final wealth when the amount allocated to asset A is q and the realized returns are x (on A) and y (on B). That is, $\pi_q(xy) = q \cdot (1 + x) + (e - q) \cdot (1 + y)$, where e denotes the endowment for allocation. Using this notation, formulate DM allocation problem as follows:

$$MAX_q \; P(hh)\cdot U\left[\pi_q(hh)\right] + P(hl)\cdot U\left[\pi_q(hl)\right] + P(lh)\cdot U\left[\pi_q(lh)\right] + P(ll)\cdot U\left[\pi_q(ll)\right]. \qquad (*)$$

By twice differentiating the objective in (*) with respect to q, it is easily verified that the function is strictly concave, which implies a unique solution. It is not difficult, however, to construct numeric examples where (for specific preferences) the optimal allocation to asset A in (*), q^*, is larger than the endowment e so that DM should short B for buying A at the optimum. Since our subjects do not short B for A (but choose $q^* = e$ in about 10% of the cases), we henceforth restrict the analysis to the case where DM solves the constrained allocation problem, where $0 \leq q \leq e$.

$$MAX_{0\leq q\leq e} \; P(hh)\cdot U\left[\pi_q(hh)\right] + P(hl)\cdot U\left[\pi_q(hl)\right] + P(lh)\cdot U\left[\pi_q(lh)\right] + P(ll)\cdot U\left[\pi_q(ll)\right]. \qquad (1)$$

Note that Equation 1 has a unique solution since we maximize a strictly concave function on a closed interval.[21] Using $(x(\mathrm{A}) - y(\mathrm{B}))$ to denote the derivative of $\pi_q(xy)$ with respect to q, write the first-order conditions for an interior ($0 < q^* < e$) solution as follows:

$$\begin{aligned} &P(hh)\cdot(h(\mathrm{A}) - h(\mathrm{B}))\cdot U'\left[\pi_{q^*}(hh)\right] + P(hl)\cdot(h(\mathrm{A}) - l(\mathrm{B}))\cdot U'\left[\pi_{q^*}(hl)\right] \\ &+P(lh)\cdot(l(\mathrm{A}) - h(\mathrm{B}))\cdot U'\left[\pi_{q^*}(lh)\right] + P(ll)\cdot(l(\mathrm{A}) - l(\mathrm{B}))\cdot U'\left[\pi_{q^*}(ll)\right] = 0. \end{aligned} \qquad (2)$$

Note, however, that the left-hand-side of Equation 2 can be positive (negative) when the optimal solution is $q^* = e$ ($q^* = 0$).[22]

In the following claims, $q^*_j(\bar{h},\bar{l},\underline{h},\underline{l})$ denotes the optimal solution to Equation 1 of the allocation problem Pj when the returns on the two assets are $h(\mathrm{A}) = \bar{h}$, $l(\mathrm{A}) = \bar{l}$, $h(\mathrm{B}) = \underline{h}$, $l(\mathrm{B})$

[21] It is easily verified that when $U'' \geq 0$, the optimal allocation is always at the point where DM allocates all the endowment to asset A.

[22] $q^* = 0$ is ruled out in Claim 2.

Table 7. Violations of Weak Condition for Consistency with EU (RDU)

	At Least One Problem	P1	P2	P3	P4	P5
MBA (N = 49)	40.8% (36.73%)	16.33% (16.33%)	12.24% (12.24%)	10.2% (8.16%)	10.2% (8.16%)	8.16% (8.16%)
BUS (N = 67)	83.6% (77.61%)	31.34% (26.86%)	29.85% (26.86%)	32.84% (25.37%)	31.34% (25.37%)	38.8% (35.82%)
Version 3 (N = 15 or N = 16)*	41.9% (19.3%)	0% (0%)	43.7% (12.5%)	NA	20% (6.7%)	18.7% (18.7%)

In Version 3, N = 16 for P2 and P5, and N = 15 for P1 and P4.

$= \underline{l}$. We present the claims before discussing their interpretation; the proofs are relegated to Appendix C.[23]

Claim 1: Monotonicity. $q^*_{j+1}(\bar{h},\bar{l},\underline{h},\underline{l}) \leq q^*_j(\bar{h},\bar{l},\underline{h},\underline{l})$ for every j = 1, 2, 3, 4. Moreover, the inequality is strict when $q^*_j(\bar{h},\bar{l},\underline{h},\underline{l}) < e$.

Claim 2: Weak Condition for Consistency with EU.

$$q^*_j(\bar{h},\bar{l},\underline{h},\underline{l}) > \frac{\underline{h} - \underline{l}}{(\bar{h} - \bar{l}) + (\underline{h} - \underline{l})} \cdot e$$

Claim 1 deals with the case where return levels are fixed in problems Pj and P($j + 1$). In this case, risk-averse subjects should strictly increase the allocation to the dominated asset B as the correlation between the returns on A and B decreases, unless they allocate all their endowment to A in both problems Pj and P($j + 1$).

Claim 2 gives a lower (upper) bound on the amount that should be allocated to asset A (B). Since the bound is independent of the curvature of the utility function and applies even to "extremely risk-averse" agents, it provides a weak condition for consistency with EU-maximization. The intuition underlying the claim is straightforward: Since hl and lh are equally likely and $h(A) - l(B) > h(B) - l(A)$, the benefits from diversification are limited and EU-maximizers must restrict the allocation to the dominated asset.

Note that the upper bound on the amount that can be allocated to B becomes more restrictive as the dominated asset gets relatively riskier; that is, as $\underline{h} - \underline{l}$ increases relatively to $\bar{h} - \bar{l}$. In our experiment $(\underline{h} - \underline{l})/(\bar{h} - \bar{l} + \underline{h} - \underline{l})$ varied between 0.4167 and 0.5625, which does not appear too restrictive. The data, however, reveal that 20 of the 49 MBA students (40.8%), 56 of the 67 BUS participants (83.6%), and 13 of the 31 participants in Version 3 (41.9%) violate the weak consistency condition in at least one problem. The higher violation rate for the BUS clearly follows from the larger allocations to asset B in this version of the experiment. The problem level violation rates of the MBA students vary between 0–16% (see Table 7) compared to violations rates around 30% for the BUS.[24]

In Appendix C, we generalize Claims 1 and 2 to the case where subjects maximize rank-dependent utility. We show that the lower bound of Claim 2 on the rational allocation to asset A, applies—with a weak inequality—to the case where DM maximizes rank-dependent utility,

[23] The supplementary Appendix is available at http://www2.colman.ac.il/business/doron/.

[24] The average difference between the relevant bound and the actual (lower) allocation to asset A, in the cases where subjects violated the consistency condition, was 13,563 (13.56% of 100,000 endowment) for the BUS and 91.8 (9.18% of 1000 endowment) for the MBA.

for typical probability weighting schemes. The large allocations to asset B therefore cannot be explained by standard probability weighting. The literature in behavioral finance, however, documents two types of bias that may intuitively explain subjects' inclination to allocate a large proportion of their funds to the dominated asset.[25] Consider first the naïve diversification phenomenon (Benartzi and Thaler 2001). Investors (in practice and laboratories) tend to allocate $1/N$ of their funds to each of N investment channels proposed in a fixed investment menu. In our experiment, the apparent dominance of A results in clear preference for asset A over asset B, but the naïve inclination to equally diversify between the assets proposed in the menu results in exaggerated allocations to B. Alternatively, consider again the typical return tables that subjects observed (Appendix A). Recall that the *lh* row in each table is the only row where asset B pays higher return than asset A and provides the only rationale for investing in B. Borrowing terminology from the judgment psychology literature (Kahneman et al. 1982), the *lh* observations may appear "salient" and attract subjects' attention as representing the advantage of allocating funds to asset B. Since decision makers tend to overweight the probability of salient events (see, for example, the discussion in Barberis and Thaler 2003), the exaggerated allocations to B may follow from over-weighting of the *lh* cases.

7. Discussion

We run three versions of an experiment on financial prediction and allocation and find that subjects in all versions exhibit three types of bias: (i) probability matching in the prediction tasks, (ii) neglect of distributional data in the allocation tasks, and (iii) significant violations of EU/RDU.[26]

It is interesting to compare subjects' prediction patterns in the current experiment with those reported by Rubinstein (2002). In Rubinstein's cards experiment, subjects are told that five cards are randomly picked from a stack of 100 (36 green, 25 blue, 22 yellow, and 17 brown) and put in sealed envelopes. Subjects are asked to guess the color of the five hidden cards. Only 40% of 124 undergraduate students submit the rational "green" prediction for all five envelopes; about 30% of the other students chose to diversify predictions in proportions that are similar to the proportion of colors in the stack (probability matching). When the same experiment is run on 49 economics graduate students, 61% are classified rational while 33% are "probability matching." Note that the proportions of rational prediction in our allocation experiments are significantly lower (7.5% for the undergrads, about 30% for the MBAs). While the difference between the experiments is too large to draw concrete conclusions, our conjecture is that the lower levels of rationality in the current application may reflect subjects' difficulty in perceiving financial returns as time-independent or "random walk." Tversky and Kahneman (1971) use the term "law of small numbers" to describe the natural inclination to assign large-

[25] The inconsistencies may also follow from error or noise in subjects' allocations. The magnitude of violations (especially in the BUS version), however, motivates "direct" explanations.

[26] The pool of subjects used in our experiments is relatively sophisticated and experienced in financial decision. The literature suggests that "experts" (stock brokers or investment consultants) frequently produce similar biases to the ones observed with students. For interesting relevant references see Reyna (2004) and Torngren and Montgomery (2004); Reyna (2004) in particular shows that rates of "base-rate neglect" of medical doctors are similar to rates of high school students.

sample properties to small series. In our application, subjects expect that the empirical distribution would reappear in an independent sample of four draws.

The comparison of allocations across problems suggests that subjects neglect probabilistic information while strongly responding to differences in return levels when allocating funds to the two assets. Insensitivity to distributional information in portfolio selection was also observed in the experiments of Kroll et al. (1988) and Anderson and Settle (1996). Neglect of probabilistic information was documented in other completely different contexts in many preceding studies. The large literature on "base rate neglect" (Tversky and Kahneman 1982, for example) documents the neglect of prior probability distributions in deriving post-signal posteriors. To take a completely different type of example, the literature on catastrophic events (Sunstein 2003) suggests that when outcomes are especially bad, people tend to focus on the badness of the outcome while ignoring the probability that the outcome will occur. Using ideas from recent psychological literature on the cognitive process of decision making (Reyna 2004), we suggest that in both types of bias, subjects' attention is drawn to the salient elements in the formal description of the problem (the conditional probabilities of signals in cases of base-rate neglect; the catastrophic damage levels in studies of disaster) while it intuitively neglects other features of the problem (prior frequency rates; specific probabilities of bad outcomes) in spite of their normative relevance. In the current application, subjects' attention is drawn to the return levels on the two assets (and to the possibility to decrease risk by investing in asset B) while neglecting specific data on the joint distribution of returns.[27] The fact that nonprofessional financial media frequently analyze return levels, trends, and cycles while almost ignoring higher-level statistics like volatility and correlation seems to complement our experimental findings. The hypothesis that insensitivity to joint-distributional data is actually reflected in market prices of options and other financial assets seems an intriguing challenge for empirical research.

Appendix A: Typical Prediction/Allocation Problem (MBA Version 1)

Table A1 summarizes the empirical distribution of returns on the virtual assets A and B in a sample of 12 observations (as explained in the general instructions):

Task 1

Table A2 reveals the realized returns of one asset in four additional observations; we ask you to fill in the return for the other asset. Remember that correct guessing may give you a bonus payment of 20 N.I.S.

Task 2

For the current task, we endow you with 1000 N.I.S. You are requested to allocate the endowment between the financial assets A and B under the assumption that we will pay the realized return of your portfolio in observation 17.

I chose to allocate ___________ N.I.S. to asset A and _________ N.I.S. to asset B

(please verify that the amounts sum up to 1000).

[27] Using Reyna (2004) terminology, the objective probability distribution is not captured as part of the "gist" of the problem; Reyna uses similar ideas to explain the "base-rate neglect."

Table A1. Empirical Distribution of Returns on the Virtual Assets A and B

Number of Observations	Return on A	Return on B
4	14%	13%
2	3%	13%
4	3%	1%
2	14%	1%

Table A2. Realized Returns of One Asset in Four Additional Observations

Observation #	Return on A	Return On B
13	14%	
14		1%
15		1%
16		1%

Appendix B: Translated Instructions (MBA experiment)

The experiment is composed of five problems. In the next paragraphs we outline the structure of a typical problem. The problem starts by presenting information on the returns earned by two virtual assets (say X and Y) in a sample of 12 observations. The information will be presented in Table B1.

The data on the table suggests that in 4 of 12 observations asset X has earned 2% while asset Y earned 1%; in 2 of 12 observations asset X earned 15% while asset Y earned 1%, etc.

The return on each of the virtual assets can only take two values: high or low (in the table above, for example, the high return level on X is 15% while the low return level on X is 2%). The table would thus refer to four possible cases: the case where both assets earned the low return level; the case where asset X earned the low return level and asset Y earned the high return level, etc.

Assume that the information summarized in the table describes the joint distribution of returns (that is, the probability that the return on X is 2% and the return on Y is 1% is 4/12, etc.). To determine your realized payoff we have sampled five additional observations from the same distribution (observations 13–17). In the experiment you would be asked to complete two tasks:

1. To fill in the return on one asset from data on the return earned by the other asset—for four of the additional observations (observations 13–16).
2. Allocate an investment budget of 1000 N.I.S. between the two assets under the assumption that you will receive the returns earned by the two assets in observation 17.

After collecting your questionnaires and coding the data we will randomly select one of the five problems included in your questionnaire (“the selected problem”). The actual payoff that you will receive at the end of the experiment will be calculated from your answers to “the selected problem” in two steps:

1. You will receive the return earned by the portfolio that you have constructed. For example, if you have decided to allocate 600 N.I.S. to asset X and 400 N.I.S. to asset Y and the realized returns in observation 17 were 15% on X and 1% on Y, then the return that your portfolio has earned is $600 \times 0.15 + 400 \times 0.01 = 94$ N.I.S. You would thus receive 94 N.I.S.
2. We will randomly select one of the observations 13–16 for which you were asked to fill in the return table and check if the return that you have filled in was equal to the actual return (recall that observations 13–17 were already sampled by the experiment organizers). If the return that you have filled in was equal to the actual

Table B1. Returns Earned by Two Virtual Assets

Number of Observations	Return on X	Return on Y
4	2%	1%
2	15%	1%
2	2%	12%
4	15%	12%

return, we will pay you 20 more N.I.S. (in addition to the amount calculated in step 1). If the return that your have filled in was not equal to the actual return, you will not receive an additional payment.

Important remark: The questionnaire that you have received is ''individual'' in the sense that it is different from the questionnaires of the other participants. In addition we would like to remind you that investment allocation decisions of the type examined in the questionnaire depend on the individual tastes of each investor. The problems that you are asked to answer do not have a correct solution; the ''solution'' depends on your personal tastes. We ask you to fill in the questionnaire with attention and give us your candid answers. Students who are not interested to participate may return their questionnaire now.

Payments: At the end of the experiment we will give you an internet address where you may examine the payments received by the participants (subjects would be identified by their ID number for anonymity) and inform you of the payment collection procedure.

Appendix C: Supplementary Material

Rank-Dependent Utility

The rank-dependent utility model (Quiggin 1982) has become a prominent alternative to von Neumann Morgenstern expected utility framework when no negative payoffs are involved (see the surveys by Starmer 2000 and Schmidt 2004). The model was applied to various financial applications and used to shed light on portfolio selection patterns (Quiggin 1993; Haliassos and Hassapis 2001; Bleichrodt and Eeckhoudt 2005; Polkovnichenko 2005), the home-bias puzzle (Aizenman 1997), and equilibrium prices of financial assets (Tallon 1997).

Rank-dependent utility departs from expected utility in two major assumptions. The decision maker is assumed to consider the gains/losses from investment rather than considering the final wealth levels obtained. Moreover, the probabilities of different scenarios are weighted, where the weighting of given probability depends on the rank of the associated payoff. Formally, the model assumes an increasing utility function $U[\cdot]$ and a strictly increasing probability weighting function $w : [0, 1] \to [0, 1]$ satisfying $w(0) = 0$ and $w(1) = 1$ such that the utility from a lottery that pays $x_1 \geq x_2 \ldots \geq x_n$ with corresponding probabilities $p_1, p_2, \ldots, p_n$ (such that $\sum_{i=1}^{n} p_i = 1$) is

$$\sum_{i=1}^{n} \pi_i \cdot U(x_i),$$

where

$$\pi_1 = w(p_1) \quad \text{and} \quad \pi_i = w\left(\sum_{j=1}^{i} p_j\right) - w\left(\sum_{j=1}^{i-1} p_j\right) \quad \text{for } i = 2,3\ldots n.$$

To fit the model to our specific application, let $g_q(xy)$ denote the realized gain of a subject that allocated q to asset A and $(e - q)$ to asset B when the return on A was x and the return on B was y; that is, $g_q(xy) = q \cdot x + (e - q) \cdot y$. Note that when $0 \leq q \leq e$, $g_q(ll) \leq g_q(hl) \leq g_q(hh)$ and $g_q(ll) \leq g_q(lh) \leq g_q(hh)$ so that $g_q(hh)$ represents the highest possible payoff while $g_q(ll)$ represents the lowest payoff. The ranking of $g_q(hl)$ and $g_q(lh)$, however, depends on the specific value of q. Let

$$\bar{q}(\bar{h},\bar{l},\underline{h},\underline{l}) = \frac{\underline{h} - \underline{l}}{(\bar{h} - \bar{l}) + (\underline{h} - \underline{l})} \cdot e;$$

recall that the expression represents the threshold level for consistency with EU maximization as defined in Claim 2. It is easily verified that $g_q(\bar{h}\underline{l}) > g_q(\bar{l}\underline{h})$ when $q > \bar{q}(\bar{h},\bar{l},\underline{h},\underline{l})$ while $g_q(\bar{h}\underline{l}) < g_q(\bar{l}\underline{h})$ when $q < \bar{q}(\bar{h},\bar{l},\underline{h},\underline{l})$; the two payoffs are equal at $q = \bar{q}(\bar{h},\bar{l},\underline{h},\underline{l})$.

Let p_j denote the probability of hh (ll) cases in problem Pj. Note that $(0.5 - p_j)$ then represents the probability of hl (lh) cases. To concisely represent the objective function of a rank-dependent utility maximizer, we define four weights as follows:

- $w_j^1 = w(p_j)$ represents the weight of the highest-ranked payoff $g_q(hh)$;
- $w_j^2 = w(0.5) - w_j^1$ represents the weight of the second-highest payoff, which is $g_q(hl)$ when $q \geq \bar{q}$ and $g_q(lh)$ when $q < \bar{q}$;
- $w_j^3 = w(1 - p_j) - w(0.5)$ represents the weight of the third-ranked payoff: $g_q(lh)$ when $q \geq \bar{q}$ and $g_q(hl)$ when $q < \bar{q}$; and
- $w_j^4 = 1 - w(1 - p_j)$ represents the weight of the lowest-ranked payoff $g_q(ll)$.

The rank-dependent utility of DM is then written as follows:

$$RDU(q) = \begin{cases} w_j^1 \cdot U\left[g_q(hh)\right] + w_j^2 \cdot U\left[g_q(lh)\right] + w_j^3 \cdot U\left[g_q(hl)\right] + w_j^4 \cdot U\left[g_q(ll)\right] & 0 \le q \le \overline{q} \\ w_j^1 \cdot U\left[g_q(hh)\right] + w_j^2 \cdot U\left[g_q(hl)\right] + w_j^3 \cdot U\left[g_q(lh)\right] + w_j^4 \cdot U\left[g_q(ll)\right] & \overline{q} \le q \le e \end{cases}.$$

DM's objective is $\mathrm{MAX}_{0 \le q \le e}\ RDU(q)$.

As in section 6, we assume that $U[x]$ is twice continuously differentiable, strictly monotone ($U' > 0$) and strictly concave ($U'' < 0$). We present a convexity condition on the probability weighting function under which Claims 1 (monotonicity) and 2 (weak consistency with EU) generalize to the case of RDU maximization.

Definition: Convexity around 0.5. Say that the probability weighting function $w(\cdot)$ is convex around 0.5 iff $w(0.5 - p) + w(0.5 + p) \ge 2 \cdot w(0.5)$ for every $0 \le p \le 0.5$.

Put differently, convexity around 0.5 implies that $w(0.5 + p) - w(0.5) \ge w(0.5) - w(0.5 - p)$ for every $0 \le p \le 0.5$. It thus guarantees that $w_j^3 \ge w_j^2$ for each of the five problems examined in the experiment.

Claim B.1. When the probability weighting function is convex around 0.5, each Pj has a unique solution $q_j^*\left(\overline{h},\overline{l},\underline{h},\underline{l}\right)$ satisfying

(I) $q_{j+1}^*\left(\overline{h},\overline{l},\underline{h},\underline{l}\right) \le q_j^*\left(\overline{h},\overline{l},\underline{h},\underline{l}\right)$ for $j = 1, 2, 3, 4$, where the inequality is strict for $\overline{q}\left(\overline{h},\overline{l},\underline{h},\underline{l}\right) < q_j^*\left(\overline{h},\overline{l},\underline{h},\underline{l}\right) < e$.
(II) $q_j^*\left(\overline{h},\overline{l},\underline{h},\underline{l}\right) \ge \overline{q}\left(\overline{h},\overline{l},\underline{h},\underline{l}\right)$.[28]

Claim B.2. The Tversky and Kahneman (1992) probability weighting function

$$w(p) = \frac{p^\gamma}{\left[p^\gamma + (1-p)^\gamma\right]^{1/\gamma}}$$

is convex around 0.5.

Claims B.1 and B.2 imply that the inconsistency rates described in brackets on Table 7 apply to the case where DM is a rank-dependent utility maximizer with the frequently used probability weighting function proposed by Tversky and Kahneman (1992); 36.73% of the MBA students, 77.61% of the BUS students, and 19.3% of the participants in Version 3 violated the weak condition for consistency with rank-dependent maximization in at least one of the 5/2 problems. The proofs are provided below.

Proofs

Proof of Claim C.1. Since each asset i pays $h(i)$ or $l(i)$ with equal probabilities, the expected payoff on asset i is $E(i) = 0.5 * h(i) + 0.5 * l(i)$ and $h(i) - E(i) = E(i) - l(i) \equiv d_i$. Let p denote the probability of hh (ll) realizations in the joint distribution data; the probability of hl (lh) realizations is then equal to $(1 - 2p)/2$. The covariance between the returns is calculated directly: $\mathrm{Cov(A, B)} = p \cdot d_A \cdot d_B - (1 - 2p) \cdot d_A \cdot d_B + p \cdot d_A \cdot d_B$; and it follows that the Pearson coefficient of correlation between the returns is $4p - 1$. Substituting the appropriate p-values for problems P1–P5 gives the levels of correlation on the right column of Table 1. *QED*.

Proof of Claim 1. Let $q_j^* \equiv q_j^*\left(\overline{h},\overline{l},\underline{h},\underline{l}\right)$ denote the solution to Pj.

Let p_j denote the probability of hh (ll) return-combinations in problem Pj; $(1/2 - p_j)$ then represents the probability of hl (lh) cases.

Consider first the interior case where $q_j^* < e$ (note that $q_j^* = 0$ is impossible by Claim 2).

Rewrite the first-order conditions in Equation 2 as follows:

$$\begin{aligned} &p_j \cdot \left(\overline{h} - \underline{h}\right) \cdot U'\left[\pi_{q_j^*}(hh)\right] + \left(1/2 - p_j\right) \cdot \left(\overline{h} - \underline{l}\right) \cdot U'\left[\pi_{q_j^*}(hl)\right] \\ &+ \left(1/2 - p_j\right) \cdot \left(\overline{l} - \underline{h}\right) \cdot U'\left[\pi_{q_j^*}(lh)\right] + p_j \cdot \left(\overline{l} - \underline{l}\right) \cdot U'\left[\pi_{q_j^*}(ll)\right] = 0. \end{aligned} \tag{B1}$$

That is,

$$\begin{aligned} &p_j \cdot \left[\left(\overline{h} - \underline{h}\right) \cdot U'\left[\pi_{q_j^*}(hh)\right] + \left(\overline{l} - \underline{l}\right) \cdot U'\left[\pi_{q_j^*}(ll)\right]\right] \\ &+ \left(1/2 - p_j\right) \cdot \left[\left(\overline{h} - \underline{l}\right) \cdot U'\left[\pi_{q_j^*}(hl)\right] + \left(\overline{l} - \underline{h}\right) \cdot U'\left[\pi_{q_j^*}(lh)\right]\right] = 0. \end{aligned} \tag{B2}$$

[28] Note that we use a weak inequality in the case of rank-dependent utility (Claim B.1) while we were using a strong inequality in the case of EU maximization (Claim 6.1.2). The difference is discussed in the proof of Claim B.1.

Since

$$(\bar{h} - \underline{h}) \cdot U'\left[\pi_{q_j^*}(hh)\right] + (\bar{l} - \underline{l}) \cdot U'\left[\pi_{q_j^*}(ll)\right] > 0, \tag{B3}$$

it follows from Equation B2 that

$$(\bar{h} - \underline{l}) \cdot U'\left[\pi_{q_j^*}(hl)\right] + (\bar{l} - \underline{h}) \cdot U'\left[\pi_{q_j^*}(lh)\right] < 0. \tag{B4}$$

Since $p_j > p_{j+1}$ for $j = 1, 2, 3, 4$, Equations B2–B4 imply that

$$\begin{aligned} & p_{j+1} \cdot \left[(\bar{h} - \underline{h}) \cdot U'\left[\pi_{q_j^*}(hh)\right] + (\bar{l} - \underline{l}) \cdot U'\left[\pi_{q_j^*}(ll)\right]\right] \\ & + (1/2 - p_{j+1}) \cdot \left[(\bar{h} - \underline{l}) \cdot U'\left[\pi_{q_j^*}(hl)\right] + (\bar{l} - \underline{h}) \cdot U'\left[\pi_{q_j^*}(lh)\right]\right] < 0, \end{aligned}$$

so that the derivative of the objective (1) for problem $j + 1$ is negative at q_j^*. Since (1) is strictly concave it immediately follows that $q_{j+1}^* < q_j^*$ when $q_j^* < e$.

To complete the proof, note that when $q_j^* = e$, the left-hand-side (LHS) of Equation B1 may be positive; it follows from the arguments above that $q^*_{j+1} \leq e$ in this case. *QED*.

Proof of Claim 2. Note that if $q_j^*(\bar{h},\bar{l},\underline{h},\underline{l}) = e$, the claim holds.

Consider next the case where $0 < q_j^*(\bar{h},\bar{l},\underline{h},\underline{l}) <$ e:

From Equation B4, we know that

$$(\underline{h} - \bar{l}) \cdot U'\left[\pi_{q_j^*}(lh)\right] - (\bar{h} - \underline{l}) \cdot U'\left[\pi_{q_j^*}(hl)\right] > 0;$$

so that

$$(\underline{h} - \bar{l}) \cdot U'\left[\pi_{q_j^*}(lh)\right] > (\bar{h} - \underline{l}) \cdot U'\left[\pi_{q_j^*}(hl)\right]. \tag{B5}$$

Because $\bar{h} - \underline{l} > \underline{h} - \bar{l} > 0$ and $U' > 0$, it follows from Equation B5 that

$$U'\left[\pi_{q_j^*}(lh)\right] > U'\left[\pi_{q_j^*}(hl)\right]; \tag{B6}$$

so that q_j^* must satisfy

$$q_j^* \cdot (1 + \bar{l}) + \left(e - q_j^*\right) \cdot (1 + \underline{h}) < q_j^* \cdot (1 + \bar{h}) + \left(e - q_j^*\right) \cdot (1 + \underline{l}), \tag{B7}$$

which gives the inequality stated at the claim.

To complete the proof, note that $q_j^*(\bar{h},\bar{l},\underline{h},\underline{l}) = 0$ is impossible since the LHS of Equation B1 is positive at $q_j^* = 0$. *QED*.

Proof of Claim B.1. Note that the ranking of payoffs $g_q(hh) > g_q(lh) > g_q(hl) > g_q(ll)$ is fixed and independent of q on the interval $(0, \bar{q})$. Since the second derivative of $RDU(q)$ with respect to q is negative on this interval, it follows that $RDU(q)$ is strictly concave on $(0, \bar{q})$.

A similar argument implies that RDU is strictly concave on $(\bar{q}, e)$.

The RDU function however is nondifferentiable at $\bar{q}$.

The left-derivative at $\bar{q}$ is

$$\begin{aligned} & w_j^1 \cdot (\bar{h} - \underline{h}) \cdot U'\left[g_{\bar{q}}(hh)\right] + w_j^2 \cdot (\bar{l} - \underline{h}) \cdot U'\left[g_{\bar{q}}(lh)\right] \\ & + w_j^3 \cdot (\bar{h} - \underline{l}) \cdot U'\left[g_{\bar{q}}(hl)\right] + w_j^4 \cdot (\bar{l} - \underline{l}) \cdot U'\left[g_{\bar{q}}(ll)\right]; \end{aligned} \tag{B8}$$

while the right-derivative at $\bar{q}$ is

$$\begin{aligned} & w_j^1 \cdot (\bar{h} - \underline{h}) \cdot U'\left[g_{\bar{q}}(hh)\right] + w_j^2 \cdot (\bar{h} - \underline{l}) \cdot U'\left[g_{\bar{q}}(hl)\right] \\ & + w_j^3 \cdot (\bar{l} - \underline{h}) \cdot U'\left[g_{\bar{q}}(hl)\right] + w_j^4 \cdot (\bar{l} - \underline{l}) \cdot U'\left[g_{\bar{q}}(ll)\right]. \end{aligned} \tag{B9}$$

Recall, however, that when the probability weighting function is convex at 0.5, $w_j^3 \geq w_j^2$; since $\bar{h} - \underline{l} > \underline{h} - \bar{l}$ and $U'[g_{\bar{q}}(hl)] = U'[g_{\bar{q}}(lh)]$, we may conclude that

$$w_j^2 \cdot (\bar{l} - \underline{h}) \cdot U'\left[g_{\bar{q}}(lh)\right] + w_j^3 \cdot (\bar{h} - \underline{l}) \cdot U'\left[g_{\bar{q}}(hl)\right] > 0.$$

This implies that the left-derivative at $\bar{q}$ as presented in Equation B8 is positive. From the strict concavity of RDU on $(0, \bar{q})$ it follows immediately that if q_j^* solves Pj, then $q_j^* \geq \bar{q}$. The strict concavity of RDU on $(\bar{q}, e)$ then implies that a unique solution $\bar{q} \leq q_j^* \leq e$ exists, which proves part (II) of the claim. (Note that when $w_j^3 > w_j^2$, the right derivative [Equation B9] is lower than the left derivative [Equation B8]; it is thus possible to construct examples where the left derivative [Equation B8] is positive while the right derivative [Equation B9] is negative, in such cases the optimal solution to the RDU problem is obtained at $\bar{q}$.)[29]

To show that monotonicity holds in the case of RDU maximization, assume that the first order conditions for problem Pj are satisfied at some $\bar{q} \leq q_j^* \leq e$ (the possibility that $q_j^* < \bar{q}$ has been ruled out in the preceding argument). Consider first the interior case where $\bar{q} < q_j^* < e$.

From the strict concavity of RDU on $(\bar{q}, e)$, it follows that the derivative of RDU must be 0 at q_j^*; that is,

$$\begin{aligned} & w_j^1 \cdot (\bar{h} - \underline{h}) \cdot U'\left[g_{q_j^*}(hh)\right] + w_j^2 \cdot (\bar{h} - \underline{l}) \cdot U'\left[g_{q_j^*}(lh)\right] \\ & + w_j^3 \cdot (\bar{l} - \underline{h}) \cdot U'\left[g_{q_j^*}(lh)\right] + w_j^4 \cdot (\bar{l} - \underline{l}) \cdot U'\left[g_{q_j^*}(ll)\right] = 0. \end{aligned} \tag{B10}$$

For simplicity, rewrite Equation B10 as follows:

$$w_j^1 \cdot K1 + w_j^2 \cdot K2 + w_j^3 \cdot K3 + w_j^4 \cdot K4 = 0, \tag{B11}$$

where

$$\begin{aligned} K1 &= (\bar{h} - \underline{h}) \cdot U'\left[g_{q_j^*}(hh)\right] \geq 0, \\ K2 &= (\bar{h} - \underline{l}) \cdot U'\left[g_{q_j^*}(hl)\right] > 0, \\ K3 &= (\bar{l} - \underline{h}) \cdot U'\left[g_{q_j^*}(lh)\right] < 0, \text{ and} \\ K4 &= (\bar{l} - \underline{l}) \cdot U'\left[g_{q_j^*}(ll)\right] > 0. \end{aligned}$$

This gives an explicit expression for $K2$:

$$K2 = \frac{w_j^1 \cdot K1 + w_j^3 \cdot K3 + w_j^4 \cdot K4}{-w_j^2}. \tag{B12}$$

Consider now the derivative of the objective function for problem $P(j+1)$ with respect to q_{j+1}:

$$\begin{aligned} & w_{j+1}^1 \cdot (\bar{h} - \underline{h}) \cdot U'\left[g_{q_{j+1}}(hh)\right] + w_{j+1}^2 \cdot (\bar{h} - \underline{l}) \cdot U'\left[g_{q_{j+1}}(hl)\right] \\ & + w_{j+1}^3 \cdot (\bar{l} - \underline{h}) \cdot U'\left[g_{q_{j+1}}(lh)\right] + w_{j+1}^4 \cdot (\bar{l} - \underline{l}) \cdot U'\left[g_{q_{j+1}}(ll)\right]. \end{aligned} \tag{B13}$$

At $q_{j+1} = q_j^*$, Equation B13 can be represented as follows:

$$w_{j+1}^1 \cdot K1 + w_{j+1}^2 \cdot K2 + w_{j+1}^3 \cdot K3 + w_{j+1}^4 \cdot K4. \tag{B14}$$

Substituting $K2$ from Equation B12 into Equation B14 gives Equation B15 below:

$$w_{j+1}^1 \cdot K1 + w_{j+1}^3 \cdot K3 + w_{j+1}^4 \cdot K4 - \frac{w_{j+1}^2}{w_j^2} \cdot \left[w_j^1 \cdot K1 + w_j^3 \cdot K3 + w_j^4 \cdot K4\right]. \tag{B15}$$

Since $p_{j+1} < p_j$, it follows from the monotonicty of w that $w_{j+1}^1 < w_j^1$, $w_{j+1}^2 > w_j^2$, $w_{j+1}^3 > w_j^3$, and $w_{j+1}^4 < w_j^4$. Since $K1 \geq 0$, $K3 < 0$, and $K4 > 0$, we conclude that

[29] For an arbitrary example, take the Tversky and Kahneman (1992) probability weighting function with $\gamma = 0.6$, a power utility function $u(x) = x^a$ with $a = 0.6$, and consider P5 with the return levels 0.14, 0.02, 0.14, and 0.01. When $e = 100{,}000$, the point of non-differentiability is 52,000. Simple calculations show the left derivative at this point is positive 0.000457 while the right derivative is -0.00001.

$$0 < \left[w_{j+1}^{1}\cdot K1 + w_{j+1}^{3}\cdot K3 + w_{j+1}^{4}\cdot K4\right] < \left[w_{j}^{1}\cdot K1 + w_{j}^{3}\cdot K3 + w_{j}^{4}\cdot K4\right], \tag{B16}$$

and that

$$\frac{w_{j+1}^{2}}{w_{j}^{2}} > 1 \tag{B17}$$

From Equations B16 and B17, it follows that the expression in Equation B14 is negative. That is, the derivative of the objective function for problem $P(j+1)$ is negative at q_j^*. The strict concavity of the objective function on $(\bar{q}, e)$ then implies that $q_{j+1}^* < q_j^*$ when $\overline{q} < q_j^* < e$.

Finally note that when $q_j^* = e$, $q_{j+1}^* \leq q_j^*$ by definition. When $q_j^* = \overline{q}$, on the other extreme, the LHS of Equation B10 may be negative and the proof above implies that $q_{j+1}^* = q_j^* = \overline{q}$ in this case. *QED*.

Proof of claim B.2. Convexity around 0.5 holds if $w(0.5 + p) + w(0.5 - p) > 2 \cdot w(0.5)$.

Substituting the weighting function

$$w(x) = \frac{x^{\gamma}}{[x^{\gamma} + (1 - x)^{\gamma}]^{1/\gamma}}$$

into the convexity condition and rearranging gives the inequality

$$[(0.5 + p)^{\gamma} + (0.5 - p)^{\gamma}]^{\frac{\gamma - 1}{\gamma}} > [2\cdot 0.5^{\gamma}]^{\frac{\gamma - 1}{\gamma}}. \tag{B18}$$

To prove Equation B18, consider first the case where $\gamma < 1$. Since $(\gamma - 1)/\gamma < 0$ in this case, we have to show that $[(0.5 + p)^{\gamma} + (0.5 - p)^{\gamma}] < [2 \cdot 0.5^{\gamma}]$, but this follows immediately from the concavity of x^{γ} for $\gamma < 1$, as $0.5 \cdot (0.5 + p)^{\gamma} + 0.5 \cdot (0.5 - p)^{\gamma} < 0.5^{\gamma}$ when $\gamma < 1$. The proof for the case where $\gamma > 1$ similarly follows from the convexity of x^{γ} for $\gamma > 1$. *QED*.

References

Aizenman, J. 1999. International portfolio diversification with generalized expected utility preferences. *Canadian Journal of Economics* 32:995–1008.

Anderson, B. F., and J. W. Settle. 1996. The influence of portfolio characteristics and investment periods on investment choice. *Journal of Economic Psychology* 17:343–59.

Barberis, N., and R. Thaler. 2003. A survey of behavioral finance. In *Handbook of the economics of finance,* volume 1, edited by G. M. Constantinides, M. Harris, and R. M. Stulz. Amsterdam: Elsevier Science, pp. 1053–1128.

Benartzi, S., and R. H. Thaler. 2001. Naive diversification strategies in defined contribution saving plans. *American Economic Review* 91:79–98.

Benzion, U., E. Haruvy, and T. Shavit. 2004. Adaptive portfolio allocation with options. *Journal of Behavioral Finance* 5:43–56.

Bleichrodt, H., and L. Eeckhoudt. 2005. Saving under rank-dependent utility. *Economic Theory* 25:505–11.

Bosch-Domènech, A., and J. Silvestre. 1999. Does risk-aversion or attraction depend on income? *Economic Letters* 65:265–73.

Haliassos, M., and C. Hassapis. 2001. Non-expected utility, savings and portfolio. *Economic Journal* 110:69–102.

Harrison, W. H., E. Johnson, M. M. McInnes, and E. E. Ruström. 2005. Risk aversion and incentive effects: A comment. *American Economic Review* 3:897–901.

Holt, C. A., and S. K. Laury. 2002. Risk aversion and incentive effects. *American Economic Review* 92:1644–55.

Kahneman, A., P. Slovic, and A. Tversky. 1982. *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.

Kroll, Y., H. Levy, and A. Rapoport. 1988. Experimental tests of the separation theorem and the capital asset pricing model. *American Economic Review* 78:500–19.

Markowitz, H. 1952. Portfolio selection. *Journal of Finance* 7:77–91.

Polkovnichenko, V. 2005. Household portfolio diversification: A case for rank-dependent preferences. *Review of Financial Studies* 18:1467–1502.

Quiggin, J. 1982. A theory of anticipated utility. *Journal of Economic Behavior and Organization* 3:323–43.

Quiggin, J. 1993. *Generalized expected utility theory: The rank-dependent utility model*. Amsterdam: Kluwer Academic Publishers.

Reyna, V. F. 2004. How people make decisions that involve risk. *Current Directions in Psychological Science* 13:60–6.

Rubinstein, A. 2002. Irrational diversification in multiple decision problems. *European Economic Review* 46:1369–78.

Schmidt, U. 2004. Alternatives to expected utility: Some formal theories. In *Handbook of utility theory,* volume 2, edited by P. J. Hammond, S. Barabera, and C. Seidl. Boston: Kluwer, pp. 757–838.

Starmer, C. 2000. Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature* 38:332–82.

Sunstein, R. C. 2003. Terrorism and probability neglect. *Journal of Risk and Uncertainty* 26:121–36.

Tallon, J. M. 1997. Micro-economic risk and asset prices in general equilibrium model with rand-dependent-expected utility. *Finance* 18:139–53.

Torngren, G., and H. Montgomery. 2004. Worse than chance? Performance and confidence among professionals and laypeople in the stock market. *Journal of Behavioral Finance* 5:148–53.

Tversky, A., and D. Kahneman. 1971. Belief in the law of small numbers. *Psychological Bulletin* 76:105–10.

Tversky, A., and D. Kahneman. 1982. Evidential impact of base rates. In *Judgment under uncertainty: Heuristics and biases*, edited by D. Kahneman, P. Slovic, and A. Tversky. Cambridge: Cambridge University Press, pp. 153–60.

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

Von Neumann, J., and O. Morgenstern. 1944. *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

Vulkan, N. 2000. An economist's perspective on probability matching. *Journal of Economic Surveys* 14:101–18.