

THE EFFECT OF MYOPIA AND LOSS AVERSION ON RISK TAKING: AN EXPERIMENTAL TEST*

RICHARD H. THALER
AMOS TVERSKY
DANIEL KAHNEMAN
ALAN SCHWARTZ

Myopic loss aversion is the combination of a greater sensitivity to losses than to gains and a tendency to evaluate outcomes frequently. Two implications of myopic loss aversion are tested experimentally. 1. Investors who display myopic loss aversion will be more willing to accept risks if they evaluate their investments less often. 2. If all payoffs are increased enough to eliminate losses, investors will accept more risk. In a task in which investors learn from experience, both predictions are supported. The investors who got the most frequent feedback (and thus the most information) took the least risk and earned the least money.

In an innovative paper Mehra and Prescott [1985] announced the existence of a new anomaly that they dubbed the equity premium puzzle. The equity premium is defined as the difference in returns between equities (stocks) and a risk-free asset such as treasury bills. The puzzle about the historic equity premium is that it has been very large. Over the time period Mehra and Prescott studied (1889–1978) the annual real return on the S&P 500 was about 7 percent while the return on T-bills was less than 1 percent. Since 1978, stocks have done even better. When these large return differences are cumulated, the differentials become staggering. For example, a dollar invested in the S&P 500 on January 1, 1926, was worth over \$1100 by the end of 1995, while a dollar invested in T-bills was worth only \$12.87. Mehra and Prescott show that it is difficult to explain the combination of a high equity premium and a low risk-free rate within a standard neoclassical model. The implicit coefficient of relative risk aversion needed to produce such numbers was over 30, while most estimates put it close to 1.¹

In a recent paper Benartzi and Thaler [1995] offer an explanation of the equity premium using two behavioral concepts:

*We are grateful for financial support from the Dreman Foundation. Special thanks are due to an anonymous referee whose comments greatly improved the paper. This research was conducted prior to Amos Tversky's death. He saw an earlier version of this paper, but did not participate in this revision. The paper, like many other things, would have been better if Amos had lived longer.

1. See Siegel and Thaler [1997] for a recent review of the literature on the equity premium puzzle.

mental accounting and loss aversion (specifically Tversky and Kahneman's [1992] cumulative prospect theory). Although they could not formally test their hypothesis, the value of the one parameter they estimate (the frequency with which portfolios are evaluated) appears to be plausible (roughly one year). They call their explanation myopic loss aversion since it is the combination of short horizons and strong distaste for losses that produces their results. Still, they did not offer any direct evidence in support of myopic loss aversion.

This paper reports the results of an experiment designed to fill the gap. Section I offers a brief review of myopic loss aversion. Section II describes the experimental design. Section III reports the results. Section IV discusses the implications of the results.

I. MYOPIC LOSS AVERSION

The myopic loss aversion explanation rests on two behavioral principles: loss aversion and mental accounting. Loss aversion refers to the fact that people tend to be more sensitive to decreases in their wealth than to increases. This is reflected in the prospect theory value function that has a kink at the origin. Empirical estimates find that losses are weighted about twice as strongly as gains (e.g., Tversky and Kahneman [1992]; Kahneman, Knetsch, and Thaler [1990]). The disutility of losing \$100 is roughly twice the utility of gaining \$100.

Mental accounting is the set of (implicit or explicit) cognitive activities that individuals and households engage in to serve the same function that regular accounting serves in an organization. In the context of financial transactions, the key mental accounting issues concern aggregation: how transactions are grouped both cross-sectionally (e.g., are securities evaluated one at a time or as portfolios) and intertemporally (how often are portfolios evaluated). A financial investor can be modeled as making a series of decisions about the allocation of his assets. Mental accounting determines both the framing of decisions and the experience of the outcomes of these decisions. An investor who frames decisions narrowly will tend to make short-term choices rather than adopt long-term policies. An investor who frames past outcomes narrowly will evaluate his gains and losses frequently. In general, narrow framing of decisions and narrow framing of outcomes tend to go together, and the combination of both tendencies defines a myopic investor.

Loss aversion and mental accounting are well illustrated by a famous example due to Samuelson [1963]. Samuelson once offered a colleague the following bet: flip a coin, heads you win \$200 and tails you lose \$100. Samuelson reports that his colleague turned this bet down but said that he would be happy to take 100 such bets. Samuelson then proved a theorem that showed that this pair of choices is irrational. That is, someone should not be willing to play a bet many times if he is not willing to play it just once.

Of more interest here is what the colleague offered as his rationale for turning down the bet: "I won't bet because I would feel the \$100 loss more than the \$200 gain." This sentiment is the intuition behind the concept of loss aversion. One simple utility function that would capture this notion is the following:

$$(1) \quad U(x) = \begin{array}{ll} x & x \geq 0 \\ 2.5x & x < 0, \end{array}$$

where x is a *change* in wealth relative to the status quo. With this loss-averse utility function, the utility of a single play of the gamble is negative (less than the status quo of not playing). However, the distribution of outcomes created by the portfolio of two bets $\{\$400, .25; 100, .50; -\$200, .25\}$ yields positive expected utility. Whether or not a decision maker accepts playing the gamble twice depends on mental accounting. A myopic decision maker will first determine whether he likes the prospect of the initial gamble in the series, will conclude that he does not, and consequently reject the entire series. Samuelson's friend was evidently loss averse, but not myopic. The argument developed by Benartzi and Thaler [1995] is that the price of financial assets reflects the preferences of investors who are both loss averse and myopic.

To appreciate the effect of myopia on risk attitudes, consider an investor with the utility function defined above who must choose between a risky asset that pays an expected 7 percent per year with a standard deviation of 20 percent (like stocks) and a safe asset that pays a sure 1 percent. By the logic we applied to Samuelson's colleague, the attractiveness of the risky asset depends on the time horizon of the investor. An investor who is prepared to wait a long time before evaluating the outcome of the investment as a gain or a loss will find the risky asset more attractive than another investor (equally loss averse, but more myopic) who expects to evaluate the outcome soon. Furthermore, investors who differ in the frequency with which they evaluate

outcomes will not derive the same utility from owning stocks. The probability of observing a loss is higher when the frequency of evaluation is high. If losses cause more mental anguish than equivalent gains cause pleasure, the experienced utility associated with owning stocks is lower for the more myopic investor [Kahneman, Wakker, and Sarin 1997]. Over time, the myopic investor is expected to gravitate to a lower level of risk.

The goal of the current paper is to provide more direct tests of the separate effects of myopia and loss aversion on risk preferences. We also ask what investors who vary in myopia will learn from their experience. The basic situation that we investigate simulates a series of decisions about an allocation of resources between two assets that differ in risk, in a multiperiod game. All subjects face the same allocation problem, but we protect some subjects against their own myopia in two ways: (i) by compelling them to commit themselves to the same allocation for several periods, and (ii) by providing them with information about their outcomes at relatively infrequent intervals. We expect this simulation of reduced myopia to induce a higher willingness to take risks. In the standard model of a rational utility-maximizing investor, of course, the framing of outcomes should have no effect at all.

The second major goal of the paper is to test the role of loss aversion in risk aversion. Notice that a decision maker with the simple preferences described by (1) is risk neutral for gambles strictly in the domain of either gains or losses but risk averse for mixed gambles (gambles that have some positive and negative outcomes). Suppose that an individual with these preferences is indifferent between a mixed gamble (or investment) and some sure outcome S . Let $-L$ be the largest loss possible for playing G . Consider what happens if L is added to both G and S . Now G is preferred to S . Reality is slightly more complex. According to cumulative prospect theory [Tversky and Kahneman 1992], the shape of the value function for monetary incomes induces slight risk aversion in the domain of gains (and an equivalent level of risk seeking in the domain of losses). Specifically, the cumulative prospect theory value function is

$$(2) \quad v(x) = \begin{cases} x^\alpha & \text{if } x \geq 0 \\ -\lambda(-x)^\beta & \text{if } x < 0, \end{cases}$$

where λ is the coefficient of loss aversion. Tversky and Kahneman have estimated α and β to be 0.88 and λ to be 2.25. An individual

with the preferences described by cumulative prospect theory is only mildly risk averse for gambles involving only gains, but strongly risk averse for gambles that entail potential losses.²

This analysis yields the prediction that adding a constant to the returns of stocks and bonds increases the attractiveness of stocks by decreasing the frequency of losses. To test this prediction, a fourth condition is run in which all the returns are increased by a constant large enough to assure that the returns on stocks over the shortest evaluation period are always positive. This is accompanied by a statement that "there is a high rate of inflation so both funds will always have a positive return." Notice that the prediction that this manipulation will have an effect means that we are implicitly assuming that our subjects exhibit money illusion. While this is not an explicit assumption of either the myopic loss aversion theory or of prospect theory, prospect theory does assume that people are passive in accepting the "frames" or problem descriptions offered to them. This is, after all, what allows framing effects (different choices from alternative descriptions of the same problem) to work. Here, the subjects only see positive returns and the inflation rate is not even stated, so it would be impossible for subjects to transform the nominal returns presented into their real counterparts. We therefore think that the assumption of money illusion in the experiment is a reasonable one, whether or not real investors suffer from the same illusion.

II. EXPERIMENTAL DESIGN

The subjects in the experiment were 80 undergraduate students at the University of California at Berkeley. Subjects were asked to imagine that they were the portfolio manager for a small college, and were told that they would be required to allocate a portfolio of 100 shares between 2 investments. Responses were given as an integer from 0–100 indicating the level of investment in Fund A; the level of investment in Fund B was automatically calculated ($B = 100 - A$), and subjects were asked to verify each

2. The simple analysis in the case of piecewise linear preferences does not carry over to cumulative prospect theory with nonlinear preferences. There can be special cases where a very large increase in the payoffs to the gamble and the sure outcome can make the sure outcome relatively more attractive. However, the actual change in payoffs we used does make stocks more attractive to a prospect theory investor with the preferences given by (2). Furthermore, that will be true for any value of α as long as $\lambda > 1.55$.

allocation. Subjects were paid for their participation. Regarding their pay they were told the following: "The amount of money you will collect at the end of the experiment will depend entirely on the success of your investment decisions. The likely range of outcomes is between \$5 and \$30." In fact, subject payments ranged from \$5 to \$35. The stakes were large both in comparison to the pay for part-time work and for other experiments on campus, and we were confident that the subjects took the task seriously. The decision to make the exact mapping from portfolio returns to payments vague was an intentional one. In the version with many trials the stakes for any one trial were necessarily quite small, and we did not want to reinforce this by emphasizing that the expected payoffs to the subjects on a particular trial was only a few cents.

The experiment was designed to simulate investment over time in two hypothetical funds. Fund A was drawn from a normal distribution with a mean return per period of 0.25 percent and a standard deviation of 0.177 percent, and was truncated at 0 to prevent negative returns. Fund B was drawn from a normal distribution with a mean return of 1 percent and a standard deviation of 3.54 percent. These values correspond approximately to the actual returns of five-year bonds and a value-weighted stock index over 6.5 weeks. This information was *not* given to the subjects, however. They had to learn about the risks and returns of the two funds through experience. The subjects were also not told that Fund A was bonds and Fund B was stocks. However, to simplify the presentation, hereafter we will refer to Fund A as the Bond Fund and Fund B as the Stock Fund.

Each trial in the experiment consisted of presenting information to the subject, and recording the subject's allocation of shares to the investments. Subjects were randomly assigned to one of four conditions. In the "monthly" condition, subjects made 200 decisions, each corresponding to a single period. In the "yearly" condition, subjects made 25 decisions, each of which, they were told, was to be binding for eight periods. In the "five-yearly" condition, subjects made five decisions, each of which was binding for 40 periods. The "inflated monthly" condition was equivalent to the monthly condition, but the returns were translated upward by 10 percent so that subjects always experienced nominally positive returns from both funds. Subjects in this condition were told that there was a high rate of inflation which was responsible for returns always being positive.

After each decision, subjects saw a bar graph that displayed the aggregated returns of each fund and of their portfolio for the period(s) to which the decision applied. To avoid an inadvertent effect caused by the displays themselves, each condition's graphs were scaled so that an *average* return for each fund would appear the same size on the subject's screen across conditions.

The actual returns the subjects received were created as follows. For the "monthly" subjects, the returns displayed were, in fact, drawn at random from the distributions described above. Subjects in the other experiments were "yoked" to these subjects in order to eliminate any variation in the observed returns across conditions. So the first subject in the experiment was put in the monthly condition and got the first set of random returns. The second subject was placed in the "yearly" condition and got the same return in his or her first period as the aggregated return over the first eight³ periods for the first subject in the monthly condition. The third subject then received the same feedback but in five-yearly increments, and the fourth subject just received the same feedback as the first subject except that 10 percent was added to all the returns (so no return was negative). For the fifth subject, back in the monthly condition, a new set of returns was generated, and the process repeated.

At the conclusion of the trials, each subject was asked to make a final allocation that would be binding for 400 periods (with no intervening feedback). This decision was the same for subjects in each of the four conditions. As with the experience during the trials, this final allocation would help determine how much money the subject earned during the experiment. However, since this decision was for 400 periods and the previous trials were collectively 200 periods, in fact, the final decision determined two-thirds of the subject's payoff.

III. RESULTS

In Panel A of Table I we report the mean final (400-period) allocations to the Bond Fund for each of the four conditions. This high stakes decision is made after all the feedback has been received. The allocation to the Stock Fund is 100 percent less the

3. The yearly condition is eight periods rather than twelve because the "months" were 6.5 weeks. We chose this because we did not want the subjects in the monthly condition to have to make so many decisions that they would get bored.

TABLE I
ALLOCATIONS TO BOND FUND

Feedback group	Percent allocation to bond fund			
	<i>n</i>	Mean	<i>SD</i>	<i>SE</i>
A. Final decision				
Monthly	21	59.1	35.4	7.73
Yearly	22	30.4 ^b	25.9	5.51
Five-yearly	22	33.8 ^b	28.5	6.07
Inflated monthly	21	27.6 ^b	23.2	5.07
B. During the last five years				
Monthly	840	55.0	31.8	1.10
Yearly	110	30.7 ^a	27.0	2.57
Five-yearly	22	28.6 ^a	25.1	5.36
Inflated monthly	840	39.9	33.5	1.16

In each column, means with common superscripts do not differ significantly from one another ($p > .01$).

allocation to the Bond Fund. Myopic loss aversion implies two predictions. First, the final allocations to bonds should fall as the length of the evaluation period increases. Second, the allocation to bonds should fall when returns are transformed to eliminate losses. Both predictions receive some support. In the baseline monthly condition, subjects put 59.1 percent of their final period money in the Bond Fund. Subjects in the aggregated conditions invest less in the Bond Fund: 30.4 percent in the yearly condition and 33.8 percent in the five-yearly condition. Both of these amounts are significantly less than the allocation in the monthly condition, and are not significantly different from each other. Two factors probably contribute to the similarity of results in the two aggregated conditions: first, there was much less opportunity to learn from experience in the five-yearly condition (five trials) than in the yearly condition (40 trials). Second, the present analysis assigns most importance to the intimidating effect of trials on which the returns from stocks was negative. In the monthly condition 39 percent of the trials had negative returns; whereas only 14 percent of the yearly returns were negative, and none were negative in the five-yearly conditions. The difference in the frequency of losses between the two aggregate conditions was probably too small to have an effect. In the fourth condition of the experiment, the frequency of losses was zero, and the learning period was long (200 trials). Subjects in this condition invested only 27.6 percent of their assets in the bond fund, a

reduction of more than 30 percent relative to the original monthly condition.

To provide a sense of the dispersion of the final allocations across subjects, we have also plotted histograms, by condition, in Figure I. These reveal that there are a few clever subjects in every condition who realize that putting all of their money in equities is a good strategy in this experiment (and perhaps elsewhere). However, in the monthly condition, only six of the twenty-two subjects put as much as 80 percent of their money in stocks while twelve of the subjects put at least 80 percent of their money in bonds.

Panel B of Table I presents allocations to bonds over the last stages of the trials. To facilitate comparison across conditions, we use the average over the last five years of the experiment, which means the last forty choices in the two monthly conditions, the last five choices in the yearly condition, and the last choice in the five-yearly condition. Inspection of the table reveals that the allocations over the last five years of the trials are very similar to the final allocations. In fact, none of the means within condi-

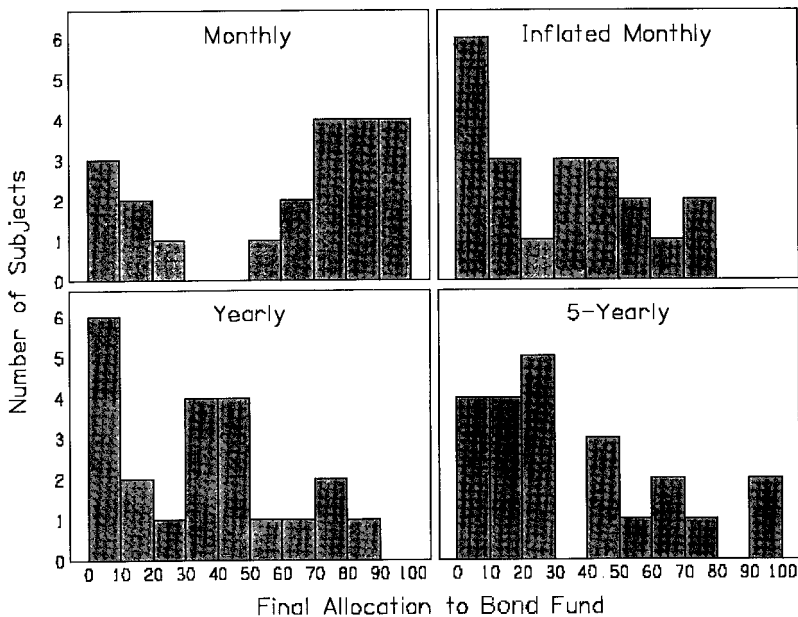


FIGURE I

tion differs between the two measures (using paired-sample *t*-tests). Evidently, subjects learned whatever they were going to learn about the situation in the first 80 percent of the trials. The observation that subjects made exactly the same choices (on average) when they were allocating assets one period at a time or when they were committing themselves for 400 trials is consistent with one of two extreme hypotheses: either the subjects had an infinite horizon or else they were radically myopic. As we already know that the first hypothesis could not be true, we conclude that our subjects were so myopic that they simply did not distinguish the tasks of deciding on an allocation for 400 trials and for a single trial. This result also shows that the subjects do not realize that they suffer from myopic loss aversion—they are unsophisticated about their malady. If they were sophisticated but unable to control themselves, then, when offered the opportunity to commit themselves for 400 trials with no interim feedback, they would take the opportunity to accept more risk.

Another way of looking at the results is to consider how the subjects change their asset allocation over the course of the experiment. Figure II plots the average percentage of funds allocated to the bond fund in each of the four conditions. The upper panel shows the two monthly conditions, while the lower panel shows the three noninflated conditions. The large dot at the end represents the average allocation in the final “long” decision. As one would expect, in all conditions subjects (on average) begin by allocating their money evenly between the two funds. (They have been given no information as yet to help them choose any other allocation.) To see whether the apparent drifts over time in these plots are statistically reliable, we also ran simple regressions. For each condition we pooled all the subjects and regressed the allocation to bonds on trial *t* for subject *j* on the trial number. (The final 400-period choice was not included.) These regressions are shown in Table II. They confirm the patterns of the plots. There is no significant drift in the monthly condition; in the other three conditions, subjects are reliably decreasing their allocations to bonds as the experiment progresses.

IV. DISCUSSION

This paper continues the investigation of myopic loss aversion. The research program involves an interplay between theory, experimental tests, and field data. Benartzi and Thaler [1995]




FIGURE II

took cumulative prospect theory and confronted it with the historic stock and bond return data. Since myopic loss aversion passed the plausibility test, this paper goes back to the laboratory to give the theory additional testing.

The test of loss aversion was straightforward. The translation of all outcomes into the positive domain eliminated the experience of loss, which we consider the most important factor in inducing risk aversion. This manipulation had the expected effect: it caused a substantial increase in the preference for the Stock Fund. The effect of myopia on risk taking was studied by compelling some subjects to adopt a nonmyopic framing of both decisions and outcomes: they had to commit themselves for multiple periods, and received feedback relatively infrequently. This manipulation also eliminated the experience of losses (by statistical aggregation), and it also increased the preference for stocks. Thus, the roles of myopia and loss aversion were both confirmed. Further support is offered by Gneezy and Potters [1997], a similar experiment conducted independently.

The results for the inflated monthly condition were quite strong: adding a constant to returns increased the allocation to

TABLE II
REGRESSIONS PREDICTING ALLOCATION TO BOND FUND FROM TRIAL NUMBER

Variable	<i>B</i>	<i>SE B</i>
Monthly condition (<i>n</i> = 4200)		
Trial number	-0.00	0.01
Intercept	55.9	0.94
Yearly condition (<i>n</i> = 550)		
Trial number	-0.71*	0.16
Intercept	46.2	2.31
Five-yearly condition (<i>n</i> = 110)		
Trial number	-4.64*	1.52
Intercept	48.1	5.04
Inflated monthly condition (<i>n</i> = 4200)		
Trial number	-0.05*	0.01
Intercept	48.9	0.93

p < .01.

stocks by over 30 percent. This manipulation had the predicted effect even though the subjects were told that the high returns were due, in part, to a high level of inflation in the economy. This result can be seen as a laboratory-induced money illusion. (For other examples see Shafir, Diamond, and Tversky [1997].) Money illusion may also be part of the explanation of one aspect of the equity premium puzzle, the “risk-free rate puzzle.” As noted by Weil [1989], since in the standard constant elasticity of substitution model the coefficient of relative risk aversion is the inverse of the elasticity of intertemporal substitution, one would expect to find a high equity risk premium accompanied by a high risk-free rate. Since the risk-free rate has been barely positive in real terms, the standard model has trouble explaining why investors are very reluctant to bear risk but are willing to save at very low real interest rates. One explanation is that money illusion helps endow fixed income assets with an illusory “no loss” attribute that is not present in equities. However, it has not been the case that inflation has been kind to stock market returns, at least in the short run. Clearly, this is a complex question that deserves more attention.

This was a learning experiment, in which subjects had an opportunity to observe the outcomes of two assets that differed in both risk and return, and to take note of the experienced utility that they derived from these outcomes. The subjects experienced the history of a market, at different levels of aggregation for the

different experimental conditions, and they made different choices in the Final period. In a variation on the same theme, Benartzi and Thaler [1996] have manipulated the manner in which historical data are displayed. One group of employees covered by a defined contribution pension plan is shown data on the distribution of annual rates of return for two assets; another group of employees from the same organization is shown a distribution of simulated 30-year rates of return. The latter group chooses to invest much more of their retirement funds in stocks (90 percent versus 40 percent). The results of the two studies are qualitatively similar in confirming the effect of aggregation on risk taking.

Finally, we turn to a practical implication of this research. Myopia and loss aversion may well be general features of human cognition; in general, these features do not produce good decision making. However, some features of the environment can make it easier or more difficult for myopic loss-averse investors to avoid the mistakes to which they are prone. Providing such investors with frequent feedback about their outcomes is likely to encourage their worst tendencies. The subjects in the monthly condition had more information and more freedom than the subjects in the yearly and five-yearly condition, but more is not always better. The subjects with the most data did the worst in terms of money earned, since those with the most frequent data invested the least in stocks (and thus earned the least). This can occur in any domain in which losses are a factor. For example, Kahneman and Lovallo [1993] show that what they call "narrow framing" (e.g., evaluating projects one at a time) can induce severe risk aversion in an organizational context.

The conclusion that the provision of aggregated data and the restriction of opportunities to change decisions help mitigate the effects of myopia, and loss aversion is relevant to some current trends in the financial markets. Defined contribution pension plans are becoming much more important, and employees are being forced to take responsibility for their own retirement asset allocation decisions. The research described here as well as the Benartzi and Thaler [1996] study shows that the decisions made by employees covered by such plans may vary considerably depending on how their investment opportunities are described and the manner and frequency with which they receive feedback on their returns. Only additional research in this domain of the psychology of investment decision making can help us determine the right mix of information to provide to employees.

APPENDIX: EXPERIMENT INSTRUCTIONS

What follows is the instructions that each subject received in the “monthly” condition, and [in brackets] the variations for the “yearly” condition. The additional sentence for the “inflated monthly” condition appears {in curly brackets}.

This experiment simulates a series of investment decisions.

In this experiment, imagine that you are the investment manager for a small college’s endowment. In every financial period [Every eight financial periods] you must decide how you wish to split the college’s assets between two investments, which we label “Fund A” and “Fund B.” You will be asked how much you wish to allocate to Fund A, an integer from 0 (allocate nothing to Fund A) to 100 (allocate all assets to Fund A). Whatever you do not allocate to Fund A is allocated to Fund B.

After each decision you will receive information about the performance of Fund A, Fund B, and your split between them since your decision. The information will be displayed on the screen as three bars showing the percentage return of Fund A, Fund B, and your decision. {There is a high rate of inflation, so both funds will always have a positive return.} You will then be asked to make the decision for the next period [for the next eight periods].

At the end of the experiment, you will be asked to make one final investment decision for a large number of future periods.

Your final payment will be determined from the overall return of all your investment decisions, including the final one.

Your responses in this experiment are very important to us. Please consider your decisions carefully. Please do not write anything down during this experiment. If you have any questions, please feel free to ask the experimenter at any time.

Thank you for your help!

GRADUATE SCHOOL OF BUSINESS, UNIVERSITY OF CHICAGO
 DEPARTMENT OF PSYCHOLOGY, STANFORD UNIVERSITY
 WOODROW WILSON SCHOOL OF PUBLIC AFFAIRS, PRINCETON UNIVERSITY
 DEPARTMENT OF PSYCHOLOGY, UNIVERSITY OF CALIFORNIA, BERKELEY

REFERENCES

- Benartzi, Shlomo, and Richard H. Thaler, “Myopic Loss Aversion and the Equity Premium Puzzle,” *Quarterly Journal of Economics*, CX (1995), 75–92.
 Benartzi, Shlomo, and Richard H. Thaler, “Risk Aversion or Myopia: The Fallacy of Small Numbers and Its Implications for Retirement Savings,” UCLA working paper, 1996.

- Gneezy, Uri, and Jan Potters, "An Experimental Test on Risk Taking and Evaluations Periods," *Quarterly Journal of Economics*, CXII (1997), 631–45.
- Kahneman, Daniel, Jack Knetsch, and Richard H. Thaler, "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy*, XCVIII (1990), 1325–48.
- Kahneman, Daniel, and Dan Lovallo, "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking," *Management Science*, XXXIX (1993), 17–31.
- Kahneman, Daniel, Peter P. Wakker, and Rakesh Sarin, "Back to Bentham? Explorations of Experienced Utility," *Quarterly Journal of Economics*, CXII (1997), 375–405.
- Mehra, Rajnish, and Edward C. Prescott, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, XV (1985), 145–62.
- Samuelson, Paul A., "Risk and Uncertainty: A Fallacy of Large Numbers," *Scientia*, XCVIII (1963), 108–13.
- Shafir, Eldar, Peter Diamond, and Amos Tversky, "Money Illusion," *Quarterly Journal of Economics*, CXII (1997), 341–374.
- Siegel, Jeremy, and Richard H. Thaler, "Anomalies: The Equity Premium Puzzle," *Journal of Economics Perspectives*, XI (1997), 191–200.
- Tversky, Amos, and Daniel Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, V (1992), 297–323.
- Weil, Philippe, "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics*, XXIV (1989), 401–21.