The Journal of Behavioral Finance
2004, Vol. 5, No. 1, 43–56

Adaptive Portfolio Allocation with Options

Uri Benzion, Ernan Haruvy and Tal Shavit

We conducted a laboratory experiment of repeated portfolio allocation choice between a bond, a stock, and a put option on the stock. The study involves two conditions: a full hedging possibility and a partial hedging possibility. Surprisingly, participants were able to converge to the mean-variance frontier in the environment with partial hedging possibilities, but were unable to do so in the full hedging condition. This suggests that subjects may not be cognizant of the mean-variance frontier. If subjects begin away from the frontier and have to adjust toward it, incentives off the frontier are critical. Simulations of adaptive dynamic models confirm this assertion. The study provides insight into the adaptive behavior of investors in the presence of hedging possibilities and implications for efficient investment strategies.

The trade-off between return and volatility is a significant concern to investors, who are typically willing to sacrifice some returns to lower risk. The usual approach of individual investors is to balance portfolios with stocks and bonds. Because these two asset classes have historically exhibited low correlation, and bonds generally have lower volatility (as well as lower expected returns), investors can substantially reduce risk by increasing the proportion of bonds in their portfolios.

Another approach to reducing volatility is to hedge stock market bets with stock options. Writing a covered call option would limit stock gains while guaranteeing some immediate return. Similarly, buying a put option as an insurance policy against a potential decline in stock price can limit losses. But small investors are typically reluctant to take the stock option approach to balancing their portfolios. Stock options add an extra level of complexity, are perceived to be risky, and can be prohibitively expensive to a small investor due to high commissions and fees per contract.

In addition to these concerns, adaptive adjustment of portfolio allocation with options is far from trivial. Investors may know their own preferences for returns and volatility, but they may be uncertain about the expected returns and volatility from various assets. Bonds and stocks (and other uncorrelated assets) allow investors to increase allocation over time to "safer" assets when they feel their portfolios have become "too risky," or to increase allocation to high-yield assets when their portfolios become "too safe."

Uri Benzion is a professor in the Department of Economics at Ben-Gurion University. **Ernan Haruvy** is an assistant professor in the School of Management at the University of Texas at Dallas. **Tal Shavit** is a doctoral candidate in the Department of Economics at Ben-Gurion University.

Requests for reprints should be sent to: Ernan Haruvy, School of Management, University of Texas at Dallas, PO Box 830688, Richardson, TX 75083. Email: eharuvy@utdallas.edu

Options, however, are not as easy to adjust. And, although stocks are risky, options pose significantly more risk. Hence, increasing the proportion of options in a portfolio must be done to arrive at an optimal stock-to-option ratio, or the portfolio's volatility can rise even above that of 100% stock allocation. From a learning perspective, learning to allocate assets when they are correlated is a far more complex task.

Learning theories have become pervasive in the business and economics literature, and, consequently, our knowledge of the way humans make decisions in repeated settings has evolved. We know that humans respond to information by increasing the frequencies of actions that have done well in the past, and decreasing the frequencies of actions that have not done well. Such learning theories explain a wide range of puzzles related to investor behavior and investment decisions, from the equity premium puzzle (Erev and Barron [2001]) to asset market bubbles (e.g., Fisher [1998], Hommes et al. [2002], and Duffy and Unver [2003]).

The puzzle we address here is the inability of investors to hedge effectively using put options. We demonstrate the puzzle by using two experiments. In both, subjects are asked to allocate a stock, a bond, and a put option for the stock. Following each period, returns for all three assets are reported. Instructions clearly spell out, in a simple non-technical way, the possible returns to each asset and the probabilities associated with each return, as well as the interdependencies between the stock and the put option. The experiments differ in that only one allows for full hedging. In the full hedging condition, volatility can be reduced to zero while extracting a yield greater than the bond (the safe asset). In the other condition, only partial hedging is possible—volatility can be lowered but not to zero. When only partial hedging is available, we find that subjects learn to hedge, while in the full hedging condition they do not.

We intend to explore three related questions:

1. Do subjects change the way they invest over time in an intelligent way, so as to yield higher returns? This question has received some attention in the behavioral finance literature (e.g., Benartzi and Thaler [1995] and Thaler et al. [1997]).
2. Do subjects become more efficient investors over time, i.e., do they minimize risk for a given expected return, or maximize expected returns for a given risk (Markowitz [1959])? Using the standard deviation of returns as a proxy for risk, the efficient combinations of risk and returns are on the mean-variance frontier. Accordingly, we evaluate subject allocations relative to the mean-variance frontier over time.
3. Are the learning patterns identified in the first two tasks responsive to past observed returns? To that end, we apply leading models in the learning literature.

Learning Models

We compare several recent studies examining human adjustment to economic incentives and discover that they share several common elements (see, e.g., Cheung and Friedman [1997]; Stahl [1996]; Erev and Roth [1998]; Erev, Bereby-Meyer, and Roth [1999]; Sarin and Vahid [1999, 2001]; Camerer and Ho [1999]; and Fudenberg and Levine [1998]). These elements are 1) a scoring rule, 2) an updating rule, and 3) a mapping rule. We briefly explain each element and describe two leading models in the learning literature, experience weighted attraction and belief learning.

Scoring and Updating Rules

Each possible action j is assigned a score by a decision maker n at time t. The score is based on past performance and is generally a function of j's past payoffs, whether realized only, or both realized and foregone. The score can be interpreted as expected payoff, expected utility, "attraction" of the alternative, or another measure of propensity to choose a given action. A general adaptive expectation process has decision maker n updating her "expected" score of choice j in period t + 1 as follows:

$$A_{nj,t+1} = (1-\mu_{ij}(t))A_{njt} + \mu_{ij}(t)X_{njt} \tag{1}$$

where A_{njt} denotes the score to choice j in period t by decision maker n, and X_{njt} denotes observed payoffs. $\mu_{ij}(t)$ is the adjustment parameter, resulting in the price expectation for a given choice being a weighted average of the previous expectation and the last actual payoff. Sarin and Vahid [1999] provide an excellent discussion of the major updating rules and the appropriate $\mu_{ij}(t)$ specification for each.

A Mapping Rule

The mapping rule of choice is a logit mapping from the "score" of each alternative to the probability of choosing that alternative. Prior to the logit transformation, the score is multiplied by a sensitivity or rescaling parameter λ, whose purpose is to smooth the mapping function so that each choice is selected with some positive probability. That is, as λ tends to 0, all choices are equally likely, and when λ tends to infinity, the best choice is chosen with certainty and no other action is chosen. A_{njt} denotes the score assigned by agent n to action j. P_{njt} denotes the probability of agent n choosing action j. Then the mapping from score to probability is:

$$P_{njt} = \frac{\exp(\lambda A_{njt})}{\sum_{k \in J} \exp(\lambda A_{nkt})} \tag{2}$$

Adjustments to the Scoring Rule

If the decision maker is restricted to integer percentages for each of the three assets, we would have 5,151 possible portfolio allocations to score.[1] The problem can get intractable quickly, but coarser binning can be performed for larger decision sets. In contrast to the extant learning literature, the expected payoff to each possible portfolio is not a good scoring rule, because volatility of returns is very important to investors.

Constant relative risk aversion (CRRA), however, is an attractive alternative that survives this criticism and is intuitively more appealing. The utility function, and therefore the score, is: $U(W) = W^{1-r}/(1 - r)$, where W denotes wealth and r is a coefficient of risk aversion. Since the earnings are already rescaled in practice (by a free parameter to be estimated), we conveniently leave only r in the exponent, so that $U(W) = W^{1-r}$.

With this utility function, investors may care about more than the mean and the variance of returns (overcoming a common criticism of CAPM). That is, changing the distribution of returns while holding the mean and variance of returns constant will affect investor behavior. This rule is easy to apply to a wide variety of situations, and it does not affect lending and borrowing ability (another criticism of CAPM). It is easily incorporated into any learning model with virtually no modification to model updating or model parameters, and the addition of just one risk aversion parameter, r.

Although the CRRA scoring rule is appropriate for estimation, mean-variance analysis is also appropriate here. Mean-variance analysis is most appropriate when returns are normally distributed (which is not the case here nor in actual applications), or when investor preferences are quadratic. Levy and Markowitz [1979] justify mean-variance analysis by showing it can be regarded as a second order Taylor series approximation of standard utility functions, such as the power utility and exponential utility. In accordance, we examine allocations at the beginning and the end of the experiment relative to the mean-variance frontier.

Another alternative is the payoff-utility mapping proposed by Thaler et al. [1997] (following Tversky and Kahneman's [1992] cumulative prospect theory), by which the shape of the value function for monetary incomes indicates slight risk aversion in the domain for gains, and an equivalent risk-seeking in the domain of losses. Specifically, the value function is

$$\begin{aligned} U(x) &= x^{a} && \text{if } x \geq 0 \\ U(x) &= -\eta(-x)^{\beta} && \text{if } x < 0 \end{aligned} \tag{3}$$

where η is a coefficient of loss aversion.

An individual with the preferences described by cumulative prospect theory is only mildly risk averse for gambles that entail only potential gains, but strongly risk averse for gambles that entail potential losses. Essentially, this model adds two parameters to the CRRA model. We found no significant improvement in the fit resulting from such parameters; for the sake of parsimony, we do not pursue this model further. However, note that both Thaler et al. [1997] and Gneezy and Potters [1997] found strong support for this model in repeated investment experiments. With more conditions and stronger emphasis on losses we might be able to replicate those findings.

Experience Weighted Attraction (EWA)

Camerer and Ho's [1999] EWA model is a hybrid of belief and reinforcement learning. This model posits that players evaluate the performance of each possible action in the last period and update their propensities to use each action accordingly. However, the action actually taken by each player receives greater attention in the evaluation process. Hence, actions are reinforced according to past performance, but actual actions receive additional reinforcement.

The choice probabilities are expressed as logistic functions of the propensity, or attraction, of each strategy, A_{njt}, where n indexes the individual and j the strategy:

$$P_{njt} = \frac{\exp(\lambda A_{njt})}{\sum_{k \in J} \exp(\lambda A_{nkt})} \tag{4}$$

where λ, as in the two previous models, can be thought of as a precision parameter, or the sensitivity of players to attractions. The attractions are updated according to the following two dynamics:

$$A_{njt+1} = \{\phi\, N_{nt}\, A_{njt} + [\delta + (1 - \delta)\, \chi_{njt}]\, u_{njt}\}/N_{nt+1} \tag{5}$$

where u_{njt} is the actual payoff, realized or foregone, to action j in period t and

$$N_{nt+1} = \phi\,(1 - \kappa)\, N_t + 1 \tag{6}$$

χ_{njt} is a characteristic function that takes the value of 1 if player n played action j in period t and 0 otherwise, and δ is an "imagination parameter" that determines the ability of a player to consider the possible payoffs if he had selected a different action. There are two decay parameters, ϕ and κ.

Belief Learning (BL)

In addition to EWA, we also study a simple belief learning model. In this model, 1) beliefs are updated with depreciation of past beliefs by a constant, and 2) expected payoffs are normalized by a measure of payoff variability. For the choice function, BL applies the logit probabilistic choice rule:

$$P_{njt} = \frac{\exp(\frac{\lambda}{S_{nt}} \bar{U}_{njt})}{\sum_{k \in J} \exp(\frac{\lambda}{S_{nt}} \bar{U}_{nkt})} \tag{7}$$

where P_{njt} is the probability of player n choosing action j in period t, and $\bar{U}_{njt}$ is the propensity toward action j, also known as "attraction" in EWA. The parameter λ is the precision or sensitivity parameter. Propensities are updated as a weighted average of past observed payoffs. The score $\bar{U}_{njt}$ is the expected payoff from action j in period t, based on weighted past observations of payoffs, with equal weight given to foregone and realized payoffs. That is,

$$\bar{U}_{njt} = (1 - \phi)\ \bar{U}_{njt-1} + \phi\ u_{njt-1} \tag{8}$$

where u_{njt} is the actual payoff, realized or foregone, to action j in period t. Finally, variability is determined by Equations (9)-(11). Let y_{nt} be the realized payoff, and x_{nt} the foregone payoff to player n (from the most attractive alternative option) at time t.

$$S_{n+1} = \frac{tS_{nt} + v_{nt}}{t+1} \tag{9}$$

$$v_{nt} = \min\{| y_{nt} - x_{nt} |, | a_{nt} - y_{nt} |\} \quad (10)$$

$$a_{nt} = \frac{(t-1)a_{nt-1} + y_{nt-1}}{t} \quad (11)$$

An Asset Allocation Experiment With Stock and Put Options

Experimental Procedure

The participants were forty-six undergraduate students, all in their second or third year. All had taken a basic course in statistics. The experiment took place at the Laboratory for Economic Experiments at Ben-Gurion University, and lasted approximately one hour. We used the zTree software, developed at the Institute for Empirical Economics at the University of Zurich. Subjects were divided into two groups: the *full hedging condition*, with twenty-four subjects, and the *partial hedging condition*, with twenty-two subjects.

In each condition, subjects chose allocations between a stock, a bond, and a put option. To avoid any biasing framing effect, we didn't use the terms stock, bond, or put, but labeled the assets neutrally as assets A, B, and C. This approach was also used by Kuon and Abbink [1998] in an experimental study on the option valuation approach.

Subjects repeated the portfolio allocation decision 200 times, with each round followed by an on-screen report of each asset's return in that period as well as the return on the portfolio chosen by the subject in that round. In each round, the subjects were asked to allocate 100 tokens among the three assets. The tokens were converted into NIS at a rate of five tokens to one NIS.[2] To avoid a possible wealth effect, one round per subject was randomly picked at completion of the experiment to determine payment.

Instructions for both conditions clearly spelled out, in a simple and non-technical way, the possible returns to each asset and the probabilities associated with each return, as well as the interdependencies between the stock and put option (see translated versions in Appendix A). After we handed out the written instructions, we read them aloud. We then gave the subjects five minutes to carefully review and understand the examples and to ask questions, at which point we explained the examples on the board and answered the questions.

The Assets

In each of the 200 rounds, the subjects were asked to allocate their endowment among the stock, the bond, and the put option. We used simple lotteries to represent the three assets' rates of returns. The put option, when combined with the stock, gives the investor the opportunity to guarantee a positive return regardless of the return on the stock, although it would be lower than the expected return on the stock alone.

Full hedging condition assets. The assets were presented as follows:

Bond (Asset A)

Probability	Return
1/3	6%
1/3	7%
1/3	8%

The above lottery represents a bond with a mean return of 7% and a standard deviation on the return of 0.816%.

Stock (Asset B) and Put Option (Asset C)

Probability	Asset B (Stock) Rate of Return	Asset C (Put Option) Rate of Return
0.45	30%	–100%
0.45	0%	60%
0.1	–30%	220%

The lottery for asset B represents a stock with a mean return of 10.5% and a standard deviation on the return of 19.61%. The lottery for asset C represents a put option with a mean return of 4% and a standard deviation on the return of 104.61%.

For a stock currently priced at 100, lottery C represents a put option with a price of 18.75 and a strike price of 130. That is, if the stock rises to 130 (a 30% return on the stock), the put option owner will lose his entire investment in the put option (-100% return on the put option). If the stock's value remains at 100 (0% return on the stock), the put option owner would be able to buy the stock in the open market for 100 and sell to the issuer of the option at the option's strike price of 130, so his profit would be 30. A profit of 30 on the original option's price of 18.75 would yield a rate of return of 60%. If the stock's value drops to 70 (a -30% return), the put option owner would be able to sell the stock at the strike price of 130 and buy it for 70 in the open market, so his return would be 220%.

The subjects were given the following possibilities for assets B and C given the randomly generated number:

Random Number	Asset B (Stock) Rate of Return	Asset C (Put Option) Rate of Return
1–45	30%	–100%
46–90	0%	60%
91–100	–30%	220%

Assets B and C are perfectly negatively correlated (a correlation of -1). An investment in assets B and C at a ratio of 5.33 (84.2% in the stock and 15.8% in the put option) would yield a zero variance portfolio that would dominate the return from the bond. Hence, there is a full hedging opportunity. The return to the full hedging portfolio is 9.47%. When a subject invests in the stock and the put option at the ratio of 5.33 to 1, he is faced with the portfolio return schedule presented in Table 1.

The rate of return for this portfolio is 9.47% regardless of the random outcome, so an investment at the stock-to-put ratio of 5.33 to 1 gives the subject a risk-free portfolio with a return of 9.47%. In comparison, the expected return to the bond is 7% with a standard deviation of 0.816%. On the mean-variance frontier for the full-hedging condition, an investment in the stock and the put option at the ratio of 5.33 to 1 gives the efficient allocation (see Figure 1).

Partial hedging condition assets. The assets are presented as follows:

Bond (Asset A)

Probability	Return
1/3	6%
1/3	7%
1/3	8%

This lottery represents a bond with a mean return of 7% and a standard deviation on the return of 0.816%. For all other ratios, the rates of return are different in each state, so the variance is above zero.

Stock (Asset B) and Put Option (Asset C)

Probability	Asset B (Stock) Rate of Return	Asset C (Put Option) Rate of Return
0.6	50%	–100%
0.1	5%	11%
0.3	–50%	215%

The lottery for asset B represents a stock with a mean return of 15.5% and a standard deviation on the return of 44.86%. The lottery for asset C represents a put option with a mean return of 5.6% and a standard deviation on the return of 140.88%.

For a stock priced at 100, lottery C represents a put option with a price of 27 and a strike price of 135. If the stock's value is 150 (a 50% return), the put option owner will lose his entire investment (a -100% return). If the stock's value is 105 (a 5% return), the put option owner will sell the stock at 135 (the strike price) and buy it at 105, so his return would be 30. A return of 30 on a buying price of 27 is a rate of return of 11%. If the stock's value is 50 (a -50% return), the put option owner will sell the

Table 1. *Portfolio Return Schedule–Full Hedging Condition*

Random Number	% of Investment			Portfolio's Rate of Return		
	Stock	Put	Total	Stock	Put	Total
1–45	84.2	15.8	100	30%	–100%	0.842 (30%) + 0.158 (–100%)= 9.47%
46–90	84.2	15.8	100	0%	60%	0.842 (0%) + 0.158 (60%) = 9.47%
91–100	84.2	15.8	100	–30%	220%	0.842 (–30%)+ 0.158 (220%) = 9.47%

FIGURE 1
Mean-Variance Frontier for Full Hedging Condition

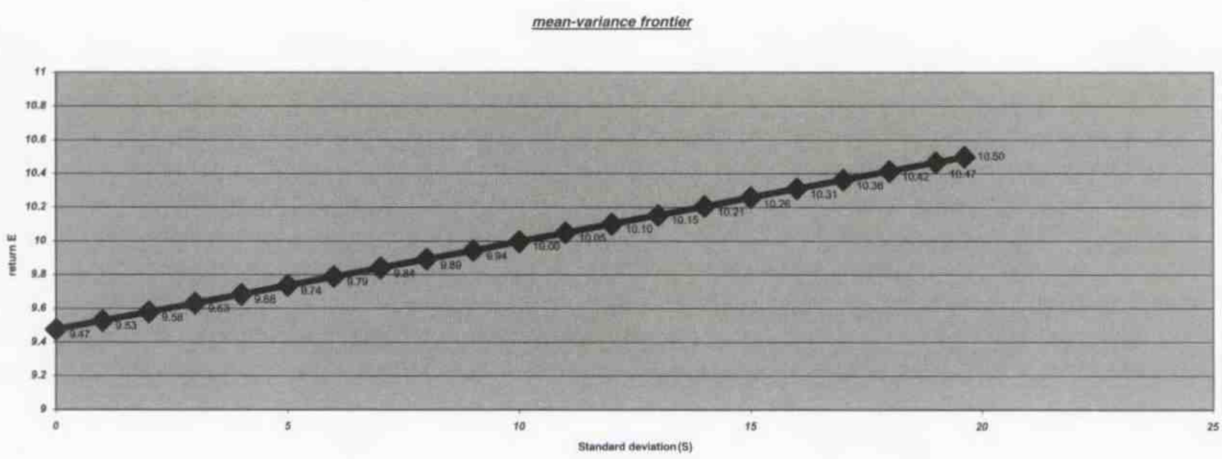

stock at 135 (the strike price), and buy it at 50, so his return would be 85. A return of 85 on a buying price of 18.75 is a rate of return of 215%.

The subjects faced the following possibilities for assets B and C, given the randomly selected number:

Random Number	Asset B (Stock) Rate of Return	Asset C (Put Option) Rate of Return
1–60	50%	–100%
61–70	5%	11%
71–100	–50%	215%

The correlation between assets B and C is –0.997. On the mean-variance frontier, the ratio of assets B and C over most of the range is 3.14 to 1. The expected return in the partial hedging condition for a minimum variance portfolio is 13.1% (the minimum standard deviation of returns is 2.22%). When a subject invests in the stock and the put option at the ratio of 3.14 to 1 (say, 75.8% in the stock, and 24.2% in the put option), he is faced with portfolio presented in Table 2.

The value of the portfolio is dependent on the state of nature, so an investment at the ratio of 3.14 to 1 gives the subject a risky portfolio with a minimum risk of 2.22% and an expected return of 13.1%. In comparison, the expected return to the bond is 7% with a standard deviation of 0.816%. On the mean-variance frontier for the partial hedging condition, an investment in the stock and the put option at the ratio of 3.14 to 1 (along most of the frontier) gives the efficient allocation (see Figure 2).

We explained to the subjects that in each round a number from 1 to 3 would be randomly generated by the computer. Asset A's rate of returns, corresponding to the randomly drawn numbers 1, 2, or 3, would be 6%, 7%, or 8%, respectively. In addition, in each round a second number from 1 to 100 would be randomly generated to determine Asset B's and C's rates of returns.

Findings

Figure 3 shows the actual allocations to the stock, the bond, and the put option in each condition, averaged over individuals and periods in each of five blocks of forty periods. The actual allocations are plotted next to simulated averages by EWA and BL. The learning models' simulated averages show that in the full hedging condition learning is expected to be virtually flat. The actual experimental data confirm this prediction. In fact, a pairwise t-test of individual allocations in the first and last blocks of forty periods finds that the average stock allocations are not significantly different between the two blocks (H_0: average stock allocation of individual i in block 1 = average stock allocation of individual i in block 5). In the full hedging condition, the t-statistic is 1.04 (d.f. 23), with a one-tailed p-value of 0.155.

In contrast, the simulated allocations in the partial hedging condition suggest slow yet steady learning toward portfolios with increased stock proportions and reduced bond and put holdings. The experimental data appear to confirm this assertion. The pairwise t-test of

Table 2. *Portfolio Return Schedule–Partial Hedging Condition*

Random Number	% of Investment			Portfolio's Rate of Return		
	Stock	Put	Total	Stock	Put	Total
1–60	75.8	24.2	100	50%	–100%	0.758 (50%) + 0.242 (–100%) = 13.7%
61–70	75.8	24.2	100	5%	11%	0.758 (5%) + 0.242 (11%) = 6.45%
71–100	75.8	24.2	100	–50%	215%	0.758 (–50%)+ 0.242 (215%) = 14.1%

FIGURE 2
Mean-Variance Frontier for Partial Hedging Condition

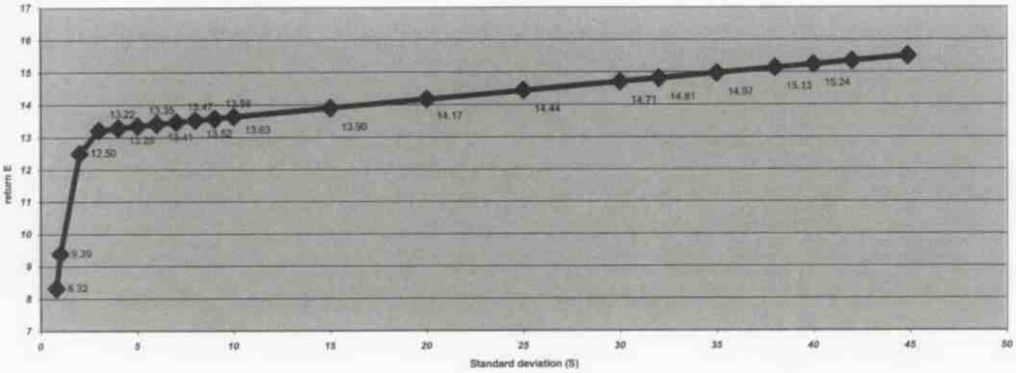

FIGURE 3
Portfolio Allocations Over Time

Partial Hedging Condition

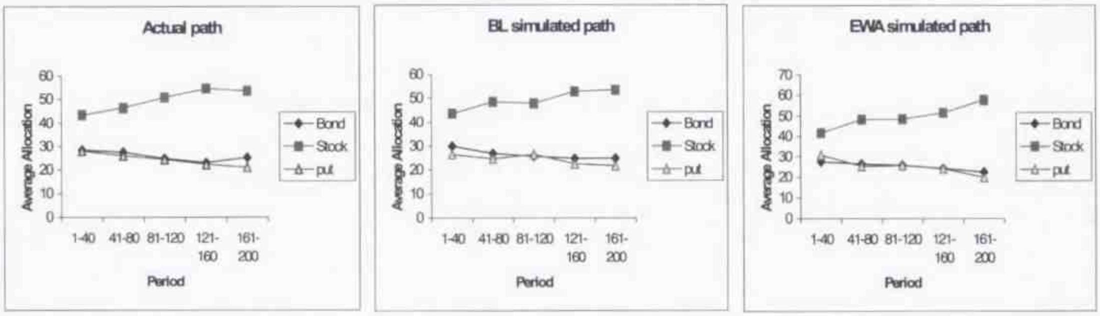

Full Hedging Condition

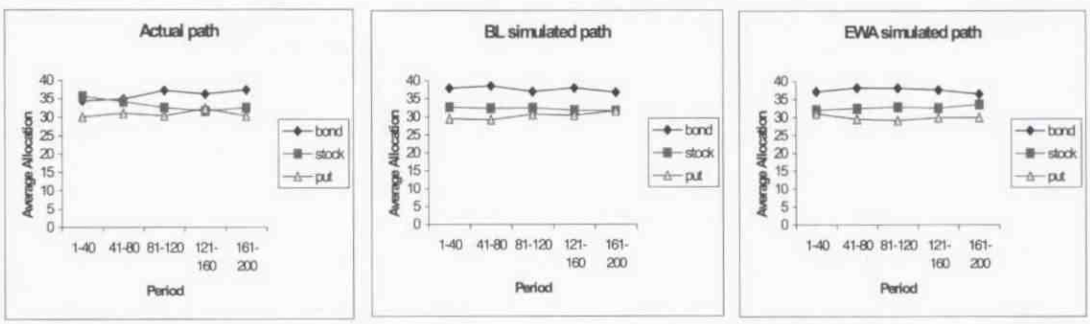

individual allocations in the first and last blocks of forty periods finds that the average allocations are significantly different. The partial hedging condition's t-statistic is 2.11 (d.f. 21), with a one-tailed p-value of 0.024.

We derived the parameter estimates for the simulations of both conditions by minimizing the mean squared deviation of actual individual allocations in each category in each period from the simulated mean response in each category in each period under the partial hedging condition (see Table 3). Note that the risk aversion parameter, r, is relatively high and significantly different from 0 at the 5% level in both models.

Second, while the EWA inertia parameters (κ, ϕ, δ) predict rapid best response movement in the beginning and some moderate flattening later on, the BL model,

Table 3. *Parameter Estimates of the Learning Models under the Partial Hedging Condition*

EWA		BL	
Parameter	**Estimate**	**Parameter**	**Estimate**
κ	0.015	ϕ	0.142
ϕ	0.979	λ	3.263
δ	0.986	r	0.313
N_0	11.38		
λ	13.10		
r	0.147		
MSD =	1,235	MSD =	1,242

with three fewer parameters and a single inertia parameter, has a diminished ability to capture the inertia structure. Yet it predicts a similar learning curve. Note also that by the EWA interpretation of the δ parameter, most of the learning is belief learning rather than reinforcement learning. That is, if one follows Camerer and Ho's [1999] interpretation of EWA parameters, subjects pay close attention to foregone payoffs — payoffs from alternative portfolio allocations not selected. In fact, the interpretation would suggest that attention to foregone payoffs is virtually identical to attention to realized payoffs.

Figure 4 plots the average individual allocations in the first and last blocks of forty periods next to the two learning models' predictions for twenty-four simulated individuals. The models do well at capturing the average allocation paths in Figure 3 (also see the goodness of fit statistics in Table 4), but they do relatively poorly at characterizing the distribution of allocations over individuals. There are two main reasons for this failure.

First, the models' parameters were estimated by minimizing distance to the mean simulated path, so the objective function did not take other moments into consideration. Second, the models themselves do not have parameters for heterogeneity, nor do they allow for heterogeneity in learning or in risk aversion. For example, gender effects were marginally significant in the full hedging condition in the last block,

FIGURE 4A

Stock-Put Allocations for Each Subject in the First and Last Forty Periods of the Partial Hedging Condition, Plotted Against the Mean-Variance Frontier

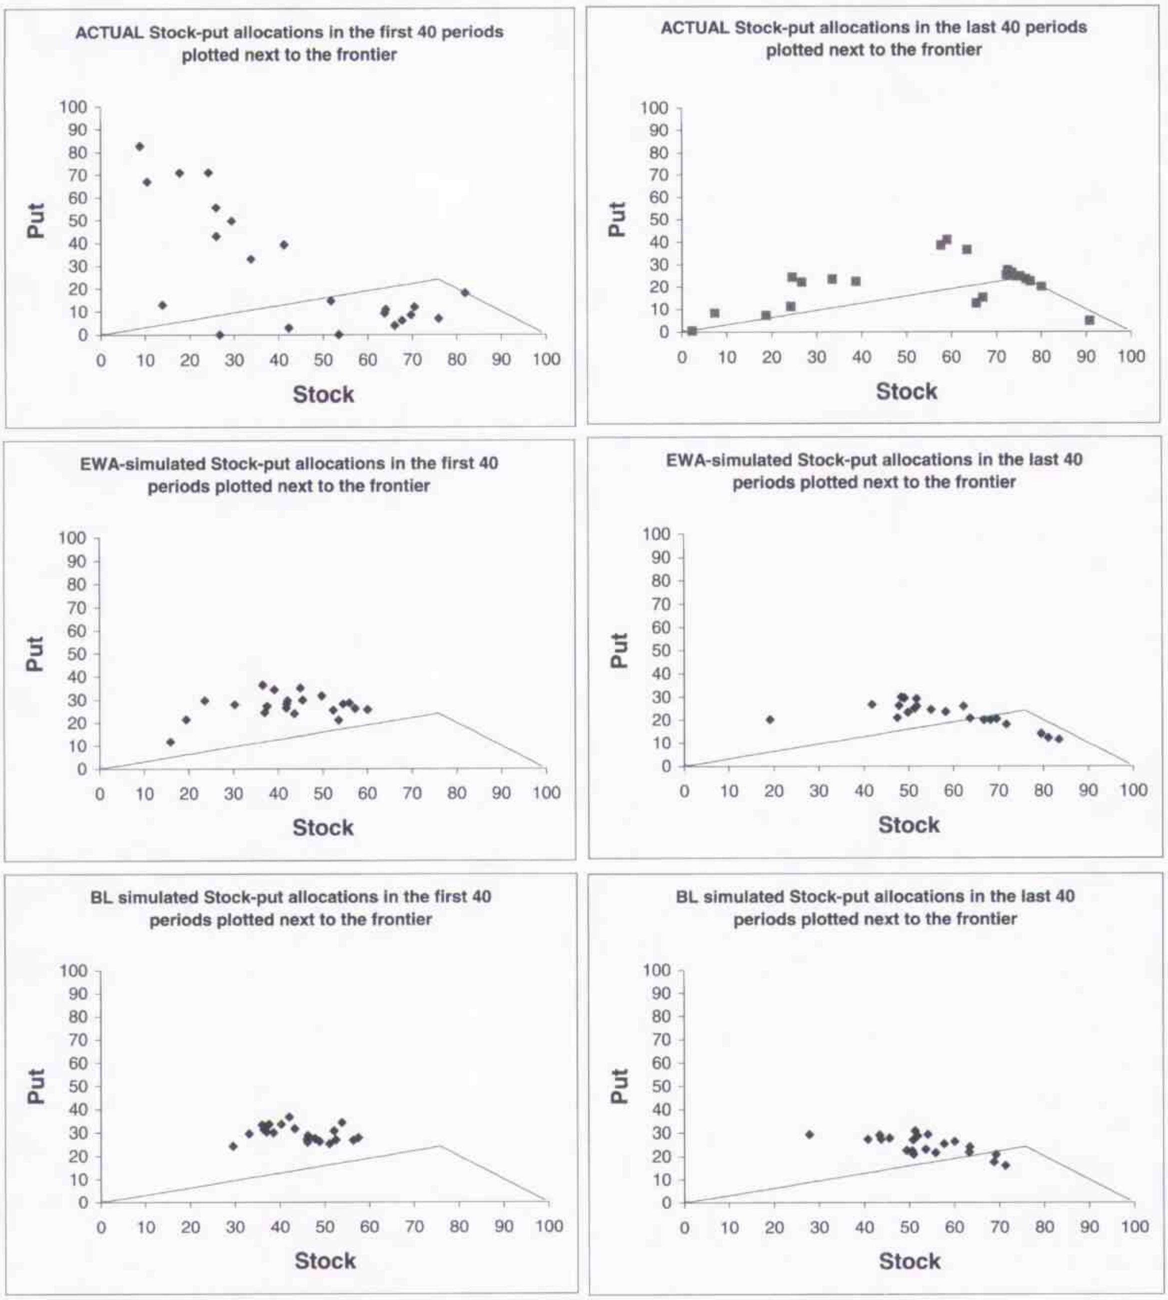

using stock allocations as the dependent variable (two-tailed t-test p-value: 0.064; two-sided rank sums test p-value: 0.030). But they were not significant in the partial hedging condition (two-tailed t-test p-value: 0.227; two-sided rank sums test p-value: 0.139). Based on past research (e.g., Barber and Odean [2001]; Powell and Ansic [1997]), with more observations these gender differences may become more significant. This is a limitation of the current specification, but heterogeneity in risk preferences and learning is outside the scope of this work.

To examine whether subjects learn to become more efficient investors, we estimate the Sharpe [1964] index for the first and last ten periods in each condition. We substitute the bond's risk and return for the risk-free asset in the Sharpe index, as follows:

$$\textit{Sharpe Index} = [E(R) - E(B)]/[S(R) - S(B)]$$

where E(R) is the portfolio's expected return, S(R) is the portfolio's standard deviation, E(B) is the bond's expected return, and S(B) is the bond's standard deviation.

FIGURE 4B

Stock-Put Allocations for Each Subject in the First and Last Forty Periods of the Full Hedging Condition, Plotted Against the Mean-Variance Frontier

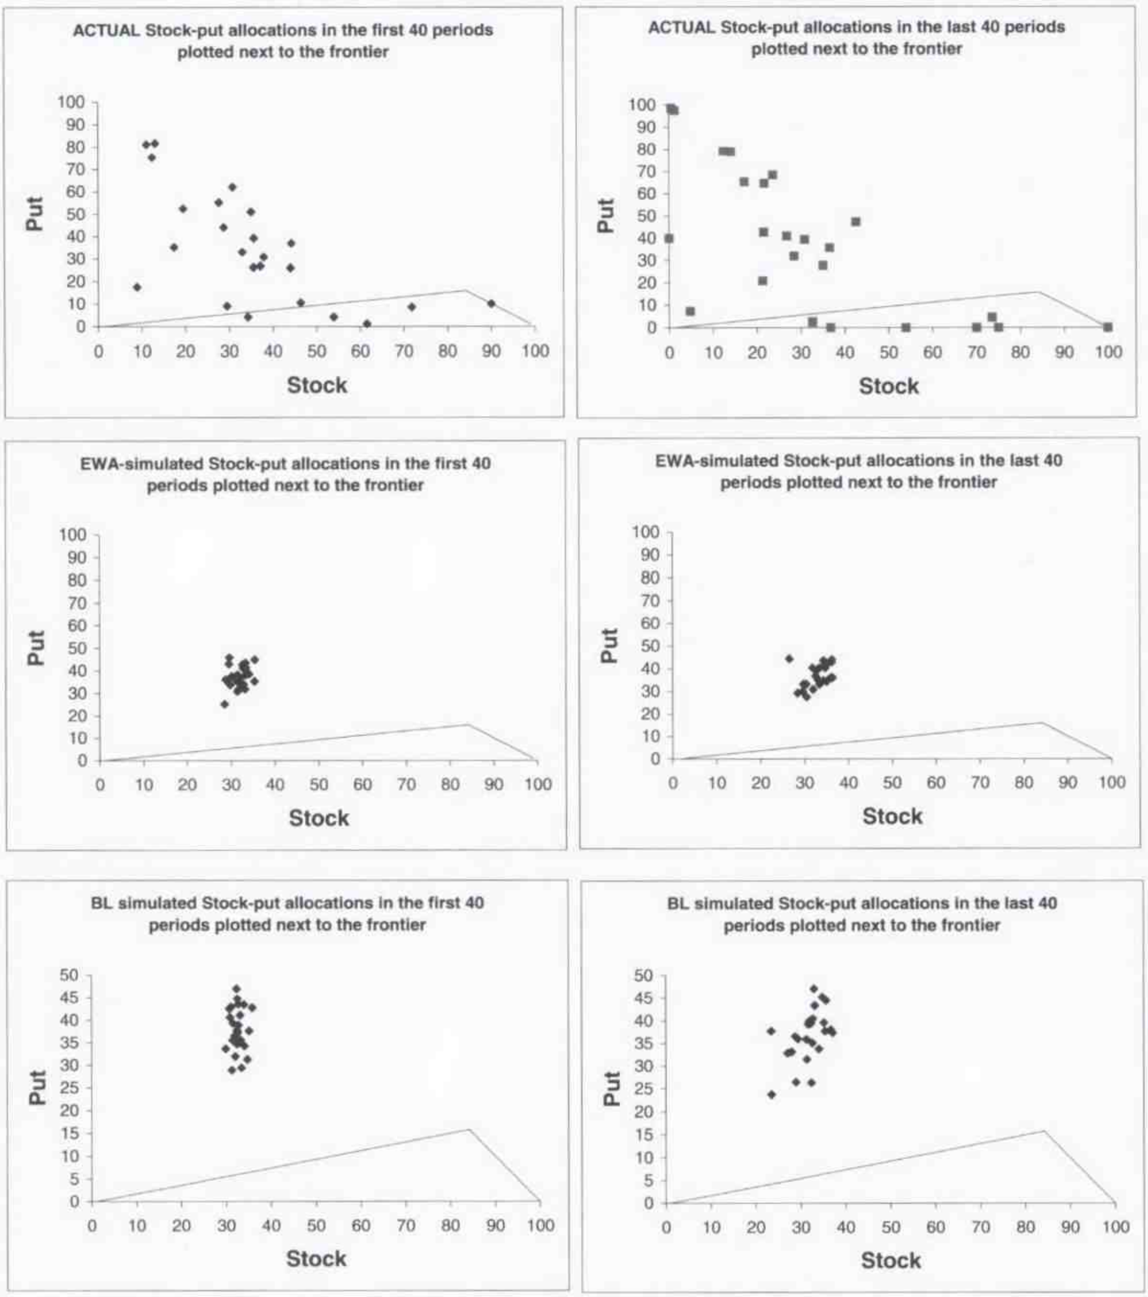

We calculate the expected return and average standard deviation in the first and last ten periods for each subject in each condition. In the full hedging condition, the average Sharpe index in the first ten periods is 0.0376. In the last ten periods, the average is 0.0782. The t-statistic is 0.89 (d.f. 23), and the one-tailed p-value is 0.19.

Figure 5 shows the average Sharpe index for each subject in the first and last ten periods. It is clear that in the full hedging condition subjects did not learn to become more efficient investors. In the partial hedging condition, the average Sharpe index in the first ten periods is 0.295, and 0.68 in the last ten periods. The t-statistic is 3.59 (d.f. 21), and the one-tailed p-value is 0.0008. Figures 5 and 6 show the average Sharpe index for each subject in the first and last ten periods of each condition.

We also determined the ratio between the average investment in stock and the average investment in the put option over individuals and periods in each of five blocks of forty periods. Figures 7 and 8 show the ratio for the average over individuals and periods in each of

Table 4. *Root Mean Squared Error (RMSE) for each model's prediction relative to the experiment under each setting (using proportions rather than percentages as the units of measurement). The top number is the RMSE of the simulation prediction relative to the experiment's average choices. The bottom number is the RMSE of the simulation prediction relative to individual subjects' allocation choices.*

$$\mathrm{RMSE}_1 = \sqrt{\frac{\sum_{t=1}^{200}(bond_t^{pred} - \frac{\sum_i bond_{it}^a}{N})^2 + (stock_t^{pred} - \frac{\sum_i stock_{it}^a}{N})^2 + (put_t^{pred} - \frac{\sum_i put_{it}^a}{N})^2}{200 \times 3}}$$

$$\mathrm{RMSE}_2 = \sqrt{\frac{\sum_{t=1}^{200}\sum_i (bond_t^{pred} - bond_{it}^a)^2 + (stock_t^{pred} - stock_{it}^a)^2 + (put_t^{pred} - put_{it}^a)^2}{N \times 200 \times 3}}$$

[a]denotes actual choice in the experiment. [pred]denotes predicted choice by the simulation. [i]denotes the individual. N is the number of individuals in the experiment.

	Partial Hedging	Full Hedging
EWA	RMSE_1 =0.044	RMSE_1 =0.061
	RMSE_2 =0.334	RMSE_2 =0.309
BL	RMSE_1 =0.046	RMSE_1 =0.065
	RMSE_2 =0.334	RMSE_2 =0.309

Note. For reference, a uniform prediction (1/3, 1/3, 1/3) would yield an RMSE_2 of 0.481 in the partial hedging condition, and 0.451 in the full hedging condition.

FIGURE 5
Full Hedging Sharpe Index for Each Subject

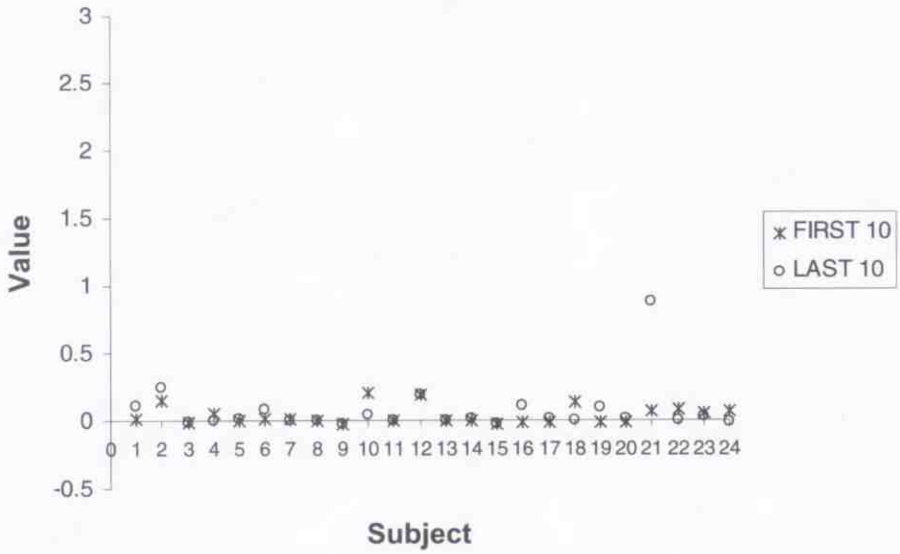

five blocks of forty periods. In the full hedging condition, subjects do not change the ratio; in the partial hedging condition, subjects increase their ratio and appear to converge to the optimal ratio. Clearly, subjects in the partial hedging condition learn to become efficient investors over time, while those in the full hedging condition do not significantly improve over time.

Conclusions

We have shown, both theoretically and experimentally, that adaptive behavior in portfolio allocation can lead to puzzling implications. In a setting with the ability to fully hedge, subjects do not converge to optimal allocations; in a setting with no possibility for full

FIGURE 6
Partial Hedging Sharpe Index for Each Subject

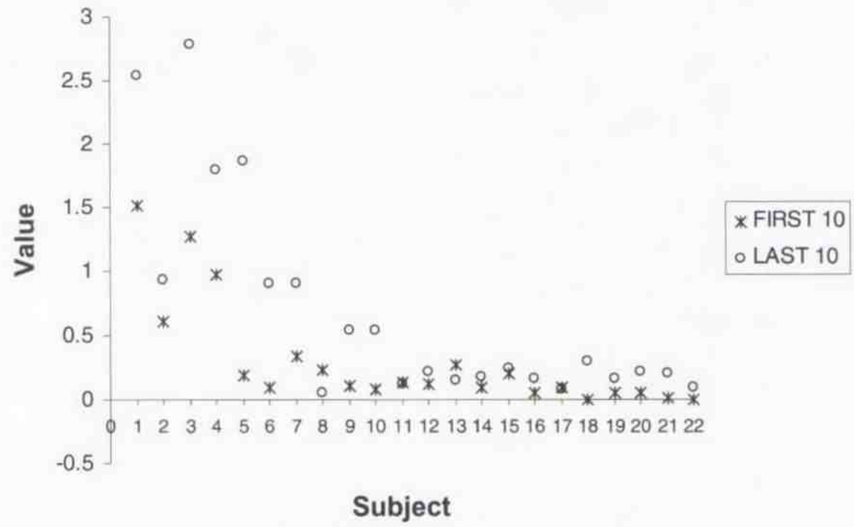

FIGURE 7
Stock-Put Ratio in the Full Hedging Condition

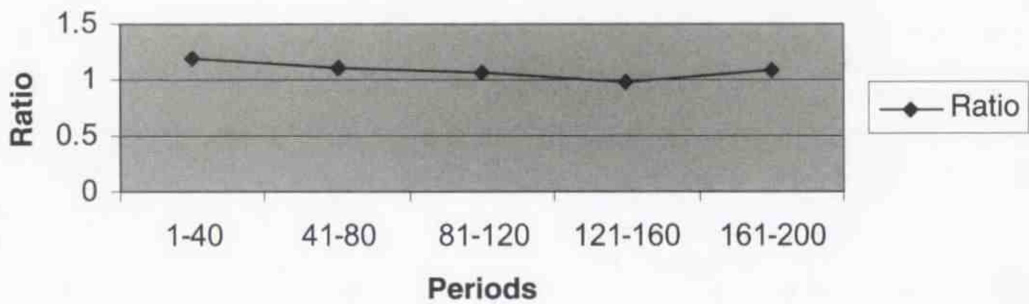

FIGURE 8
Stock-Put Ratio in the Partial Hedging Condition

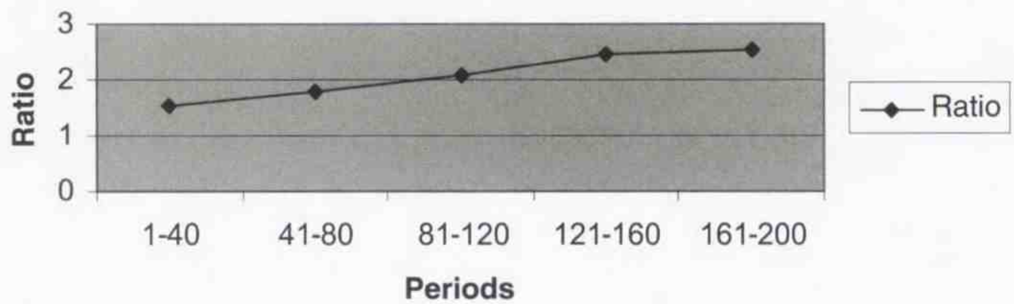

hedging, subjects quickly reach the mean-variance frontier. This suggests that returns on the mean-variance frontier are only critical if subjects are cognizant of that frontier. If subjects begin away from the frontier and have to adjust toward it, incentives off the frontier are critical. In the treatments shown here, incentives were flatter in the full hedging condition, thus the subjects' learning curve was flatter as well. By examining two prominent learning models in the economics literature we find evidence to support this explanation.

Note that the experimental subjects were presented with the return information in a clear manner, both verbally and in a tabulated format, with the correlations apparent and plenty of opportunities to digest the information and ask questions. They were all trained in statistics, although such training was not essential to make informed decisions. Nevertheless, their actions were boundedly rational and strongly adaptive.

One criticism about inference from the laboratory to real-life investment environments might be that the stakes were relatively low in this experiment compared to the real-life stakes of one's life savings or retirement plan. We cannot dismiss such criticism. Experimental data sets are not perfect substitutes for field data, and future empirical research should collect such data. Note that higher stakes are clearly expected to raise the level of rationality and the amount of calculation performed in the initial period. However, small periodic adjustments to the portfolio are not likely to be of the same magnitude. Furthermore, the information encountered in real investment environments is far more complex than what the participants encountered here, and much harder for the average small investor to process. Thus, dynamic adjustment of the type proposed here, rather than an alternative to perfectly rational optimization, may be just as relevant relative to the laboratory environment.

It is important to note that learning models and their insights should not be applied indiscriminantly. Several authors have stressed the critical difference between frequent and infrequent investing due to the differing frequency of portfolio evaluation (Benartzi and Thaler [1995]; Thaler et al. [1997]; Odean [1998, 1999]). Benartzi and Thaler [1995] and Thaler et al. [1997] note that due to myopic loss aversion the attractiveness of the risky asset depends on the frequency of portfolio evaluation. The more frequently the investment portfolio is evaluated, the less attractive the risky asset appears. Odean [1998, 1999] investigated investors with discount brokerage accounts and found that they are prone to trade too frequently, leading to what appears to be a *disposition effect*—an irrational reluctance to realize losses manifested by simultaneously holding on to losing stocks and rushing to sell winning stocks.

Learning theory, as we have outlined it, should be more readily applied to investors who frequently evaluate and update their portfolios. Reinforcements that are not immediate and frequent might be learned, but they may not necessarily extend in a straightforward way from models derived from laboratory studies, with all the biases and boundedly rational behavior observed there.

The same caution would hold for the disposition effect and myopic loss aversion, as the authors of those respective studies were careful to note. Before the advent of the internet, relatively few individual investors could be classified as "frequent." Most individual investors were long-run investors through retirement programs and mutual funds, and reviewed their investment programs relatively infrequently. Today, with easy internet access to retirement portfolios, mutual funds, and brokerage accounts, the landscape of small investing has changed dramatically. Small investors can easily review their investment decisions on a weekly, daily, or even hourly basis, with little effort or expense.

In light of the continuing and rapid trend toward greater individual command over investment decisions, and more frequent evaluation of these decisions, models of dynamic adjustment based on reinforcements should be increasingly appropriate and useful.

Acknowledgments

The authors are grateful to Ido Erev and Farshid Vahid for helpful comments and suggestions.

Notes

1. This is a right triangle in two dimensions (x, y), with 101 integers (0-100) in each dimension, excluding all coordinates for which $x + y > 100$.
2. The average payment in the full hedging condition was 23.6 NIS. The average payment in the partial hedging condition was 22.8 NIS. The conversion rate at the time of the experiment was approximately 4.7 NIS per dollar.

References

Barber, B., and R. Odean. "Boys will be Boys: Gender, Overconfidence, and Common Stock Investment." *Quarterly Journal of Economics,* 116, (2001), pp. 261–292.

Benartzi, S., and R.H. Thaler. "Myopic Loss Aversion and the Equity Premium Puzzle." *Quarterly Journal of Economics*, 110, (1995), pp. 75–92.

Brown, G.W. "Iterative Solution of Games by Fictitious Play." *Activity Analysis of Production*. New York: John Wiley & Sons, 1951.

Camerer, C., and T. Ho. "Experience-Weighted Attraction Learning in Normal Form Games." *Econometrica*, 67, (1999), pp. 827–874.

Cheung, Y., and D. Friedman. "Individual Learning in Normal Form Games: Some Laboratory Results." *Games and Economic Behavior,* 19, (1997), pp. 46–76.

Duffy, J., and U. Unver. "Asset Price Bubbles and Crashes with Near-Zero-Intelligence Trader." Working paper, Harvard Business School, 2003.

Erev, I., and G. Barron. "On Adaptation, Maximization, and Reinforcement Learning Among Cognitive Strategies." Working paper, Columbia Business School, 2001.

Erev, I., Y. Bereby-Meyer, and A. Roth. "The Effect of Adding a Constant to All Payoffs: Experimental Investigation, and Implications for Reinforcement Learning Models." *Journal of Economic Behavior and Organization,* 39, (1999), pp. 111–128.

Erev, I., and A.E. Roth. "Predicting How People Play Games: Reinforcement Learning in Games with Unique Strategy Equilibrium." *American Economic Review,* 88, (1998), pp. 848–881.

Fisher, E. "Explaining Bubbles in Experimental Asset Markets." Working paper, Ohio State University, 1998.

Fudenberg, D., and D.K. Levine. "The Theory of Learning in Games." MIT Press, 1998.

Gneezy, U., and J. Potters. "An Experimental Test on Risk Taking and Evaluation Periods." *Quarterly Journal of Economics*, 112, (1997), pp. 631–645.

Hommes, C., J. Sonnemans, J. Tuinstra, and H. van de Velden. "Expectations and Bubbles in Asset Pricing Experiments." Working paper, University of Amsterdam, 2002.

Kuon, B., and K. Abbink. "An Option Pricing Experiment with Professional Traders." Discussion Paper b-426, University of Bonn, 1998.

Levy, H., and H.M. Markowitz. "Approximating Expected Utility by a Function of Mean and Variance." *American Economic Review*, 69, (1979), pp. 308–317.

Markowitz, H. *Portfolio Selection: Efficient Diversification of Investments*. John Wiley & Sons, 1959.

Mehra, R., and E.C. Prescott. "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, 15, (1985), pp. 145–162.

Odean, T. "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance*, 53, (1998), pp. 1775–1798.

Odean, T. "Do Investors Trade Too Much?" *American Economic Review*, 89, (1999), pp. 1279–1298.

Powell, M., and D. Ansic. "Gender Differences in Risk Behavior in Financial Decision Making: An Experimental Analysis." *Journal of Economic Psychology*, 18, (1997), pp. 605–628.

Sarin, R., and F. Vahid. "Payoff Assessments without Probabilities: A Simple Dynamic Model of Choice." *Games and Economic Behavior*, 28, (1999), pp. 294–309.

Sarin, R., and F. Vahid. "Predicting How People Play Games: A Simple Dynamic Model of Choice." *Games and Economic Behavior*, 34, (2001), pp. 104–122.

Sharpe, W.F. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance*, 19, (1964), pp. 425–442.

Stahl, D. "Boundedly Rational Rule Learning in a Guessing Game." *Games and Economic Behavior*, 16, (1996), pp. 303–330.

Thaler, R.H., A. Tversky, D. Kahneman, and A. Schwartz. "The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test." *Quarterly Journal of Economics*, 112, (1997), pp. 647–661.

Tversky, A., and D. Kahneman. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, 5, (1992), pp. 297–323.

Appendix A. *Translated Instructions for the Full Hedging Condition*

- Welcome to an experiment in decision-making.
- In this experiment, you will be asked to allocate tokens among three assets.
- There will be 200 investment rounds.
- In each round, you will get the same initial endowment to invest in the assets.

Asset A

In each round, the experiment's software will randomly generate a number from 1 to 3 to determine the payoff for asset A.
If the chosen number is 1, you will get a 6% return on investment in asset A.
If the chosen number is 2, you will get a 7% return on investment in asset A.
If the chosen number is 3, you will get an 8% return on investment in asset A.

Assets B and C

In each round, the experiment's software will randomly generate a number from 1 to 100. This random number will determine payoffs for both assets B and C.

Asset B

If the chosen number is from 1 to 45, you will get a 30% return on investment in asset B.
If the chosen number is from 46 to 90, you will get a 0% return on investment in asset B.
If the chosen number is from 91 to 100, you will get a –30% return on investment in asset B.

Asset C

If the chosen number is from 1 to 45, you will lose all of your investment in asset C (a return of –100% on investment in asset C).
If the chosen number is from 46 to 90, you will get a 60% return on investment in asset C.
If the chosen number is from 91 to 100, you will get a 220% return on investment in asset C.

The Initial Endowment

In each round, you will get an initial endowment to allocate among the three assets. You must invest all of the initial endowment.

Hypothetical Example

You have forty tokens to allocate among assets A, B, and C.

Total amount to A {A} ____________________

Total amount to B {B} ____________________

Total amount to C {C} ____________________

The Portfolio Final Value in Each Round

At the end of each round, the portfolio's final value will be calculated according to your investments and returns on each asset.

(continued)

Appendix A. *(continued)*

Suppose you invest ten tokens in asset A, twenty in asset B, and ten in asset C. If the randomly generated number for asset A is 2, and the randomly generated number for assets B and C is 47, at the end of the round the value and return from each asset will be as follows:

A: Initial investment was 10, value of A is now 10.7, rate of return is 7%.
B: Initial investment was 20, value of B is now 20, rate of return is 0%.
C: Initial investment was 10, value of C is now 16, rate of return is 60%.

In addition, you will receive information on your portfolio's final value as follows:

Your total worth at the end of the period is 46.7 tokens.
{Final value calculation: 10.7 (asset A value) + 20 (asset B value) + 16 (asset C value) = 46.7}

Hypothetical Example 2
Suppose you invested ten tokens in asset A, twenty in asset B, and ten in asset C. If the random number for asset A is 3, and the random number for assets B and C is 93, the value and return from each asset will be as follows:

A: Initial investment was 10, value of A is now 10.8, rate of return is 8%.
B: Initial investment was 20, value of B is now 14, rate of return is -30%.
C: Initial investment was 10, value of C is now 32, rate of return is 220%.

In addition, you will receive information about your portfolio final value as follows:

Your total worth at the end of the period is 56.8 tokens.
{Final value calculation: 10.8 (asset A value) + 14 (asset B value) + 32 (asset C value) = 56.8}

After seeing the information on the current round's payoff you will begin a new round. In the new round you will receive a new initial endowment with no relation to your payoffs in previous rounds. The payoffs at the end of each round do not accumulate.

Payment for Experiment Participation
At the end of the experiment, one of the rounds you participated in will be randomly selected by the experiment's software, and payment for your participation will be that portfolio's final value (in tokens) in the selected round. The tokens will be converted into NIS at the rate of five tokens to one NIS.

Hypothetical Example
Suppose the portfolio final value of the selected round is 56.8 tokens. Payment for participation will then be 11.36 NIS (56.8/5 = 11.36).

Personal Details
Please provide the following:
Gender: ________________
Age: ______