


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.


In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Journal of Economic Psychology

journal homepage: www.elsevier.com/locate/joep

Adaptive behavior leads to under-diversification

Uri Ben Zion^a, Ido Erev^b, Ernan Haruvy^{c,*}, Tal Shavit^d^a Department of Economics, Ben-Gurion University, Beer-Sheva 84105, Israel^b Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel^c School of Management, University of Texas at Dallas, SM 32, P.O. Box 830688, Richardson, TX 75083, USA^d The School of Business Administration, College of Management, Rishon LeZion 75190, Israel

ARTICLE INFO

Article history:

Received 23 November 2009

Received in revised form 28 July 2010

Accepted 18 August 2010

Available online 20 September 2010

JEL classification:

C99

D81

D83

PsycINFO classification:

2343

Keywords:

Diversification

Learning

Adaptation

ABSTRACT

In a given period, a diversified fund, by virtue of being a weighted average, will perform somewhere in the middle range of its components' respective performances. This means that adaptive investors who look to the past to adjust expectations about future returns will shun diversified funds. That is, adaptive reaction to feedback implies under-diversification when the investor gets complete feedback on the performance of the diversified fund as well as its components in a given period. Three laboratory experiments and one quasi field experiment explore this possibility and its implications. We find that the availability of complete feedback drastically reduces diversification. Under-diversification is observed even when the decision makers receive a complete description of the payoff distributions and when under-diversification lowers expected return.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Studies of adaptation in economic settings have resulted in different models (see overview in Erev & Haruvy, 2010) that are based on different assumptions and processes, but the common element to nearly all adaptive models is the law of effect (Thorndike, 1898). The law of effect states that choices that have led to good outcomes in the past are more likely to be repeated in the future, whereas choices that have led to unpleasant outcomes in the past are likely to decline in the future. We use the term adaptive behavior in this work to describe behavior that follows the law of effect.

In environments with probabilistic outcomes, it is not necessarily the case that the adaptive behavior guarantees maximization of expected return. This is most likely to be the case when the payoff-maximizing choice has a low variance and the alternative choices have high variances. This type of scenario is especially prevalent in financial markets where risky assets can lead to really high payoffs for a time and individual investors or even institutions that exhibit adaptive behavior will invest in these assets (e.g., subprime mortgage pools), sometimes resulting in a global financial meltdown as happened recently.

In financial markets, adaptive behavior is sometimes referred to as chasing of past returns. There is extensive evidence that investors chase past returns (Chevalier & Ellison, 1997; Hendricks, Patel, & Zeckhauser, 1993; Ippolito, 1992; Sirri &

* Corresponding author. Tel.: +1 972 883 4865; fax: +1 972 883 6727.

E-mail addresses: benzionu@bgu.ac.il (U.B. Zion), erev@technion.technion.ac.il (I. Erev), eharuvy@utdallas.edu (E. Haruvy), ShavitT@colman.ac.il (T. Shavit).

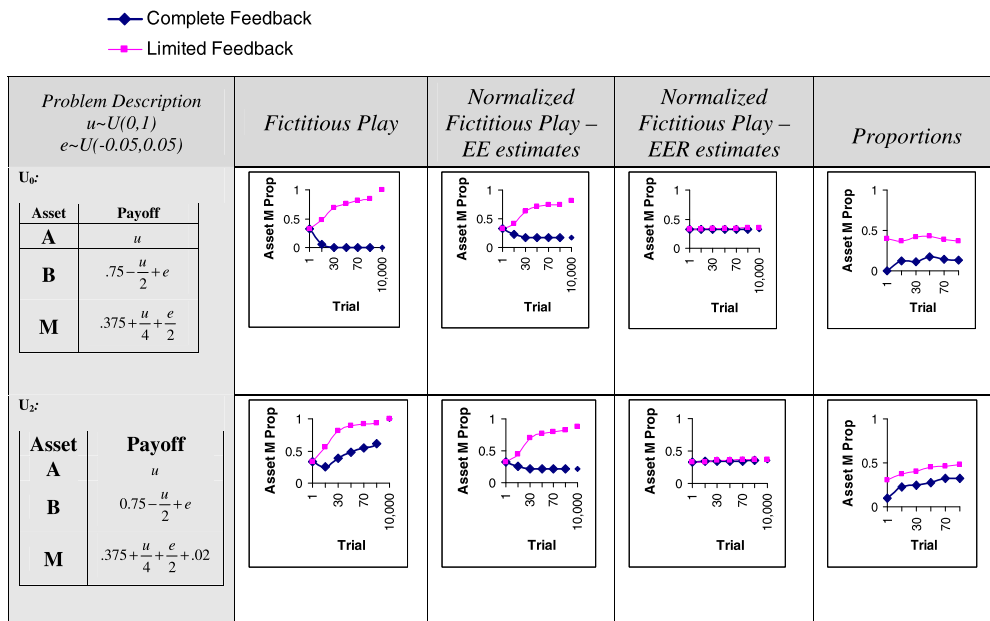


Fig. 1. Conditions U_0 and U_2 . Choice proportions of diversified fund M . The left-hand column presents the three alternatives. The central columns present the predicted choice proportions of M under two learning models as a function of feedback and time. The right column present the results of Experiments 1 and 2 described below.

Tufano, 1998). Recent advances in behavioral finance (e.g., Choi, Laibson, Madrian, & Metrick, 2009) have established adaptive behavior as a key factor in such behavior. Specifically, Choi et al. (2009) showed that investors who experience higher returns on their retirement savings increase those savings more than investors who experience less rewarding outcomes and that the link between returns and adaptation is independent of competing explanations that invoke rationality. Chasing of past returns is also closely related to market overreaction (Chopra, Lakonishok, & Ritter, 1992; De Bondt & Thaler, 1985, 1990; Nasic and Weber, 2009; Offerman & Sonnemans, 2004) where demand overshoots in response to positive recent returns.

A pattern of adaptive behavior can potentially resolve various investment puzzles. For example, Choi, Laibson, and Madrian (2010) showed that the willingness to pay excessive fees was driven by the chasing of returns. The main goal of the current paper is to highlight the role of adaptive behavior in under-diversification. Under-diversification is commonly observed (e.g., Blume & Friend, 1975; Goetzmann and Kumar, 2004; Jacobs, Muller and Weber, 2009; Kelly, 1995; Odean, 1999; Polkovnichenko, 2005; Statman, 1987, 2004) but contradicts standard portfolio theory since diversification can reduce volatility without reducing expected returns.¹

The reason adaptive behavior, particularly chasing of returns, leads to under-diversification is that a diversified asset (yielding a weighted average of other assets' returns) can never have a greater yield than the maximum of its components. Note that the implication that return-chasers will shun the diversified funds depends on the assumption of a particular feedback structure. Specifically, if an investor receives feedback each period about the returns on all investment alternatives, then the diversified fund will never be the "winner" in any one period. However, if the investor pays attention to returns only on assets that he owns, then the diversified fund may look much more attractive. This prediction about the relationship between feedback and diversification is what we set out to test in this work.

In a series of experiments, we show that investors are indeed adaptive and that adaptive behavior leads to more under-diversification when feedback is given over all alternatives relative to when feedback is limited to purchased alternatives. Experiments 1 and 2 examine this pattern under imperfect information conditions, where investors learn about asset returns by observing period-by-period outcomes. Experiment 3 examines an environment in which decision makers receive a complete prior description of the payoff distributions, and complete feedback (about the performance of all alternatives) after each trial. We find that the pattern predicted by the adaptive models emerges even under this setting: experience was found to decrease diversification. Experiment 4 highlights the robustness of the results in a quasi field experiment.

2. A simplified problem and adaptation

The current analysis focuses on a simplified multi-period investment problem. In each period, a single investor is asked to invest one unit in one of three assets: A , B or M . Assets A and B represent stocks with the same expected return and negative

¹ Uppal and Wang (2003) have a framework that explains under-diversification with a model that allows for ambiguity regarding the joint distribution of returns.

correlation between their returns. Asset M represents a diversified fund. The return from asset M is the mean of the returns from A and B plus a premium $c \geq 0$.

The left-hand column in Fig. 1 presents two numerical examples. The exact payoffs in these examples depend on two draws from uniform distributions: u is a draw from a uniform distribution between 0 and 1, and ε is uniform between -0.05 , and 0.05 . Asset A pays u , asset B pays $0.75 - \frac{u}{2} + \varepsilon$, and asset M pays the mean of A and B plus a premium $c \geq 0$. The value of c is 0 in condition U_0 , and 0.02 in condition U_2 .

It is natural to begin the investigation of adaptive behavior with a focus on the fictitious play rule (Brown, 1951; Fudenberg & Levine, 1998). Fictitious play assumes that the agent computes the average payoff from each asset, and selects the asset with the highest historically weighted average. Specifically, the historically weighted average of asset k in period $t + 1$ is

$$W_k(t+1) = \frac{t-1}{t} W_k(t) + \frac{1}{t} V_k \quad (1)$$

where $V_k(t)$ is the payoff from k in period t . As noted by Fudenberg and Levine (1998, p. 204) this simple model implies Bayesian reasoning under uniform priors and the belief that payoffs come from a fixed but unknown distribution. Note that this model is parameter-free and thus does not depend on any particular model parameterization.

The dark curve in the top-left cell of Fig. 1 presents the predictions of the fictitious play rule in condition U_0 under complete feedback. It shows that the fictitious play rule moves away from diversified asset M after the first trial. The intuition is simple: Fictitious play selects the alternative with the highest historical weighted average, and since the weighted value of M is the mean of the weighted values of A and B it cannot be the maximum. The bottom left cell considers condition U_2 . In this case, diversified asset M maximizes expected value. The results show slow movement towards the diversified asset. The simulated rate of M choices is about 25% after 10 trials, 65% after 100 trials, and almost 100% after 10,000 trials.

The light curve in Fig. 1 presents the prediction of fictitious play under limited feedback. Eq. (1) still applies, but t represents only periods in which the return to asset k was observed, and the historical weighted average for asset k is updated only after selecting k .

The simulation reveals that when the feedback is limited to the obtained payoff the fictitious play implies a strong preference to select the diversified asset. Denrell and March (2001) refer to this pattern as the hot stove effect. Denrell (2007) shows that a similar pattern is predicted even under the optimal rule assuming a bandit problem.

The center columns in Fig. 1 presents the predictions of the stochastic variant of the fictitious play rule that was found (see Erev et al. (2010), Ert & Erev (2007), and Marchiori & Warglien (2008); and see related idea in Grosskopf, Erev, & Yechiam (2006)) to provide good approximation of human behavior. The exact assumptions of this model, referred to as normalized fictitious play (NFP), are presented in Appendix A. The predictions in the second column were derived with the parameters estimated by Ert and Erev (2007, referred to in the graph as EE) to fit learning with complete feedback.

The third column presents the prediction with the parameters estimated by Erev (2009, hereafter EER) to fit adaptive behavior with limited feedback. The three models offer very different quantitative predictions from one another. But the important result is that all models, despite their differences, capture the key patterns and the direction of the divergence between the conditions. Specifically, the diversified asset M is predicted to be the least popular choice under complete feedback, and the most popular choice under limited feedback.

3. Experiments

The main goal of the experiments presented below is to clarify the relationship between human adaptation and under-diversification. Experiments 1 and 2 examine U_0 and U_2 under the conditions assumed by the learning analysis. Experiments 3 and 4 examine the robustness of the observed pattern.

3.1. Experiment 1

The basic task in Experiment 1 was a 100-trial repeated play of condition U_0 . Participants were 40 paid volunteer undergraduate subjects who had taken at least one course in statistics. The experiments took place at a computer laboratory in a major Israeli university and lasted approximately half an hour. The average final payment was 25 NIS (approximately \$5).

The choice task was presented as an investment problem. The participants were asked to invest 100 tokens in one of the assets in each trial. All 100 tokens had to be invested in one single investment (participants could not split the 100 tokens between assets in any given trial). Participants were told that the tokens they earn would be converted to money at the end of the experiment. The conversion rate was 1 NIS = 200 tokens. Thus, the average payoff per trial was 0.25 NIS. Ten practice trials were conducted before the 100 experimental trials.

The instructions (Appendix B) spelled out, in a simple and non-technical way, the experimental procedure with no feedback about the actual payoff distribution of each asset. After handing out the written instructions, we gave participants time to carefully review and understand the experimental procedure and ask questions.

The participants were divided into two feedback conditions – full feedback and limited feedback – of 20 subjects each. The groups differed with respect to the feedback provided after each trial. After each trial in the full feedback condition, par-

ticipants observed their earnings from all assets. In the limited feedback condition, the feedback was limited to the payoff obtained from the selected asset only.

3.1.1. Experiment 1 Results

The top-right cell in Fig. 1 presents the choice proportions as a function of experience in the two conditions. The diversified option was selected in 14% of the trials in the full feedback condition, and in 39% of the trials in the limited feedback conditions. For the purpose of a statistical test, only the individual proportions can be considered as independent observations. That is, our sample consists of 40 independent individual proportions, 20 in each feedback condition. Under a two-tailed *t*-test with a pooled variance, this difference is significant ($t[38] = 4.9, p < 0.01$). Thus, we reject the null hypothesis of zero difference in favor of the “sensitivity to feedback” hypothesis (positive difference).

In order to clarify the way in which the subjects responded to their past payoffs it is convenient to distinguish between the predictions of two best reply models: the fictitious play model that assumes “best reply to the average”, and the prediction of a model that assumes “best reply to the most recent experience.” In the full (limited) feedback condition, 55% (52%) of individual choices are to the option with the highest (observed) recent payoff. The fictitious play model does almost as well: 51% (47%) of choices are to the option that yielded the highest average (observed) payoff in past trials (fictitious play). These two best reply models do not always agree. In fact, they only intersect 56% (49%) of the trials (largely the earlier trials). When they intersect, 63% (60%) of choices correspond to that prediction. The percentages corresponding to the two separate models are both significantly different from 1/3 (uniform choice) at the 1% significance level. In other words, subjects are more likely than not to exhibit some form of best reply to the past.

Comparison of the observed learning curves and the predictions of the three models reveals high correspondence with previous studies: the results in the full feedback condition are closest to the prediction of the NFP model with parameters estimated on previous studies with full feedback (EE), and the results in the limited feedback conditions are closest to the prediction of the NFP model with estimated parameters from previous studies with limited feedback (EER).²

3.2. Experiment 2

Experiment 2 involved 40 new undergraduate participants. This experiment is a direct replication of Experiment 1 with one exception: it examines behavior in condition U_2 . Whereas in U_0 , the three assets have equal expected returns, U_2 makes the diversified asset's return strictly higher in expectation (the value of c is 0.02).

3.2.1. Experiment 2 Results

The bottom-right cell in Fig. 1 presents the choice proportions as a function of experience in the two conditions. The diversified option was selected in 27.8% of the trials in the full feedback condition, and in 43.3% of the trials in the limited feedback conditions. The difference reflects a negative effect of complete feedback on earnings. Under a two-tailed *t*-test with a pooled variance, under the null hypothesis of zero difference, this difference is significant ($t[38] = 2.19, p = 0.035$).

Comparison of the observed learning curves and the predictions of the three models reveals the pattern observed in Experiment 1: the results in the full feedback condition are closest the prediction of the NFP model with the parameters that best fit previous studies with full feedback (EE), and the results in the limited feedback conditions are closest to with the parameters that best fit previous studies with limited feedback (EER).

As in Experiment 1, best reply to the past was clearly evident in individual choices. However, the percentages of choices corresponding to best reply, while significantly different from being coincidental, are smaller than the corresponding percentages in Experiment 1. This is to be expected given the higher incentive to choose the diversified option. The full (limited) feedback condition had 45% (49%) of individual choices match the option with highest recent (observed) payoff. Likewise, 44% (47%) of individual choices corresponded to the option with the highest average (observed) payoff up to that point.

3.3. Experiment 3

The significance of the analysis presented above can be questioned on the ground that the assumption of uninformed priors is not likely to hold in the context of investment decisions. Investors are likely to know that the diversified assets are associated with lower variability, and are likely to hold other beliefs that can mask the effect of the feedback. In order to evaluate this possibility, Experiment 3 examines behavior in an extreme case in which the decision makers receive a complete prior description of the payoff distributions of the different assets.

Since Experiment 3 has full description of the underlying payoff structure, it can be thought of as an extreme case of full information condition. However, in this condition, the feedback does not provide any additional information and should be ignored by rational decision makers.

A second goal of Experiment 3 was to relate the current results to the loss aversion assertion (Kahneman & Tversky, 1979). In order to achieve this goal Experiment 3 focuses on problem that involves gains and losses. The loss aversion implies a

² We do not estimate and/or test models here because we feel that such an exercise is likely to over-fit learning models on a small set of conditions. Readers interested in model comparison are encouraged to see Erev et al. (2010) that compares learning models based on 120 experimental conditions without forgone payoffs, and Erev, Ert, and Roth (2010) that compare models based on 80 experimental conditions with forgone payoffs.




Fig. 2. Condition N_5 . Choice proportions of diversified fund M . The left-hand column presents the three alternatives. The central columns present the predicted choice proportions of M under two learning models as a function of feedback and time. The right-hand column presents the results of Experiment 3.

strong tendency to prefer the safer alternative in this setting. The new problem (condition N_5) and the predictions of the three learning models are presented on the left-hand side of Fig. 2.

3.3.1. Method

The experiment involved a new group of 30 undergraduate subjects from the same subject pool. The experimental procedure was similar to the procedure in the previous two experiments with the following changes:

- (1) In the instructions (Appendix B), subjects were told the distribution of payoffs, verbally and in a table, and the correlation of the assets' returns was stressed as well as the fact that the diversified asset yields in each period a return higher than the average of the other two assets.
- (2) Payoffs could be negative, thus allowing for loss aversion considerations.
- (3) Only one condition, full information, was run. Thus, the feedback after each trial included the payoffs from all three assets.
- (4) The participants were paid for only one trial, randomly selected at the end of the study. The payoff was in agorot, where 100 agorot = 1 NIS. The final payoff was a show up fee of 30 NIS plus the payoff in this random trial. This change was introduced to insure that the difference between the early and late trials is not a result of the accumulation of wealth, gains or losses.
- (5) The practice trials were eliminated.

Other aspects remained the same, including the number of assets and the number of periods, and the administration of the instructions with time to ask questions.

3.3.2. Experiment 3 Results

The right-hand column in Fig. 2 presents the choice proportions as a function of experience in blocks of 20 periods³. The diversified option was selected in 23.6% of the trials. Recall that this asset yielded greater expected payoff and lower volatility than the other two assets. Also recall that this aspect was given and stressed to participants in unambiguous terms.

Evaluation of the effect of experience reveals a decrease in M choices with experience. The choice rate of the payoff-maximizing diversified-asset was 28% in the first block of 20 trials, and only 21% in the last block. The difference between the two rates, accounting for individual differences with a pairwise t -test (testing the individual difference between blocks is zero), is significant ($t = 2.38$, $df = 29$, $p = 0.024$).

Notice that the results reflect a reverse loss aversion bias: The observed deviation from maximization implies that people behave "as if" gains loom larger than losses⁴. We believe that this pattern suggests that in the current context the effect of feedback is larger than the effect of loss aversion. Additional indication of the limited effect of loss aversion is provided by an evaluation of the choice rate in the very first trial. The proportion of M choices in this trial is 33% (standard deviation 48%), which is 10 people out of 30. This value is as expected under random choice. Ert and Erev (2008) explain similar results with the assertion that loss aversion is better described as a heuristic that affects choice behavior in certain settings rather than as a general tendency.

3.4. Experiment 4

3.4.1. Method

Experiment 4 was designed to examine the robustness of the current results in a more realistic investment scenario. It focuses on investment decisions of upper-division (juniors and seniors) economics students and MBA business students in an online environment over time. The subjects had all taken courses in finance and statistics.

³ Blocks are used to facilitate comparison with the noisy experimental results.

⁴ Notice that this observation does not rule out the possibility that a certain proportion of the population exhibits loss aversion. It only means that the typical deviation from maximization does not reflect loss aversion.

In addition to the main study, the participants were asked to answer a short optional survey (selected questions Appendix C) that examined their investment and risk attitudes. The results (Appendix D) show that 40% of the participants have purchased stocks or mutual funds (answer to Q4), and the vast majority of respondents to Q6 (95.6%) exhibit risk aversion in portfolio choices based on a description.

Of the subjects who completed the experiment, 36 subjects participated in the limited feedback condition and 42 participated in the full feedback condition.

Participants were asked to log into an investment web site and to make investment choices among five Vanguard mutual funds: Vanguard Total Stock Market Index Fund, Vanguard Total Bond Market Index Fund, Vanguard European Stock Index Fund, Vanguard Pacific Stock Index Fund, and Vanguard Target Retirement 2015 Fund. The last fund was a composite of the other funds with the following percentages: 51.2% in Total Stock Market Fund, 36.0% in Total Bond Market Fund, 7.4% in European Stock Index Fund, 3.4% in Pacific Stock Index Fund and 2.0% in other. The composition of each of the funds was available in each fund's prospectus, which was available via a link to the Vanguard site.

To complete the experiment, participants had to make 30 separate investment decisions on 30 different days. Each day they logged in they received feedback regarding the most recent previous decision. The feedback on that decision was determined by the actual stock market performance of the different funds. Note that payoffs could be negative and loss aversion could play a role. There were two treatments: limited feedback and full feedback. The feedback common to the two treatments was the percentage return from the last decision, the cumulative return, and the number of periods remaining. In the full feedback treatment, subjects received feedback on the main screen about the returns on all five mutual funds on the last day visited, whereas in the limited feedback condition, subjects received feedback on the main screen only about the return on the mutual fund last chosen.

The mutual fund choices on the experiment's main screen all contained links to the prospectus and historical performance of the funds. To make it clear, the performance of any Vanguard fund is publically available with some search. We also provided a link on the main page next to each fund where this information was available. So the difference between the full information and limited information treatments is not in the availability of information (it was available in both), but in whether the information was presented on the main screen or required additional search by subjects (we do not have data on whether subjects searched for recent performance information on other funds).

Subjects received cash payment upon completion of the experiment for the returns they made in the experiment. To avoid scale effects (1% return could be different depending on the capital it is multiplied by), we paid subjects the sum of the percentage daily returns instead of the multiplicative returns (so returns of 1%, 3% and 2% on consecutive days would give a payoff of 6% instead of 6.1%). The exchange rate was 10 NIS per 1% gain, for an average payment of 28.6 NIS. There was no show up fee. The summation of percentage returns took away some of the realism but helped make the incentives comparable across subjects.

3.4.2. Experiment 4 Results

Over the duration of the experiment, the best performers were the stock funds. The best performing fund was the European Stock Index Fund (average daily return of 0.12%), followed by the Total Stock Market Index Fund (0.11%) and the Pacific Stock Index Fund (0.07%). The most volatile by a large difference was the Pacific Stock Index Fund (standard deviation of 0.82% compared to the second highest standard deviation of 0.66% for the European Stock Index Fund). In the limited feedback condition, the Total Stock Index and European Stock Index were most popular. In the full feedback condition, the Total Stock Index and Pacific Stock Index were most popular.

Fig. 3 presents the choice proportions as a function of experience in the two conditions. The diversified option, Vanguard Target Retirement 2015 Fund, was selected in 10% of the trials in the full feedback condition, and in 13% of the trials in the limited feedback conditions. This difference, while in the predicted direction, was not significant.

The demand for the diversified fund was low and not surprisingly since this fund was a poor performer through the duration of the experiment. Thus, we get no significant difference between treatments in the proportion of choices of this fund. While this is unfortunate, we do observe strong differences in choices of the most volatile fund and these differences are




Fig. 3. The observed choice proportions of the five assets of Experiment 4 as a function of experimental condition and time (three blocks of 10 periods).

Table 1

Pearson correlation between three of the questions in the questionnaires and the proportion of choices of the diversified assets. Bottom number denotes *p*-value.

	Q6	Q7	Q9	Choice of diversified asset
Q6 (risk aversion in hypothetical choice)	1.00	−0.421 <0.001	0.509 <0.001	0.341 0.005
Q7 self description as risk loving	−0.421 <0.001	1.00	−0.375 0.002	−0.388 0.001
Q9 tendency to exit after loss	0.509 <0.001	−0.375 0.002	1.00	0.145 0.250
Choice of diversified asset	0.341 0.005	−0.388 0.001	0.144 0.249	1.00

consistent with the adaptive explanation. Specifically, the biggest difference between the conditions occurred in the demand for the Pacific Index Fund which was the most volatile fund. In the limited feedback condition the volatile Pacific Index Fund was unpopular in the fourth place in terms of demand (16%), whereas in the full feedback condition it was the most popular (32%), particularly following strong performances. This difference is significant ($t[76] = 2.6, p = 0.01$). Thus, adaptive behavior, namely the chasing of returns, leads to higher choice proportion of the high variance choice in the case of full feedback.

While the choice of the diversified fund was not significantly different between treatments it was significantly correlated with the answers to two of the questions in the survey (see Table 1), it was not correlated with others.

The survey questions that exhibited the least correlation with portfolio choices are questions pertaining to lottery valuations. This is a result that is consistent with findings by Nasic and Weber (2010) and others. Specifically, Nasic and Weber (2010) found that neither risk perceptions in lotteries nor risk attitudes as inferred from certainty equivalences are sufficient to predict individual risk taking behavior in portfolio choice. Lack of correlation between financial portfolio decisions and risk aversion and risk attitude measures outside of the specific investment problem are generally referred to as domain specificity (Blais & Weber, 2006; Weber, Blais, & Betz, 2002) or content effect (Rettinger & Hastie, 2001).

As Table 1 shows, diversification was positively correlated with the survey-measured risk aversion in hypothetical choice (Q6), and negatively correlated with the survey-measured self description as risk loving (Q7). There was also positive but insignificant correlation of diversification with loss aversion as captured by individuals' self-described tendency to exit after a loss (Q9).

4. Summary

Past empirical evidence on individual portfolio choice highlights a tendency by some investors towards under-diversification. This tendency is inconsistent with mainstream finance theory.

The current research takes two steps toward clarifying this puzzle. First, it demonstrates the robustness of under-diversification in controlled experimental settings. Under-diversification is evident in the choice of less efficient portfolios with equal or lower return and higher risk. The pattern is robust even when under-diversification is known to reduce expected returns. In other words, in the laboratory, as in real-life, investors are not particularly adept at identifying and realizing opportunities for risk-reduction.

Second, the current analysis showed that under-diversification might be a product of a reasonable attempt to reply to feedback. Specifically, investors attempting to maximize their returns may opt to choose the alternative that yielded the highest return in the past. When they follow such a strategy, the diversified fund is less likely to be the highest yielding alternative, and so it is chosen with lower frequency over time. This pattern depends on feedback being given on all alternatives in each period. If the feedback on alternatives is more limited, the diversified fund looks much more attractive over time. In accordance with this explanation, we observed under-diversification when the feedback was complete and diversification when the feedback was limited to obtained payoffs.

The implications of the present set of observations may be intuitive, but they mount a serious challenge to economic theory as it pertains to investment settings. Learning theory has long been seen as complementary to rationality assumptions in that if economic agents are not perfectly rational, experience will take them in that direction. Here, we have shown a set of instances where experience can be detrimental to rational choice. This evidence is consistent with recent events in financial markets, where some of the most seasoned members of the financial community exhibited serious allocation errors (Becker, 2008). It appears as if experience was in many cases detrimental to prudence. While the current research offers no solutions to financial mismanagement by top executives, it does offer simple solutions for financial mismanagement by individual investors, namely the organization of feedback on returns in a way that is less conducive to chasing of returns.

Appendix A. Normalized fictitious play (source: Ert and Erev, 2007)

The model can be summarized by two assumptions: (1) the adjusted propensity to select alternative *j* at time *t* + 1 is:

$$q_{j,t+1} = (1 - w)q_{jt} + wx_{jt}, \quad (\text{A.1})$$

where x_{jt} is the payoff of choice j in trial t , and w is a parameter that determines the weight of this payoff. (2) The probability that the decision maker will select strategy j in trial t is:

$$p_{jt} = \frac{\exp(q_{jt}\lambda/S_t)}{\sum_{k=1}^n \exp(q_{kt}\lambda/S_t)} \tag{A.2}$$

where λ is a parameter that determines the relative importance of the different propensities, and S_t is a normalization factor. Under full feedback, S_t is the experienced regret level of the decision maker:

$$S_{t+1} = (1 - w)S_t + w|\max_t - x_{jt}| \tag{A.3}$$

where \max_t is the maximal payoff obtained in trial t over the k alternatives. Under limited feedback, S_t captures the difference between the current payoff (x_t) and the last payoff (x_{t-1}):

$$S_{t+1} = (1 - w)S_t + w|x_{t-1} - x_t| \tag{A.4}$$

The initial level is set to equal $S_1 = \lambda$ (and the first adaptation under limited feedback occurs in trial $t = 2$). When the feedback is limited to the obtained payoffs, the decision makers sample each of the k alternatives in the first k trials.

Appendix B. Experiment instructions (translated)

B.1. Experiments 1–2 instructions sheet (translated from Hebrew)

Welcome to an experiment in decision making. The experiment includes 100 rounds. In each round you will get 100 tokens and will be asked to invest all your tokens in one of three assets. The asset returns come from a distribution that is unknown to you. The profit in each round is calculated as follows:

$$\text{Profit} = \text{Rate of return on the chosen asset} \times 100$$

Hypothetical example: You have 100 tokens to invest in one of the assets A, B or C. Please select one of the assets. After choosing the asset, you will see the following feedback:

Asset	Profit in NIS for this round	Profit for 100 token investment	Rate of return	Investment
A				
B				
C				

After receiving the feedback you will begin a new round with 100 tokens. You are not able to reinvest your profit. In each round you will know your cumulative profit. Payment for participation in the experiment: 200 tokens = 1 NIS.

B.2. Experiment 3 instructions sheet (translated from Hebrew)

Welcome to an experiment in decision making. In this experiment, you will make decisions in 100 rounds. In each round, you will invest in one of three assets, A, B, and C.

In each round, the computer will randomly sample X and Y as follows: X comes from a normal distribution with a mean of 0 and standard deviation of 300. Y comes from a normal distribution with a mean of 0 and standard deviation of 10. The payoff will be rounded up to the nearest integer above it. The payoff on each asset will be determined as follows:

Asset	Payoff in agorot
A	$4X$
B	$2Y - 2X$
C	$X + Y + 5$

B.2.1. Attention!!!

The payoff you obtain from asset C is the average of the payoffs of assets A and B, plus an additional 0.05 NIS. To see this, note that $\frac{A+B}{2} = \frac{4X+2Y-2X}{2} + 5 = X + Y + 5$.

Hypothetical example:

Step A. Select an asset. Please press the button corresponding to the selected asset.

Step B. Summary of results from the current round.

Asset	Payoff
A	
B	
C	

Earnings are not cumulative and you may not reinvest earnings from previous rounds.

Payment from the experiment:

One round will be selected in random for payment and you will receive your earnings from that round plus 30 NIS show up fee.

Appendix C. Selected questions from survey administered to participants

Q4. Have you ever invested in stocks or mutual funds? Yes/No.


Q6. The table below demonstrates a trade off between average return and the likelihood of losing money in any 1 year. Please select the portfolio that, for you, best balances these tradeoffs between risk and return.

	Potential average return (%)	Chance of losing \$ in any 1 year	Worst year of 70 years (%)	Worst year of 20 years (%)
Portfolio A	17.6	1 in 3	-58	-33
Portfolio B	13.3	1 in 4	-43	-18
Portfolio C	11.7	1 in 5	-27	-9
Portfolio D	8.5	1 in 6	-19	-5
Portfolio E	6.2	1 in 10	-6	-1
Portfolio F	5.0	1 in 50	-2	0

Q7. Which of the following best describes your attitude towards balancing a desire for returns relative with the risk you feel you can tolerate?

- (a) My primary goal is preservation of principal and risk avoidance. I will accept lower returns in an effort to avoid investment risk.
- (b) I want to avoid risk, but will accept a relatively small amount to achieve a slightly higher return.
- (c) I can tolerate a moderate amount of risk in an effort to achieve a moderate amount of growth.
- (d) I want to achieve potentially high returns, and I am willing to accept the high amount of risk associated with this goal.

Q9. The following graph shows the potential range of results of six portfolios in any 1 year. The best potential (top 5% return) and worst potential (worst 5%) returns are represented. Please note that the highest potential returns also have the greatest potential losses. Which of these portfolios would you prefer to hold?



Appendix D. Survey raw data for selected questions

ID	1	3	4	5	7	8	9	10	11	12	13	15	17	18	19	20	21
Q4	1	2	2	1	2	2	2	2	1	2	2	1	2	1	1	1	1
Q6	1	4	4	3	3	4	4	4	5	4	4	3	3	5	1	2	3
Q7	2	3	3	3	3	3	2	2	2	2	2	3	2	3	3	3	3
Q9	2	4	4	3	3	4	4	4	5	4	4	4	3	3	3	2	5
ID	23	24	25	26	27	28	29	31	32	33	34	35	36	37	38	39	42
Q4	2	2	2	2	2	2	2	1	1	1	1	2	2	1	2	2	1
Q6	3	3	2	3	5	4	5	4	3	5	6	3	6	6	3	6	5
Q7	3	2	4	2	1	2	2	2	3	2	3	2	2	3	3	2	2
Q9	1	5	3	3	6	4	5	4	2	4	4	4	5	4	2	5	5
ID	43	44	45	46	47	48	49	50	53	54	55	56	57	58	59	60	61
Q4	1	1	1	2	2	2	2	1	2	1	1	2	2	2	2	2	1
Q6	4	3	6	3	3	2	4	5	6	5	4	4	3	5	4	3	3
Q7	2	3	2	2	3	3	2	3	2	2	3	2	3	2	2	3	4
Q9	4	5	4	2	4	1	5	4	4	4	2	1	3	5	4	5	4
ID	62	64	65	66	67	69	70	71	73	74	75	76	77	78	79	80	81
Q4	2	1	1	2	1	1	1	2	2	2	2	2	1	2	2	2	2
Q6	3	3	2	4	5	6	5	6	1	3	3	6	3	6	5	3	2
Q7	2	3	3	2	2	2	3	2	2	3	3	2	3	2	2	3	3
Q9	3	2	2	4	5	3	2	4	3	3	4	5	4	5	4	4	2

References

- Becker, G. (2008). *Risk-taking by top executives*. The Becker-Posner Blog.
- Blais, Ann-Rene, & Weber, Elke (2006). A domain-specific risk taking (dosPERT) scale for adult populations. *Judgment and Decision Making*, 1(1), 33–47.
- Blume, M. E., & Friend, I. (1975). The asset structure of individual portfolios and some implications for utility functions. *Journal of Finance*, 30, 585–603.
- Brown, G. W. (1951). Iterative solutions of games by fictitious play. In T. C. Koopmans (Ed.), *Activity analysis of production and allocation*. New York: Wiley.
- Chevalier, Judith, & Ellison, Glen (1997). Risk taking by mutual funds as a response to incentives. *Journal of Political Economy*, 105, 1167–1200.
- Choi, James, Laibson, David, & Madrian, Brigitte (2010). "Why does the law of one price fail? An experiment on index mutual funds". *Review of Financial Studies*, 23(4), 1405–1432.
- Choi, J., Laibson, D., Madrian, B., & Metrick, A. (2009). Reinforcement learning and savings behavior. *Journal of Finance*, 64(6), 2515–2534.
- Chopra, N., Lakonishok, J., & Ritter, J. R. (1992). Measuring abnormal performance – Do stocks overreact? *Journal of Financial Economics*, 31, 235–268.
- De Bondt, W. F. M., & Thaler, R. (1985). Does the stock market overreact? *Journal of Finance*, 40, 793–805.
- De Bondt, W. F. M., & Thaler, R. (1990). Do security analysts overreact? *American Economic Review*, 80, 52–57.
- Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review*, 114(1), 177–187.
- Denrell, J., & March, J. G. (2001). Adaptation as feedback restriction: The hot stove effect. *Organization Science*, 14, 227–243.
- Erev, I., & Haruvy, E. (2010). *Learning and the economics of small decisions*. Working paper.
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S., Hau, R., et al (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23, 15–47.
- Erev, I., Ert, E., & Roth, A. E. (2010b). *A choice prediction competition for market entry games: An introduction*. Working paper. Available from: <http://sites.google.com/site/gpredcomp/>.
- Ert, E., & Erev, I. (2007). Replicated alternatives and the role of confusion, chasing, and regret in decisions from experience. *Journal of Behavioral Decision Making*, 20, 305–322.
- Ert, E., & Erev, I. (2008). The rejection of attractive gambles and the heuristic interpretation of loss aversion. *Journal of Economic psychology*, 29, 715–723.
- Fudenberg, D., & Levine, D. (1998). *Learning in games*. Cambridge, Mass: MIT Press.
- Goetzmann, W. N., & Kumar, A. (2004). *Why do individual investors hold under-diversified portfolios?* Working paper. Yale University.
- Grosskopf, B., Erev, I., & Yechiam, E. (2006). Foregone with the wind: Indirect payoff information and its implications for choice. *International Journal of Game Theory*, 34(2), 285–302.
- Hendricks, Darryll, Patel, Jayendu, & Zeckhauser, Richard (1993). Hot hands in mutual funds: Short-run persistence in relative performance. *Journal of Finance*, 48, 93–130.
- Ippolito, Richard A. (1992). Consumer reaction to measures of poor quality: Evidence from the mutual fund industry. *Journal of Law and Economics*, 35, 45–70.
- Jacobs, Heiko, Muller, Sebastian, & Weber, Martin (2009). *How should private investors diversify? – An empirical evaluation of alternative asset allocation policies to construct a "world market portfolio"*. Working paper.
- Kahneman, Daniel, & Tversky, Amos (1979). Prospect theory: Analysis of decision under risk. *Econometrica*, 47, 263–291.
- Kelly, M. (1995). All their eggs in one basket: Portfolio diversification of US households. *Journal of Economic Behavior and Organization*, 27, 87–96.
- Marchiori, Davide, & Warglien, Massimo (2008). Predicting human interactive learning by regret-driven neural networks. *Science*, 319, 1111–1113.
- Nosic, Alen, & Martin Weber (2009). *Overreaction and asset allocation: An experimental analysis*. Working paper.
- Nosic, Alen, & Weber, Martin (2010). How risky do I invest: The role of risk attitudes, risk perceptions and overconfidence. *Decision Analysis*, 7(3).
- Odean, T. (1999). Do investors trade too much? *American Economic Review*, 89, 1279–1298.
- Offerman, Theo, & Sonnemans, Joep (2004). What's causing overreaction? An experimental investigation of recency and the hot hand effect. *Scandinavian Journal of Economics*, 106, 533–553.
- Polkovnichenko, V. (2005). Household portfolio diversification: A case for rank-dependent preferences. *Review of Financial Studies*, 18, 1467–1502.
- Rettinger, David A., & Hastie, Reid (2001). Content effects on decision-making. *Organizational Behavior and Human Decision Processes*, 85(2), 336–359.
- Sirri, Erik R., & Tufano, Peter (1998). Costly search and mutual fund flows. *Journal of Finance*, 53, 1589–1622.
- Statman, M. (1987). How many stocks make a diversified portfolio? *Journal of Financial and Quantitative Analysis*, 22, 353–363.
- Statman, M. (2004). The diversification puzzle. *Financial Analysts Journal*, 60, 44–53.

- Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychological Review Monograph* (Suppl. 2).
- Uppal, R., & Wang, T. (2003). Model misspecification and under-diversification. *Journal of Finance*, 58, 2465–2486.
- Weber, Elke, Blais, Ann-Rene, & Betz, Nancy E. (2002). A domain specific risk attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, 15(4), 263–290.