Risk aversion, Under-diversification, and the Role of Recency and Probability Matching

Tal Shavit[a], Uri Ben Zion[b], Ido Erev[c], Ernan Haruvy[d]

August, 2005

Abstract

Two experiments are presented that compare alternative explanations to the coexistence of risk aversion and under-diversification (the tendency to focus on few stocks) in investment decisions. The participants were asked to select one of three assets under two feedback conditions. The returns from two assets were negatively correlated, and the third asset was equally weighted combination of the other two, allowing for lower volatility. The popularity of the weighted asset was highly sensitive to feedback condition. It was the most popular asset when the feedback was limited to the obtained payoff, but the least popular asset when the feedback included information concerning forgone payoffs. This pattern supports the assertion that under-diversification can be a product of momentum trading and/or probability matching.

JEL Classification Numbers: *C99, D81; D83*

Key words: Risk, Diversification

[a] *Department of Economics, Ben-Gurion University, Beer-Sheva 84105, Israel*
[b] *Department of Economics, Ben-Gurion University, Beer-Sheva 84105, Israel*
[c] *Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel*
[d] *School of Management, University of Texas at Dallas.*

Corresponding author: Tal Shavit, Hagiborim 84/18 Hadera, 38380, Israel.
Mail: shavittal@gmail.com
Phone: 972-52-2920868.
Fax: 972-3-5706433.

1. Introduction

Standard portfolio theory (e.g., Markowitz, 1952) suggests that investors facing two portfolios with equal expected returns will prefer the portfolio with the lower volatility. This assertion is consistent with common applications of expected utility theory (e.g., Tobin 1958, Sharpe 1966) and with experimental studies of choice behavior and portfolio allocation over safe and risky assets (e.g., Kahneman & Tversky, 1979; Thaler et al., 1997; Gneezy & Potter, 1997).

However, there are exceptions to this predictable and robust pattern. One important exception is the empirical observation that many individual investors hold fewer individual stocks than necessary to eliminate idiosyncratic risk (e.g., Blume and Friend, 1975; Statman, 1987; Kelly, 1995; Odean, 1999; Polkovnichenko, 2004; Goetzmann and Kumar, 2004). Statman (1987), Meulbroek (2002), and Goetzman and Kumar (2004) show empirically that failure to diversify is costly. This under-diversification contradicts standard portfolio theory in cases where diversification can reduce volatility without reducing expected returns.[1] The main goal of the current paper is to compare three explanations to this seemingly risk-seeking behavior and its inconsistency with the commonly accepted risk-averse behavior of investors.

The first explanation, referred to as the "*contingent risk attitude hypothesis*", suggests that the empirical results reflect large differences between different classes of investors. Some investors prefer individual stocks and have little interest in reducing their risk by selecting index funds or using diversification. Thus, the empirical exception described above reflects the risk tendencies of a small minority that is over-represented in analysis of the behavior of private stock market investors.

Under a second explanation, referred to as the "*contingent recency hypothesis*," under-diversification is a result of a tendency to rely on recent outcomes. The contingent recency hypothesis does not specify the direction of the recency effect; it allows for the possibility of large between investors differences. Whereas some investors exhibit positive recency (see Barron & Erev, 2003), other investors may exhibit negative recency (Kroll, Levy & Rapoport, 1988). Positive recency implies momentum trading behavior (e.g., Grinblatt et al., 1995). Under momentum

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

investment strategies, investors buy recent winning stocks and sell recent losing stocks. Negative recency implies contrarian trading. This pattern is also known as the gambler fallacy (see e.g., Rabin, 2002). The contingent recency hypothesis captures under-diversification because both positive and negative recency imply a bias toward the high variability (under-diversified) options.

Finally, under the "*probability matching hypothesis*" (see Estes, 1950) investment decision deviate from expected utility maximization in the direction of probability matching. In certain cases investors match the probability of selecting each option to the proportion of times in which this option has led to the best outcome. Like the contingent recency hypothesis, this hypothesis captures under-diversification as a product of the fact that diversified options are rarely associated with the highest outcomes.

Notice that the current hypotheses cannot be easily compared based on evaluation of sequential dependencies in the stock market. Neither hypothesis has clear predictions concerning sequential dependencies. For example, the positive and negative recency effects predicted by the contingent recency hypothesis can cancel each other at the market level. Focusing on individual investors cannot solve this problem, because the contingent recency hypothesis allows for the possibility that the exact recency effect exhibited by a particular investor will change between decisions. In order to address this difficulty the current analysis relies on laboratory experiments.

2. Experiments

Two laboratory experiments were run in order to evaluate the three hypotheses. In each trial of each experiment, participants were asked to select one of three assets. Two assets, denoted assets A and B, represent stocks with the same expected return and with high negative correlation between their returns (-0.977). The third asset represents a diversified fund. The focus on this simplified choice task was designed to minimize complexity associated with a diversification strategy. Though it is possible that investors avoid diversification due to the complexity associated with holding a large portfolio, the availability of diversified funds, particularly index funds, makes this complexity argument less relevant.

In Experiment 1, the return from the diversified fund, referred to as asset M, is the mean of the returns from A and B. Thus, it had the same expected return and much lower dispersion. Two information conditions were compared—'Full Information' and 'Limited Information'. After each trial in the Full Information Condition the investors could observe the return from all three assets. In the Limited Information Condition, the information is restricted to the return from the selected asset. Experiment 2 replicates Experiment 1 with one exception: The return from the safe asset, referred to as M+, was equal to the mean of the other assets plus 2%.

Concrete predictions

According to the contingent risk attitude hypothesis, the coexistence of risk aversion and under-diversification in the stock market is due to the fact that risk seeking decisions are overrepresented in the stock market through self-selection of risk-seekers in the stock market. This representation problem is avoided in the current experiment. Thus, if risk aversion is prevalent (as implied by the contingent risk attitude hypothesis) the diversified option will be most popular in both experimental conditions.

The remaining hypotheses, contingent recency and probability matching, predict under-diversification (less than 33% diversified choices) in the full information condition and higher diversification rate in the limited information condition. These hypotheses differ with respect to their predictions of the relationship between the observed sequential dependencies and diversification. We discuss these predictions below in the context of specific quantifications of the contingent recency and probability matching hypotheses.

2.1. Experiment 1

Participants were 40 undergraduate subjects who had taken at least one course in statistics. The experiments took place at a computer Laboratory at Ben-Gurion University, and lasted approximately half an hour. The average payment was 25 NIS, or approximately $5.5. The experiment focused on a simplified investment task. In each of 100 trials, participants were asked to choose one of three assets: A, B and M.

The participants were informed that they would be asked to invest 100 experimental tokens in one of the assets in each trial (1 NIS = 200 tokens). They were also told that their profit from previous trials could not be used for reinvestment in the assets. Ten practice trials were conducted before the 100 experimental trials.

The instructions (Appendix 1) spell out, in a simple and non-technical way, the experimental procedure with no information 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 information conditions—Full Information and Limited Information-- of 20 subjects each. The groups differed with respect to the feedback provided after each trial. After each trial in the Full Information condition, participants observed their earning from all assets. In the limited information condition, the feedback was limited to the payoff obtained from the selected asset only.

The assets' returns are constructed on the basis of two variables, U and ε that were independently drawn from uniform distributions in each trial in the ranges of 0% to 100%, and -5% to 5% respectively. The distributions' equations and implied means and standard deviations are summarized in Figure 1. The implied correlation between assets A and B is –0.977.

<Insert Figure 1>

2.1.1 Results

The left-hand column in Figure 1 presents the choice proportions as a function of experience in the two conditions. The diversified option was selected if 14% of the trials in the full Information condition, and in 39% of the trials in the limited information conditions. This difference, predicted under the contingent recency and probability matching hypothesis, is significant ($t[38] = 4.9$, $p < 0.01$).

2.1.1 Alternative models

In order to evaluate the contingent recency and probability matching hypotheses, we considered three simple models that quantify the basic ideas behind these assertions:

A Weighted adjustment model. The first model is a weighted adjustment quantification of the recency hypothesis. This model assumes a stochastic choice rule and that beliefs are updated with depreciation of past beliefs by a constant. It has two parameters: λ -- payoff sensitivity parameter and β-- the belief adjustment parameter. The equations are

$A_{njt} = (1-\beta)A_{njt-1} + \beta X_{njt-1}$, if the payoff to action j was observed in period t-1

$A_{njt} = A_{njt-1}$ otherwise (1)

$$P_{njt} = \frac{\exp(\lambda A_{njt})}{\sum_{j \in J} \exp(\lambda A_{njt})}. \quad (2)$$

The two parameters of the model were estimated by running computer simulation in which 200 virtual agents who behave according to the model's assumption face decisions and receive feedback corresponding to each experimental condition. The simulations were run with a wide set of parameters, and the estimated parameters are the parameters that minimized the mean squared distance between the observed and simulated curves. The estimated parameters are 0.008 and 3.50 for beta and lambda respectively and the MSD score is 0.013.

The aggregate predictions of the weighted adjustment model are presented in the second column of Figure 1. The results reveal that the model captures then main aggregated trends, but it over-predicts the proportion of diversified (M) choices.

A probability matching model. According to the probability matching hypothesis the probability of selecting a certain option is matched to the proportion of trials in which this option has yielded the best outcome. Following Erev and Barron (in press) the current abstraction of this hypothesis, assumes best reply to history.

When full information is available, the decision maker is assumed (at trial t>1) to randomly recall one of the (t-1) previous trials, and select the alternative that yield the best outcome in that trial. Random choice is assumed in the first trial.

In the limited information condition, the decision maker is assumed to choose randomly, until observing at least one realization from each option. At later trials, the decision maker is assumed to randomly recall one of the obtained realizations from each option, and select the option with the highest recalled value.

The predictions of this parameter-free model are presented in the third column of Figure 2. The results reveal that the probability-matching model captures the information effect but under-predicts the proportion of diversified (M) choices. The MSD score is 0.025.

A combination of weighted adjustment and probability matching.

Erev and Barron (in press) show that the main behavioral regularities observed in experimental studies of decisions from experience can be captured with models that assume a joint effect of recency and probability matching. Based on this observation the third model considered here assumes that each decision maker follows one of the two rules listed above. To avoid additional free parameters, it assumes that half the participants follow the weighted adjustment rule (with the previously estimated parameters), and half the participants follow the probability-matching rule.[2] The fourth column in Figure 2 shows this post hoc model. Its MSD score is 0.016.

2.1.2 Individual differences

In order to evaluate the role of individual differences we computed a recency index for each participant in the full information condition. The index was defined as the proportion of times that the option with the highest recent payoff was selected conditional on the fact that Option M was not selected. Figure 2 presents the proportion of underdiersification (selection of A or B) as a function of the recency index for each participant. The results show that most participants (18 of 20) exhibit positive recency (the probability of selecting the option with higher recent payoff is higher than 50%. In addition, the data reveals large between participant differences. The correlation between the recency index and under- diversification is positive (r=.32) but insignificant.

The large between-participant variability implies that the models presented above are oversimplified. In particular, the assumption that all the participants (that follow the weighted adjustment rule) can be described with the same two parameters is too strong.

2.2. Experiment 2

Experiment 2 was conducted to examine the robustness of the results of experiment 1in settings where the less volatile asset has a higher expected return.

[2] The model supported in Erev and Barron, referred to as reinforcement learning among cognitive strategies (RELACS) assumes that probability matching and weighted adjustment are among the cognitive strategies considered by the decision makers. The predictions of this model for the current setting are similar (and closer to the data) than the prediction of the simpler model presented here. RELACS outperform the current model even with the parameters estimated by Erev & Barron (2005).

Under the recency and probability matching hypotheses, information concerning forgone payoff is expected to impair earning in this setting.

The design of experiment 2 closely follows the design of experiment 1. That is, 40 undergraduate students with training in statistics are divided into two information conditions of 20 participants each. Each group chooses one of three assets repeatedly 100 times and receives feedback either for the asset chosen (Limited Information Condition) or for all assets (Full Information Condition). We replace asset M of experiment 1 with asset M+, which has 2% higher return. The return distribution is shown in the top panel in Figure 3.

2.2.1. Results.

The left-hand column in Figure 3 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 Information condition, and in 43.4% of the trials in the limited information conditions. This difference, that reflects a negative effect of information on earnings as predicted under the contingent recency and probability matching hypotheses, is significant ($t[38] = 4.9$, $p < 0.01$).

The right hand columns of Figure 3 present the predictions of the learning models to Experiment 2. The predictions were derived with the parameters estimated in Experiment 1. The results show that the combined model outperforms the two basic models. The MSD score of the combined model is 0.016. The MSD scores of the basic models are 0.027 and 0.025, for the weighted adjustment model and the probability matching model respectively.

3. Summary

Previous studies of investment decisions highlight two robust but apparently inconsistent behavioral tendencies: Investors tend to exhibit strong risk aversion in demanding higher returns for risky assets, but they also tend to prefer under-diversified portfolios, a behavior consistent with risk loving behavior. The current work compares alternative explanations to this puzzle. The results demonstrate under-diversification that can be described a byproduct of a contingent recency effect, and of probability matching. Comparison of alternative quantifications of these behavioral tendencies shows the advantage of models that assume the co-existence of the two tendencies.

It is important to emphasize that the current research does not prove that under-diversification in the stock market is driven by a contingent recency effect and/or probability matching. The main contribution of the current research is the demonstration that under-diversification can be a product of these robust psychological principles.

Among the interesting implications of this demonstration is the possibility that the contingent recency and probability matching can effect diversification even when it is impossible to detect evidence for a recency effects in real stock market data. Thus, these factors may be more important than suggested by traditional analyses that focus of sequential trends in the stock markets.

Another interesting implication of the current results involves the role of information. Previous experimental research (Thaler et al., 1997; Gneezy & Potter, 1997) suggests that complete information enhances myopic loss aversion that can lead investors to prefer low risk assets even if they are associated with low expected return. The current research suggests that complete information can trigger the opposite pattern: A bias toward high variability and low return options. We believe that this interesting difference suggests that future research should consider the joint effect of loss aversion and the recency/probability matching effects studied here.

4. References

Barber, Brad M. and Terrance Odean (2000), Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors, *The Journal of Finance*, 55, 773-806.

Barron. G. and I. Erev (2003), Small Feedback Decisions and their Limited Correspondence to Description-Based Decisions, *Journal of Behavioral Decision Making*, 16(3), 215-233.

Benzion, U., E. Haruvy, and T. Shavit (2004), Adaptive Portfolio Allocation with Options, *Journal of Behavioral Finance*, 5(1), 43-56

Black, Pam (1994), "Index *Funds*: A Safer Seat for the Long-Distance Rider," Business Week, (25,April), 138-139.

Blume, Marshall E. and Irwin Friend, 1975, The Asset structure of Individual Portfolios and Some Implications for Utility Functions, *Journal of Finance*, 30, 585-603.

Bogle, John C. (1996), "Six Things to Remember About Indexing, and One Not to Forget," Speech Presented at the 1996 AIMR Annual Conference in Atlanta, Georgia (8,May).

Campbell, J., M. Lettau, M. Burton and Y. Xu (2001), "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk," *Journal of Finance*, Vol. 56 (2001), pp. 1-46

Carhart, Mark M. (1995), Survivorship Bias and Mutual Fund Performance, Ph.D. Thesis, Graduate School of Business, University of Chicago.

Carhart, Mark M. (1997), "On Persistence in Mutual Fund Performance," Journal of Finance, 52, 57-82.

Chevalier, Judith and Glenn Ellison (1999), Are Some Mutual Fund Managers Better Than Others? Cross-Sectional Patterns in Behavior and Performance, Journal of Finance, 54, 875-900.

Cohn, Richard A., Wilbur G. Lewellen, Ronald C. Lease, and Gary G. Schlarbaum (1975), Individual Investor Risk Aversion and Investment Portfolio Composition, The Journal of Finance, 30, 605-620.

Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers (1997), Measuring mutual fund performance with characteristic based benchmarks, Journal of Finance, 52, 1035-1058.

De-Long, J.B., A. Shleifer, L. H. Summers and R. J. Waldmann, 1990, Positive Feedback Investment Strategies and Destabilizing Rational Speculation, Journal of Finance, 45(2), 379-395.

Erev, I. and Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies, forthcoming in *Psychological Review*.

Goetzmann, William N. and Alok Kumar, 2004, Why Do Individual Investors Hold Under-diversified Portfolios? Working paper, Yale University.

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

Grinblatt, M., Titman,S., Wermers, R., 1995. Momentum investment strategies portfolio performance, and herding: a study of mutual fund behavior. American Economic Review 85, 1088-1105.

Hong, Harrison, Terence Lim and Jeremy C. Stein (2000), Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies, *The Journal of Finance*, Vol. 55, 265-295.

Hong, Harrison and Jeremy C. Stein (1999), A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets, *The Journal of Finance*, Vol. 54, 2143-2184.

Jagadeesh, Narasimhan and Sheridan Titman (1993), Returns to Buying Winners and Selling Losers; Implications for Stock Market Efficiency, *The Journal of Finance*, 48, 65-91.

Kelly, Morgan, 1995, All their eggs in one basket: Portfolio Diversification of U.S. Households, *Journal of Economic Behavior and Organization*, 27, 87-96.

Kroll, Y., Levy, H., & Rapoport, A. (1988) Experimental tests of the separation theorem and the capital asset pricing model, *American Economic Review*, 78 (3) 500-519

Landerman, Jeffrey M. (1996), "The Stampede to Index *Funds*," Business Week, (1,April), 78-79.

Markowitz, H., Portfolio Selection, *Journal of Finance*, 1952.

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

Malkiel, Burton, 1995, Returns from investing in mutual funds 1971-1991, Journal of Finance, 50, 549-572.

Meulbroeck, Lisa, 2002, Company Stock in Pension Plans: How Costly Is It? Working Paper, Harvard Business School working paper.

Odean,T., 1998. Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775-1798.

Odean, Terrance, 1999, Do Investors Trade Too Much? *American Economic Review*, 89, 1279-1298.

Polkovnichenko, Valery, 2004, Household Portfolio Diversification: A Case for Rank-Dependent Preferences, forthcoming in *Review of Financial Studies*.

Rabin, Matthew (2002), Inference by Believers in the Law of Small Numbers, Quarterly Journal of Economics, 117, 775-816

Rouwenhorst, K. Geert (1998), International momentum strategies, *Journal of Finance*, 53, 267-284.

Sharpe, William F., "Mutual Fund Performance," *Journal of Business*, January 1966, pp. 119-138.

Smith, Anne Kates, Steven D. Kaye, Kerry Hannon, and Jack Egan (1996), "Going for the Gold," *U.S. News and World Report*, (8, July): 56-65.

Statman, Meir (2004). How much diversification is enough, *Financial Analysts Journal*, forthcoming

Tang, Gordon (2004), How efficient is naive portfolio diversification? an educational note. Omega (Oxford), 32(2), 155-160

Tobin, J, (1958), Liquidity preference as behavior towards risk, *Review of Economic Studies* 25, 65-86.

Uppal, Raman and Tan Wang (2003), “Model Misspecification and Underdiversification,” *Journal of Finance*, 58(6): 2465-2486.

Updegrave, Walter (1995), "Why *Funds* Don't Do Better," *Money,* (August): 58-67.

Wang, Yong (1993) Near-Rational Behaviour and Financial Market Fluctuations, *Economic Journal*, 103, 1462-1478.

Welch, Ivo (1992), Sequential Sales, Learning, and Cascades, *Journal of Finance*, 47, 695-732.

Figure 1: The top panel presents the three assets examined in Experiment 1. The left –hand column in the low panel presents the observed choice proportions of the three options as a function of experimental condition and time (5 blocks of 20 periods). The right-hand columns preset the predictions of three models.

Asset	Asset return	Mean return	Standard deviation.
A	$R_A = U$	50%	29.3%
B	$R_B = (75\%-0.5*U)+ \varepsilon$	50%	14.92%
M	$R_M = 0.5* R_A +0.5* R_B$	50%	7.46%

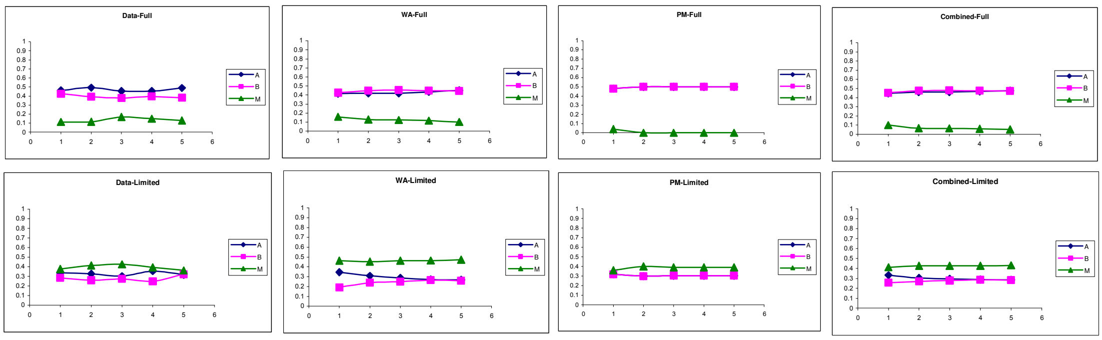

Figure 2: The proportion of under-diversification as a function of the recency index. Each data point presents one of the 20 participants in Experiment 1.

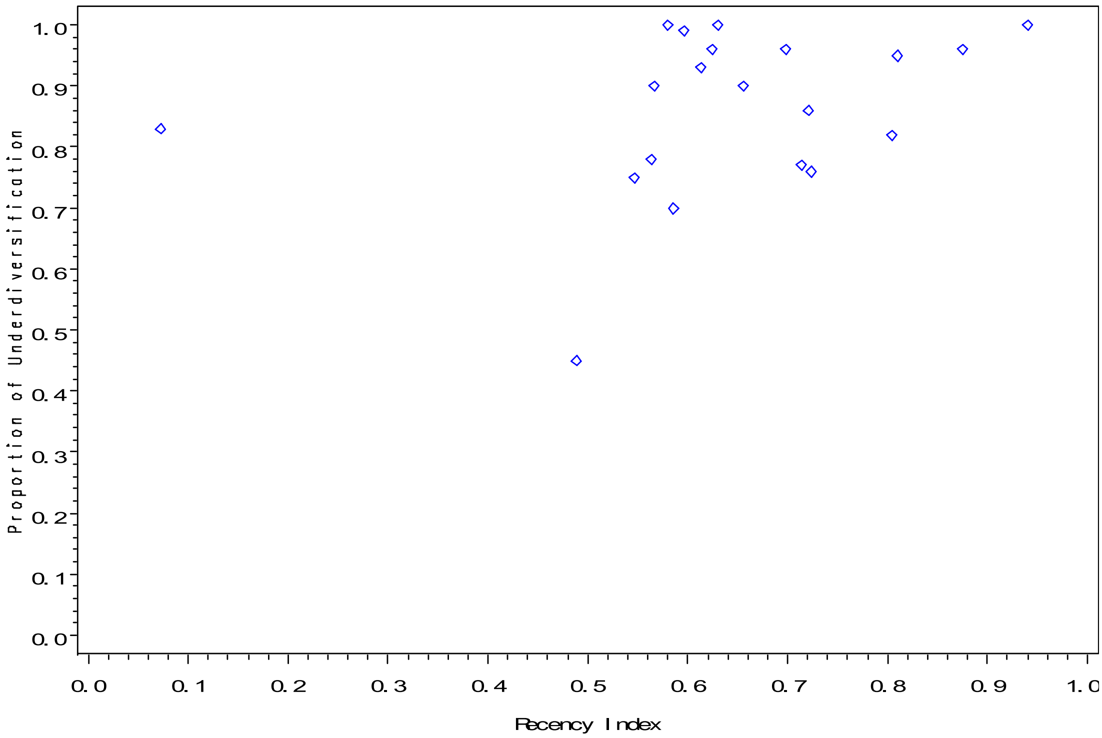

Figure 3: The top panel presents the three assets examined in Experiment 2. The left column in the low panel presents the observed choice proportions of the three options as a function of experimental condition and time (5 blocks of 20 periods). The right hand columns preset the predictions of three models.

Asset	Asset return	Mean return	Standard deviation.
A	$R_A = U$	50%	29.3%
B	$R_B = (75\%-0.5{*}U)+ \varepsilon$	50%	14.92%
M+	$R_M = 0.5{*} R_A +0.5{*} R_B + 2\%$	52%	7.46%

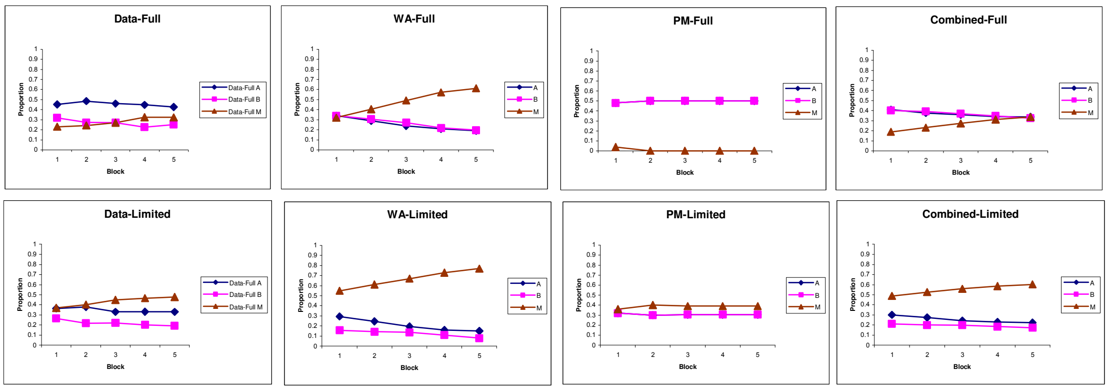

Appendix 1 – Experiment Instructions (translated).

- *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: Profit = Rate of return on the chosen asset *100*

Hypothetical example:

You have 100 tokens to invest in one of the assets A, B or C. Select one of the assets:

A	**B**	**C**

Stage 2: the feedback.

- *After choosing the asset, you will see the following information:*

ASSET	*Investment*	*Rate of Return*	*Profit for 100 token investment*	*Profit In N.I.S for this round*
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 N.I.S.

The payment is cumulative.