

Fund Choice Behavior and Estimation of Switching Models: An Experiment*

Mikhail Anufriev[†] Te Bao[‡] Jan Tuinstra[§]

March 6, 2013

Abstract

We run a laboratory experiment that contributes to the finance literature on “return chasing behavior” studying how investors switch between mutual funds driven by past performance of the funds. The subjects in this experiment make discrete choices between several (2, 3 or 4) experimental funds in multiple periods. The time series of funds’ profits are exogenously generated prior to the experiment and subjects are paid for that period according to the profit of the fund they choose. The experimental results show that the investment decision can to a large extent be explained by a discrete choice model (“switching model”) with a few lags and a predisposition effect. The intensity of choice parameter β in the discrete choice model depends on the structure of the profit time series of the funds, and there is no evidence that it is influenced by experience.

*We benefitted from helpful comments by Simone Alfarano, Cees Diks, John Duffy, Cars Hommes and participants of the 5th POLHIA Workshop 2011 at Paris, France, the “Computing in Economics and Finance” conference 2011 at San Francisco, USA, the International Meeting on Experimental and Behavioral Economics 2012 at Castellon, Spain and seminars at the Tinbergen Institute and the University of Amsterdam. An earlier version of this paper was circulated under the title “Switching Behavior in the Lab”. We are grateful to the financial support from the EU 7th framework collaborative project “Monetary, Fiscal and Structural Policies with Heterogeneous Agents (POLHIA)” Grant No. 225048 and the Research Priority Area Behavioral Economics of the University of Amsterdam.

[†]Economics Discipline Group, University of Technology, Sydney, Australia. E-mail: Mikhail.Anufriev@uts.edu.au

[‡]CeNDEF, Amsterdam School of Economics, University of Amsterdam, Amsterdam, The Netherlands. E-mail: T.Bao@uva.nl

[§]CeNDEF, Amsterdam School of Economics, University of Amsterdam, Amsterdam, The Netherlands. E-mail: J.Tuinstra@uva.nl

Keywords: Switching, discrete choice model, mutual fund, experimental economics, intensity of choice.

JEL Classification: C25, C91, D83, G11.

1 Introduction

Switching between profitable alternatives is a common phenomenon in many aspects of life, ranging from decisions about which route to take when commuting to the workplace to investment in financial markets. In the financial literature, there are many studies on “return chasing behavior” that investigate how investors switch between mutual funds when they are primarily driven by past performance of these funds. In fact, it is very common for professional financial websites (e.g. Morningstar, Yahoo Finance) to provide information on past performance of all mutual funds, and recommend or rate funds based upon this past performance. Empirical evidence suggests that the money that flows in and out of these mutual funds is indeed strongly driven by recent past performance of these funds (see e.g. Patel et al. (1994), Sirri and Tufano (1998)), although the funds that performed well in the past and thereby attracted more money do not necessarily generate a higher average return than other funds (Jain and Wu (2000), Zheng (1999)).

In this paper, we study to what extent fund choice decisions are driven by past performance by means of a laboratory experiment. The advantage of running a laboratory experiment is that it provides a controlled environment so that we can rule out other factors that may also influence investment decisions (e.g. service quality). Every participant chooses a fund and is paid for that period according to the profit of the chosen fund.¹ We understand that investors may care not only about the values of the past returns, but also about the structure of the time series (e.g. whether there is a cyclical pattern or autocorrelations in the returns). Therefore we establish an experiment where subjects have to make choices between several (2, 3 or 4) investment alternatives (“funds”), which can be from one out of three different types of time series of profits, with respectively no (autocorrelation) structure, limited structure and a high level of structure. The time series of profits of the different investment alternatives are exogenously generated prior to the experiment. Like in real life, the subjects in this experiment do not know the data

¹We use profits instead of return rates in order to make the decision problem as simple as possible for the experimental subjects.

generating process of the profits.² Subjects makes these choices for 80 consecutive periods, where the time series of profits is changed after the first 40 periods. We show the subjects different types of time series in order to make it possible to investigate whether and how decisions depend upon the features of the time series of profits.

Return chasing behavior in mutual fund choice can be studied by discrete choice models, where agents switch from one alternative to another over time on the basis of the relative past performances of alternatives. These models usually assume that people tend to switch to alternatives that generate higher profit in the recent past. Discrete choice models have been utilized to explain different economic phenomena but, following the influential work by Brock and Hommes (1997, 1998) they have become particularly popular in the literature on heterogeneous agents models of financial markets. In these models a population of traders is distributed over a small number of *trading* or *prediction heuristics* that determine the behavior of these traders. The market price, which results from the aggregation of these investment strategies, may then be used to evaluate the performance of these heuristics. When traders switch between the heuristics on the basis of their performance, the distribution of the population of traders over the heuristics will continually be subject to change. Although very promising as a behavioral model of financial market dynamics, the dynamics of these heterogeneous agents models typically depends crucially on the exact heuristics that are used. Moreover, the outcomes of heterogeneous agent models depend critically upon how the switching process is modelled exactly, which parameters are included and their precise values used for analysis and simulations. No convincing empirical validation which favours one specification of the switching process and/or values of the learning parameters over another has been provided yet. This fact significantly impedes application of heterogeneous agent models. One of the objectives of the current paper is to get a better insight into how switching between heuristics is optimally modeled.

A simple version of the discrete choice model can be written as:

$$n_{k,t+1} = \frac{\exp[\beta\pi_{k,t}]}{\sum_{j=1}^K \exp[\beta\pi_{j,t}]}, \quad (1)$$

where $n_{k,t}$ is the fraction of agents using alternative k at time t , and $\pi_{k,t}$ is the

²Anufriev et al. (2012) present a similar experiment where subjects do know the data generating process.

realized payoff generated by choosing alternative k at time t . This is a backward looking model, so the fraction of agents choosing alternative k in period $t + 1$ depends (positively) upon $\pi_{k,t}$. The idea behind this model is very intuitive: if past performance contains information about future performance, the agents will tend to choose, for the current period, the alternative that was more profitable in the recent past. In order to understand this model better and apply it to real data, it is very important to estimate the switching rule, especially the intensity of choice β in the model. The problem of identifying the main determinants of switching behavior, especially the intensity of choice parameter β has far not been satisfactorily resolved. Besides studying return chasing behavior in the choice of mutual funds, one goal of this study is to make a systematic estimation of the parameter β in this model. We want to address the following questions: (1) Is there a universal β that can fit people's decision in all environments? (2) Is there heterogeneity in β between individuals? (3) Is β affected by experience?

The experimental results show that subjects switch a lot between the funds, and the switching behavior is usually driven by past performance of the funds. The earnings are generally higher when there is structure (e.g. autocorrelation) in the data. A simple discrete choice model with one lag and a predisposition effect provides a good fit to the data when there is no cyclical/periodic pattern in the return time series. When there is a cyclical pattern in the time series of profits, the model needs to incorporate more lags in order to provide a good description of the data, which suggests that people refer to the cyclical pattern of the return when they make the investment decision. The estimated intensity of choice β is therefore heavily influenced by the structure of the time series of the profits of funds. There is some heterogeneity in the intensity of choice parameter β between agents. There is no evidence that experience has a strong impact on the switching behavior.

Our work is related to recent work of Anufriev and Hommes (2012a, 2012b) where the evolutionary path of several heuristics commonly used in 'learning-to-forecast' experiments (see e.g. Hommes et al. (2005) and Heemeijer et al. (2009)) are calibrated. They show that three different kind of aggregate outcomes (monotonic convergence, oscillatory dampened price fluctuation and persistent price oscillations) can find their driving forces in the evolutionary selection among some simple forecasting rules by the market participants within a universal framework. But as the data of the experiments are in the form of price prediction only, it is very difficult to make direct analysis on the weights of the participants following different market heuristics. In the experiment presented in this paper, the partici-

pants will be asked to choose just one alternative, which in some cases corresponds to the profit generated by a particular prediction heuristic. That will make it much easier to do direct econometric analysis on the relationship between weights the strategies gain and their relative performance. There are also empirical works that estimate the discrete choice model using data from financial markets, the inflation rate and real estate market, such as Branch (2004), Boswijk et al. (2007), Cornea et al. (2012), Frijns et al. (2010), Goldbaum and Mizrach (2008), Franke and Westerhoff (2011), de Jong et al. (2009), Jongen et al (2012), Kouwenberg and Zwinkels (2010).

There are not many experimental finance works where subjects make binary choices. To our knowledge there is only an experimental literature on herding behavior (Anderson and Holt, 1997, Drehmann et al, 2005, Alevy et al, 2007), or more specifically “information cascades” where subjects choose between two assets, one of which generates a positive payoff and where the other pays nothing. They receive a public signal as well as a private signal about which asset will pay the positive dividend. The public signal is the history of the choices of former subjects who faced the same problem. An information cascade happens when they purely base their decision on the public signal, even when the private signal suggests them to choose the other asset. Anderson and Holt (1997) also estimate a discrete choice model on the probability of choosing one asset, and report the estimated β , which usually lies between 2 and 8. While in their experiment, different subjects play in different periods and the payoff of the assets are more or less repeated (each asset has equal probability to generate positive and zero payoff), in our experiment the same subject makes decision for all periods, but the payoff of the assets are time varying.

Our work is also related to an experimental study on portfolio choice by Bossaerts et al. (2007). They investigate the relation between the weight people put on one financial asset and its past performance. The difference with our experiment is that subjects in their experiment choose the fraction of wealth to be put in each asset, which is a continuous decision variable, while in our experiment the subjects make discrete choices.

The remainder of the paper is organized as follows. Section 2 discusses the experimental design and Section 3 discusses the experimental results. In Section 4 the discrete choice model discussed above is estimated from the experimental data. Section 5 concludes. The Appendix contains the instructions given to the subject of the experiment, and some estimation results.

2 Experimental Design

The experiment was conducted on June 17-18, 2010, at the CREED laboratory of the University of Amsterdam. A total of 91 student subjects participated in four sessions, and none of them participated in more than one session. Each session lasted for about two hours, and the subjects' payoffs varied between 18 and 25 euros.

2.1 Subjects' Task

The experiment is an individual choice experiment where subjects have to repeatedly make an investment decision. In particular, the subjects observe the time series of realized profits of either two, three or four investment alternatives, denoted A , B , C and D , in two blocks of 50 periods. The investment alternatives are different in the two blocks (and the number of alternatives is also different in these two blocks). Each block of 50 periods starts with a given history of 10 periods, so that the subject can get acquainted with the investment alternatives. For the remaining 40 periods the participant has to submit an investment decision for each period sequentially by selecting one of the possible alternatives. After a decision is made for a certain period the realized profits of all available investment alternatives for that period are revealed.

The payoff for a subject in any particular period is equal to the profit of the alternative they choose in that period. This profit ranges between 0 points and 16 points, and the total number of points earned during the experiment is translated to euros, with the number of euros given by 0.05 times the number of points (that is, subjects receive one euro for every 20 points they earned by their investment decisions). In addition subjects receive a 'show-up' fee of 5 euro.

Figure 1 shows an example of the computer screen the subjects face during the experiment. The subjects are presented with a graph of past profits in the bottom left part of the screen and the same information is given, in numbers, in the table in the bottom right part. The subject then has to make a discrete choice on which alternative to invest in, using the radio boxes in the top part of the screen. After all subjects have done this the computer screen is updated with the new information (realized profits of the different alternatives of the previous period, and the choice of the subject).

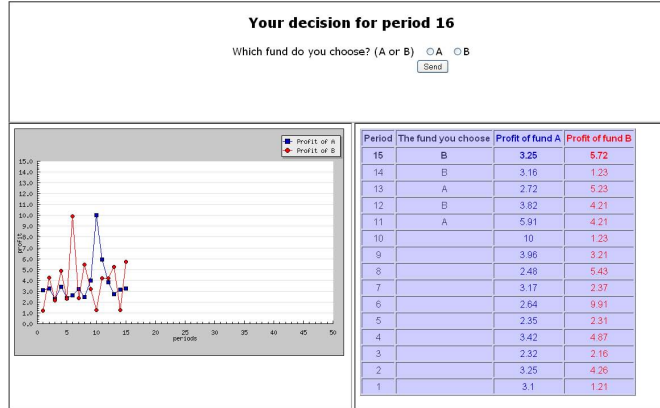


Figure 1: An example experimental screen. The subjects submit their decision by choosing between the radio icons in the top box, and can refer to information of past profits of the funds in the graph and table below.

2.2 Time series

We construct seven different sets of time series of profits for investment alternatives for 50 periods. The *data generating processes* (DGPs) used to obtain these time series, which can be divided into three categories, are explained below.

1. **White Noise time series** (one set, denoted *WN2*). A set of two white noise time series. Both time series are generated by the DGP

$$\pi_{i,t} = 5 + \varepsilon_{it}, \quad i = A, B \text{ and } t = 1, 2, \dots, 50, \quad (2)$$

where ε_{it} is independently and identically distributed, with $\varepsilon_{it} \sim N(0, 1)$. The realized time series are shown in Figure 2. Note that these two time series are independent from each other, and serially uncorrelated.

2. **Brock-Hommes time series** (three sets, denoted *BH2*, *BH3* and *BH4*, respectively). To generate these time series we simulated the two canonical heterogeneous agents models of Brock and Hommes (1997, 1998). To be specific, the set denoted *BH2* contains the two time series of profits generated by the deterministic nonlinear cobweb model presented in Brock and Hommes (1997), with parameter values $\beta = 3.9, \eta = 0.5, \delta = 0.2$ (see Brock and Hommes (1997), p. 1066-1067). One time series corresponds to the profits generated by the rational expectations heuristic (net of information costs) and the other time series corresponds to the time series generated by the naive expectations heuristic.

The three time series in *BH3* are generated by the deterministic nonlinear asset pricing model with three heuristics described in Brock and Hommes (1998, p.1240-1242), with intensity of choice parameter $\beta = 450$. We added a shock to the price in every period. This shock is independently and identically drawn from $N(0, 0.1)$. The four time series in *BH4* are generated by the asset pricing model with four heuristics in Brock and Hommes (1998), with intensity of choice parameter $\beta = 95$, and again an idiosyncratic i.i.d. shock, in this case drawn from $N(0, 0.12)$. In order to make the profits for the different time series comparable, we normalize these time series such that profits are always in the range $(0, 13)$, with a mean of around 5.³ The time series from *BH2*, *BH3* and *BH4* are shown in Figure 3. The time series from *BH2* have a clearly recognizable, repeated, quasi-periodic pattern. Their ‘cycle’ is, however, fairly unpredictable. On the other hand, the time series from *BH3* and *BH4* appear to have little structure.

3. **Stock Index time series** (three sets, denoted *SI2*, *SI3* and *SI4*, respectively). These time series are constructed from four actual stock indices, in particular the *Austrian Trade Index* (ATX), the *Belgium 20 Stock Index* (BFX), the *Dow Jones Index* (DJI) and the *FTSE 100 index* (FTSE) from October 2005 to November 2009. We use rotation based yearly return, instead of monthly returns. We normalize the data to fit in the range $(0, 13)$ by dividing the profits by 12 and adding 5.⁴ We use the ATX and BFX indices for *SI2*, the ATX, BFX and DJI indices for *SI3* and the ATX, BFX, DJI and FTSE for *SI4*. These time series are shown in Figure 4. Obviously, all of these time series are highly correlated.

We make sure that the maximum, minimum and average profit in each type of time series is similar. The descriptive statistics can be found in Table 1. This makes sure that when we estimate the discrete choice model in later sections, the estimated intensity of choice variable is free from a scaling effect.

The three types of time series do differ substantially in their autocorrelation structure. The time series from *WN2* exhibit almost no autocorrelation or periodic

³We achieved this by multiplying the time series from *BH2* by 2 and subsequently subtract 2, and by multiplying the time series from *BH3* and *BH4* by 75 and adding 3.

⁴We understand that after this adjustment the return time series does not look like the typical monthly returns of mutual funds. Here, the label “stock index” is only due to the source of the data, and the main purpose is to study the switching behavior when profits are highly autocorrelated. The results from the “white noise” time series have more relevance for the behavior of investors in actual markets for mutual funds.




Figure 2: The profit time series generated by white noise models.




Figure 3: The profit time series generated by Brock-Hommes models.

pattern. The time series from *BH2* have a strong quasi-periodic pattern, and the time series from *BH3* and *BH4* exhibit chaotic patterns. Finally, the time series from *SI2*, *SI3* and *SI4* exhibit strong autocorrelation.

Figure 5 shows the autocorrelation patterns of the different types of time series. Since we conjecture that for a subject's decisions the differences between profits of different alternatives are more important than the absolute value of these profits,

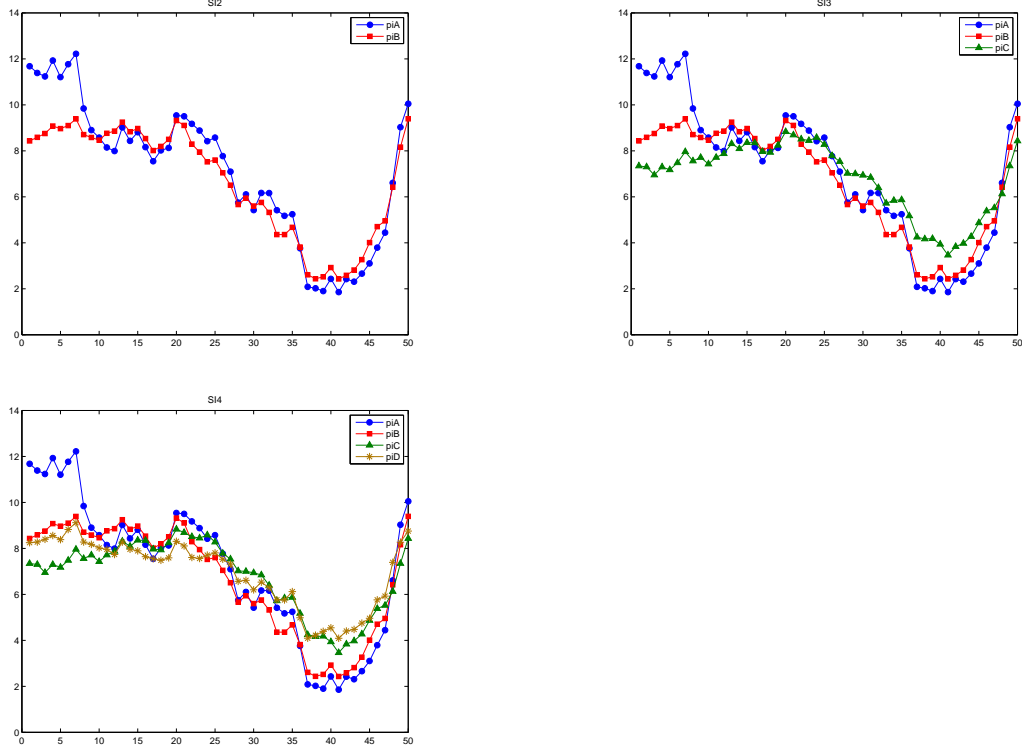


Figure 4: The profit time series from returns to stock indices.

we plot the autocorrelation of profit differences, where we take alternative A as benchmark. The profit difference is then calculated as the profit of A minus the profit of one of the other alternatives. The autocorrelation is plotted until the 10th lag. According to the Ljung-Box Q test, the autocorrelation is not significant for any of the lags of the time series from $WN2$, significant at the 5% level for all lags of the time series from $BH2$, $BH3$ and $BH4$, except for the profit difference between heuristic A and heuristic D from $BH4$, for which none of the lags is significant. As might be expected, the autocorrelation for all profit differences of the time series from $SI2$, $SI3$ and $SI4$ is significant for all lags and also higher than in the other time series (note that, since the time series from $SI2$ and $SI3$ are part of $SI4$ as well, we only include autocorrelation plots for $SI4$). The autocorrelation of profit differences is typically higher in the time series from $BH2$ than in the time series from $BH4$ and, to a lesser extent, those in $BH3$. Moreover, autocorrelations are very small for the time series from $WN2$, as they should be.

Discrete choice models, as they are used in heterogeneous agent models, are typically backward looking models, with the assumption that agents only care about the size of past profits, and do not use patterns in past profits to predict future profit levels. The implication of this is that there should be no signifi-

Time series	min	mean	max
WN2	0.60	5.64	10.25
BH2	0.47	5.58	9.42
BH3	0.87	4.55	12.69
BH4	1.33	5.64	10.64
SI2	0.36	4.66	8.04
SI3	0.36	4.27	8.04
SI4	0.36	4.89	8.04

Table 1: Descriptive statistics about profits from each set of time series.

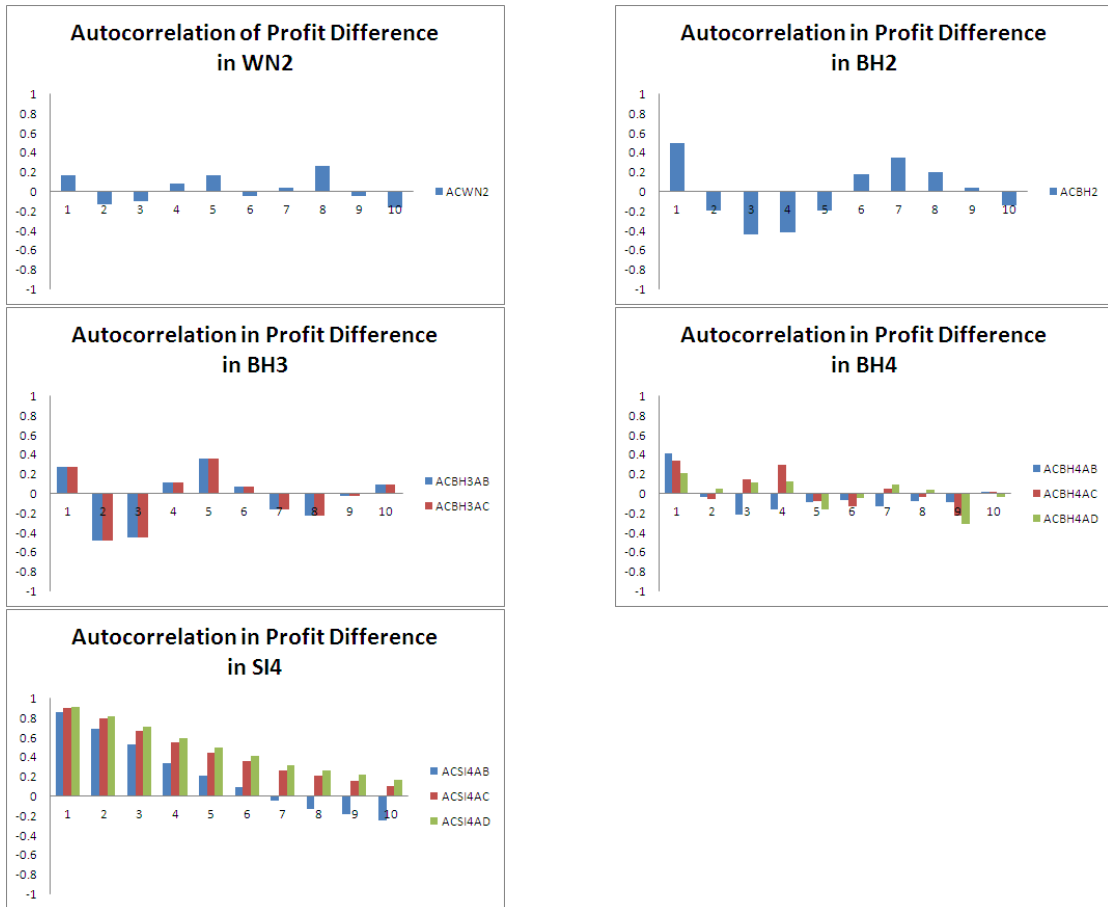


Figure 5: The autocorrelation of the difference of profits in each type of time series. The numbers on the horizontal axis indicate the order of lags.

cant differences between the discrete choice models that are estimated from the experimental data on the different time series.

On the other hand, if significant differences do exist, this may indicate that subjects try to forecast future profits by extrapolating patterns in past profits (or

Time-series	Unexperienced	Experienced	Total
WN2 (T1 and T6)	16	13	29
BH2 (T2 and T5)	15	14	29
BH3 (T3 and T1)	14	16	30
BH4 (T4)	No	14	14
SI2 (T4 and T3)	14	14	28
SI3 (T5)	14	No	14
SI4 (T6 and T2)	13	15	28

Table 2: General setting of the experiment. This table shows the number of participants in each treatment.

in profit differences) and (partly) base their investment decisions upon that.

2.3 Treatments

As discussed above, we generated seven sets of time series (*WN2*, *BH2*, *BH3*, *BH4*, *SI2*, *SI3* and *SI4*), where the range of profits and the average profit for each time series is similar, but profit differences and (autocorrelation) structure in the time series of profits is different. During the experiment each subject will be confronted with exactly two of these seven sets of time series. Note that the subjects do not know the data generating process that underlies the time series they observe.

Of the 42 different possible combinations of sets of times series, we selected six. We made sure that for every participant both the number of investment alternatives (2, 3 or 4) and the type of investment alternatives (*WN*, *BH* or *SI*) was different in the two blocks. The resulting six treatments are denoted by *T1* (first *WN2*, then *BH3*), *T2* (*BH2* and *SI4*), *T3* (*BH3* and *SI2*), *T4* (*BH4* and *SI2*), *T5* (*SI3* and *BH2*) and *T6* (*SI4* and *WN2*). The summary information of the experimental design is also shown in Table 2.

Note that each set of time series is used in exactly two treatments (except for sets *BH4* and *SI3*), once with unexperienced subjects (when it is used in the first block) and once with experienced subjects (when the set is used in the second block), where experience refers to experience with another set of time series. Table 2 also shows the number of participants in for each treatment and for each time series.

3 Experimental Results

In this section we give a, mostly, qualitative discussion of the experimental results. In Section 3.1 we investigate the dynamics of choice, and in Section 3.2 we study the number of switches on an individual level. Section 3.3. is concerned with the question whether the decisions of subjects improve over time and finally, in Section 4 we investigate to what extent subjects condition their decision on payoff information from the previous period.

3.1 Fraction of Choice

Figure 6 presents an area chart that plots the time series of the shares of subjects choosing fund A , B , C and D , respectively, in each period. The results are presented for each set of time series, and separated for experienced and unexperienced subjects (this gives 12 plots, since there are no unexperienced subjects confronted with the time series from $BH4$ and no experienced subjects confronted with the time series from $SI3$, see Table 2). There is substantial fluctuation in the shares during the course of the experiment, which suggests that at least some subjects switch between alternatives regularly during each block of the experiment. Moreover, the dynamics of shares are quite similar for unexperienced and experienced traders that face the same set of time series, as is confirmed by a visual comparison between the left and right panels of Figure 6 (this obviously does not hold for the bottom two panels, which refer to time series from different sets). This suggests that there is not a strong learning effect between the two blocks of the experiment.

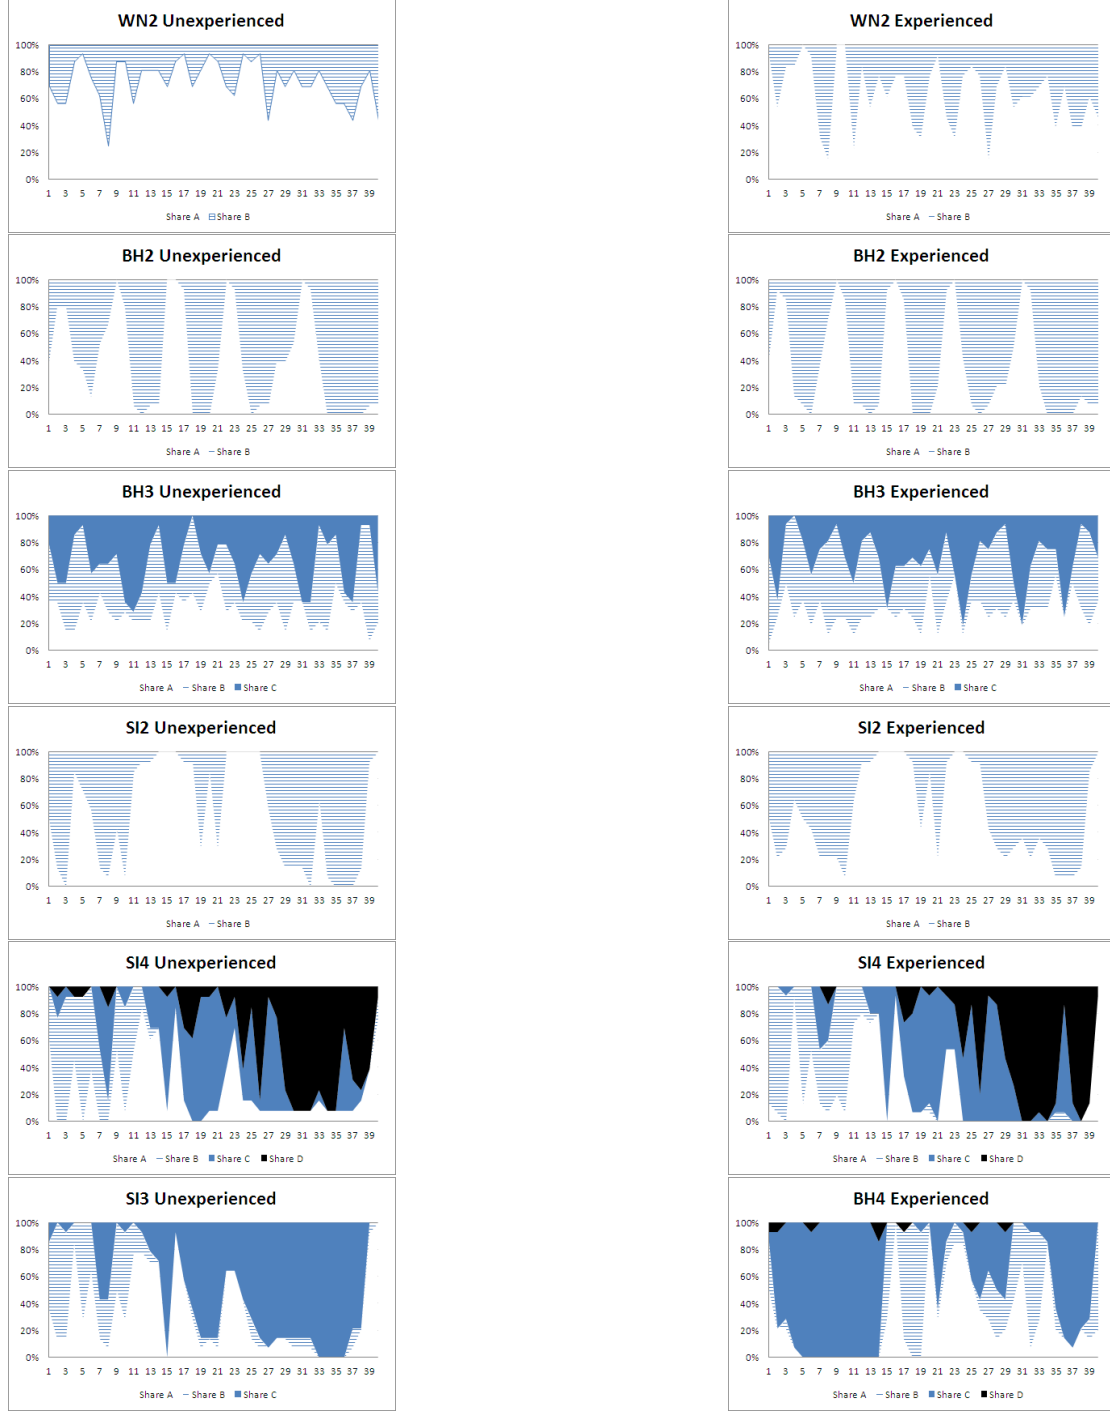


Figure 6: Shares of choices. From top to bottom: shares for *WN2*, *BH2*, *BH3*, *SI2*, *SI4*(unexperienced in left panel and experienced on right panel), *SI3* (bottom left) and *BH4* (bottom right).

It is remarkable that, for the time series from *WN2* (top two panels of Figure 6), although there is no difference in the average or expected profit of alternatives *A* and *B*, people still choose *A* with a higher frequency than *B*. This may be partly due to some “default choice bias” for alternative *A*. However, when the

Treatment	Block	Total Number of Switches	Switches per Person per Period	Time Series
1	1	210	0.33	WN2
	2	341	0.53	BH3
2	1	159	0.27	BH2
	2	254	0.42	SI4
3	1	308	0.55	BH3
	2	129	0.23	SI2
4	1	147	0.26	SI2
	2	210	0.38	BH4
5	1	169	0.30	SI3
	2	153	0.27	BH2
6	1	218	0.42	SI4
	2	228	0.44	WN2

Table 3: Number of switching in each treatment.

same subjects (from treatment $T1$) are faced with the $BH2$ time series, where alternative B is typically more profitable than alternative A , they do choose B with a higher frequency than A . Also for the time series from $SI4$, the majority of subjects learns to switch from fund A (which is initially the most profitable alternative on average) to fund D (which is the most profitable alternative, on average, for the last couple of periods). These observations suggests that the subjects do respond to profit differences in the correct way.

3.2 Frequency of Switches

To get a more precise idea of the frequency with which subjects switch between alternatives we count the number of switches for each treatment. A switch occurs when a subject chooses differently in period t than in period $t - 1$. As the number of subjects is different for different treatments, we use the number of switches per subject per period as a measure to compare different treatments. The results are shown in Table 3.

We can see from this table that the average frequency with which subjects switch in a given period ranges from 23% to 55%. The frequency of switching is very similar for subjects faced with the same time series (with the largest differences between subjects facing $WN2$, which makes sense, given that expected payoffs for alternatives are always the same for those time series), suggesting that

behavior of the subjects is quite consistent and that learning between blocks does not seem to play a role.

Figure 7 shows the empirical cumulative distribution function of the number of switches subject make for the different time series. The horizontal axis presents the number of switches made by a subject (which can range from 0 to 39) and the vertical axis gives the percentage of subjects switching a given number of time. We see most subjects switch between 10 and 20 times, and that very low and very high frequencies of switching (more than 25) are rare (the exception being subjects facing the time series from *BH3*, who change relatively often).

For a majority of cases, the distribution of frequencies of switching are similar for unexperienced and experienced subjects (the exception arising for the time series from *WN2*, where experienced subjects switch more often than unexperienced subjects, and for the time series from *SI2*, where experienced traders switch less often).

3.3 Learning Effect

We want to see whether the subjects learn during the experiment by investigating whether they perform better in later periods (within the same block) than in earlier periods. Figure 8 plots the time series of the fraction of subjects who choose the best alternative for that period (that is, the alternative that generates the highest profit in that period) for each set of time series. If subjects learn to make better decisions this fraction should increase over time.

From Figure 8 we observe that, again, experience with another time series does not improve subject's choices. Whether, within one particular set of time series there is a learning effect, depends upon the type of time series. For the time series from *BH2*, *SI2* and *SI3* we see an overall increase in the fraction of optimal choices (note that these are also the time series with the most clear structure), while for the other time series, there is not such a clear trend (for the time series from *SI4* this might be due to the fact that there is, compared to *SI3*, an additional fund which, for many periods towards the end, is not much worse than the optimal fund).

Table 4 shows the average fraction of subjects choosing the best fund in each phase. The fraction of subjects choosing the best alternative is very similar for unexperienced and experienced subjects (in fact, for three of the five cases unex-

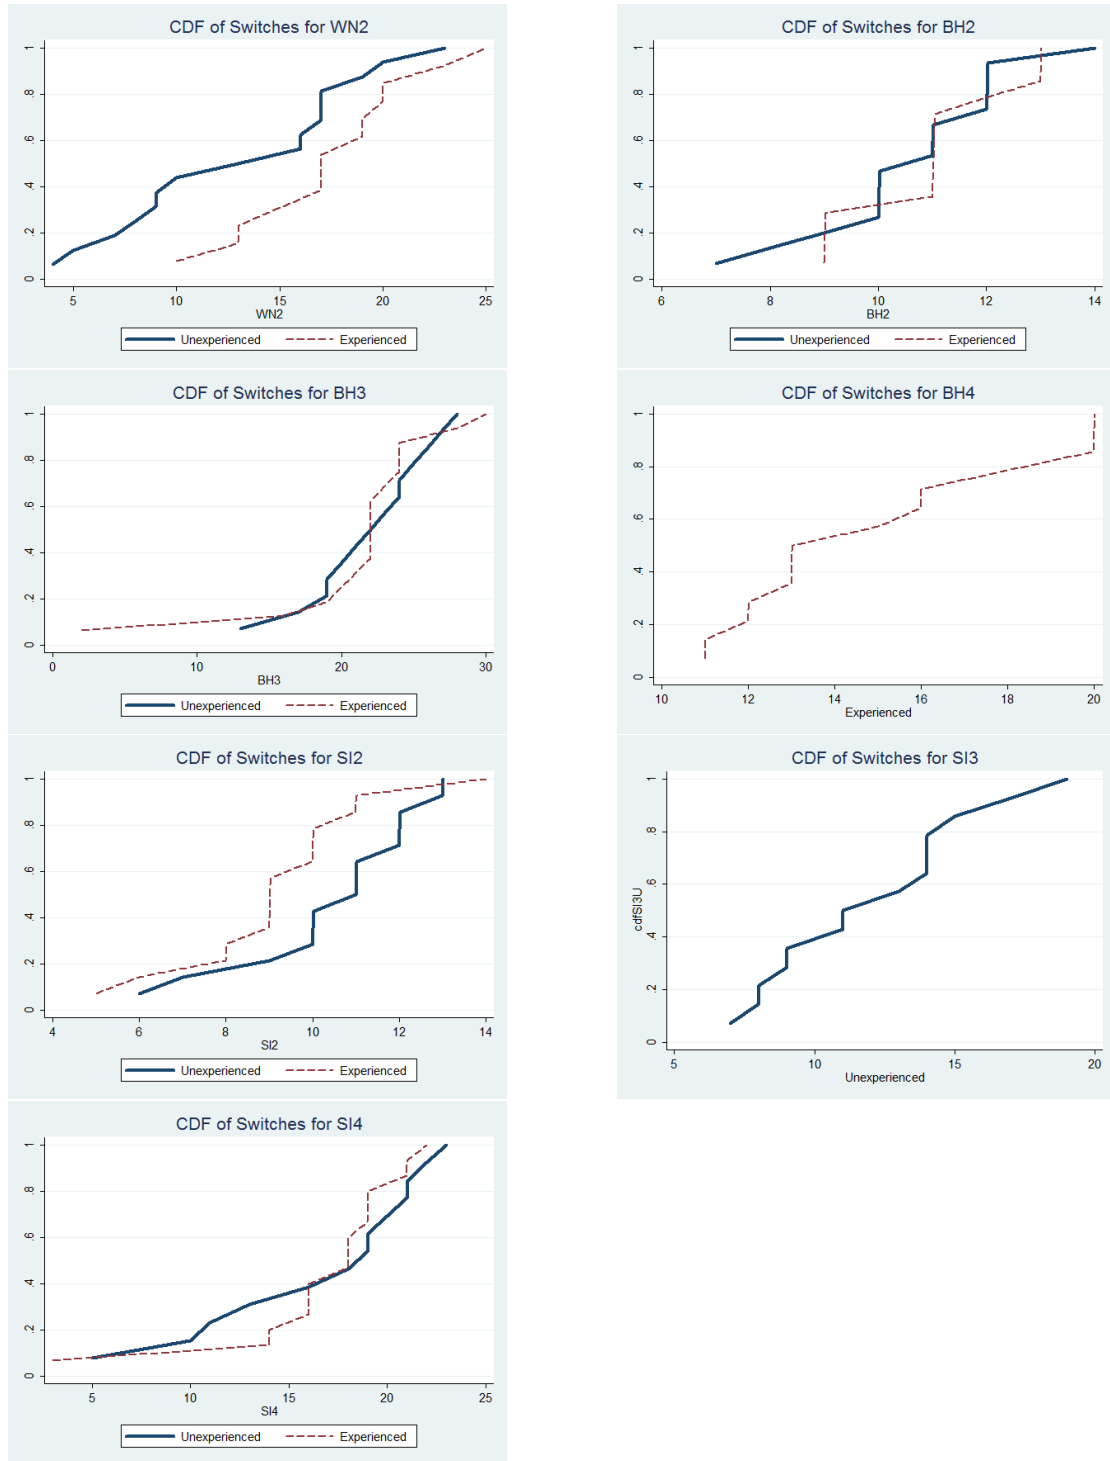


Figure 7: Empirical CDF function of the number of switches in different treatments. The vertical axis is in terms of percentage.

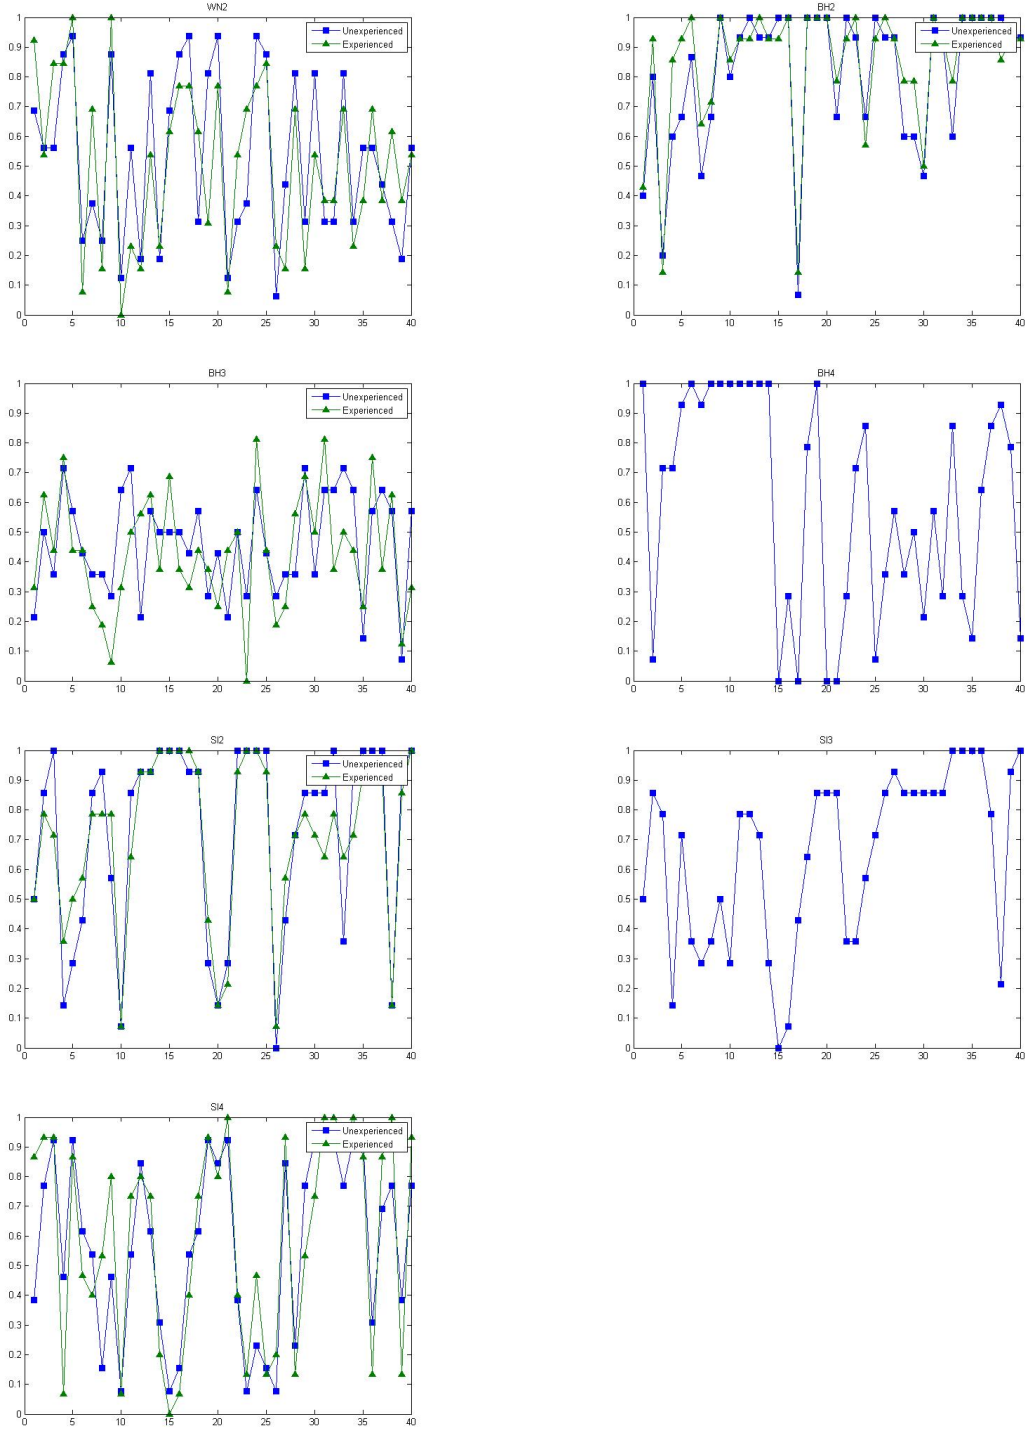


Figure 8: This figure shows the fraction of people choosing the best fund. From the top to the bottom: time series of $WN2$, $BH2$, $BH3$, $BH4$, $SI2$, $SI3$, $SI4$. Square indicates unexperienced subjects, and triangle indicates experienced subjects.

perienced subjects do slightly better than experienced subjects, but none of the differences between unexperienced and experienced subjects is significant at the 5% level according to a Mann-Whitney-Wilcoxon test). Subjects are the most

Time-series	Unexperienced	Experienced	Total
WN2 (T1 and T6)	0.5313	0.5115	0.5214
BH2 (T2 and T5)	0.8133	0.8500	0.8317
BH3 (T3 and T1)	0.4625	0.4313	0.4469
BH4 (T4)	No	0.5768	0.5768
SI2 (T4 and T3)	0.7250	0.7089	0.7170
SI3 (T5)	0.6518	No	0.6518
SI4 (T6 and T2)	0.5692	0.5967	0.5829

Table 4: Fraction of best choice in each treatment and phases.

successful in choosing the best alternative with time series from *BH2*, *SI2* and *SI3*, which have substantial structure. The fraction of optimal choices is lowest with the highly unpredictable time series from *BH3*. Remarkably, the fraction of optimal choices with the time series from *WN2* is higher than 0.5 (but it is not significantly different from 0.5 at the 5% level, according to a Wilcoxon sign rank test).

3.4 Influence of Past Performance

Figure 9 plots, for time series with only two alternatives, the fraction of subjects choosing alternative *A* against a binary variable that equals 1 if the return of *A* is larger than that of *B* in the last period, and 0. This figure illustrates whether there is past return chasing behavior. If the subjects are past return chasing, the fraction should go down (up) when the binary variable is 0 (1). This is almost always the case with the unexperienced subjects that face the time series from *WN2*, and a little less the case with the experienced subjects facing those time series. For the time series from *BH2*, the data also suggests that subjects are return chasing, although the fraction of choice of *A* goes up (down) when the return of alternative *A* has been lower (higher) for several periods. This may be due to the fact that time series from *BH2* exhibit a quasi-periodic pattern, and the cycle tends to be longer in the later periods of the experiment than in the earlier periods. The subjects might therefore choose the alternative that has generated lower return for several periods, anticipating that the order of profits will be reversed in the next period. The data for the time series from *SI2* also provide support for return chasing behavior. The fraction of subjects choosing *A* almost always goes up (down) when the return of *A* is higher (lower) in the

last period, except for the unexperienced subjects at the very beginning of the experiment.

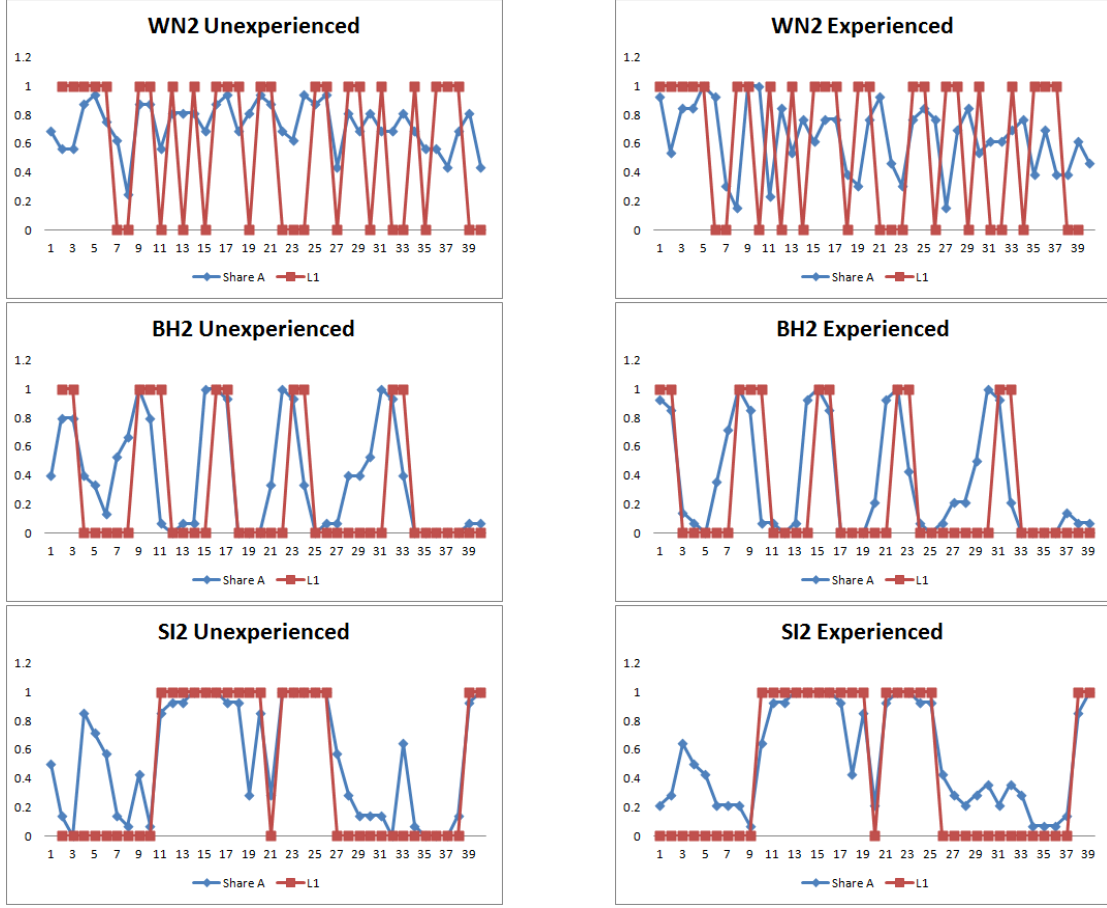


Figure 9: The fraction of subjects choosing fund A (diamonds) plotted against the binary variable that indicates whether fund A generates higher return than fund B in the last period (squares).

These observations suggest that a model where subjects condition their decision upon the profits of the different alternatives from the previous period might give a good description of the experimental data. In the next section we try to estimate such a model.

4 Describing subject behavior by the discrete choice model

In this section we try to estimate a discrete choice model, and in particular the ‘intensity of choice’ parameter β , from the experimental data. In theoretical het-

erogeneous agents models that use discrete choice dynamics, the parameter β is typically exogenously given. In this section we want to investigate the potential success of the discrete choice model as a description of the switching behavior of subjects to our experiments. Moreover, we want to investigate whether β is affected by the decision environment.

In Section 4.1 we will consider the *binary choice* model, using experimental data obtained from the choices with time series from *WN2*, *BH2* and *SI2*. In Section 4.2 we then focus on the *multiple choice model*, for the experimental data with three or four alternatives.

4.1 The Binary Choice Model

For the six blocks involving choices between two alternatives, A and B (with time series from *WN2*, *BH2* and *SI2*) we estimate the model on individual data. The binary choice model for individual i given by probabilities of choosing alternative A at time t as:

$$\Pr(c_{i,t} = A | \mathcal{I}_{t-1}) = \frac{\exp U_{t-1}^A}{\exp U_{t-1}^A + \exp U_{t-1}^B}, \quad (3)$$

where $U_{t-1}^A = \alpha + \beta \pi_{t-1}^A$ and $U_{t-1}^B = \beta \pi_{t-1}^B$.

Here π_{t-1}^A (π_{t-1}^B) is the profit that would have been generated by choice A (B) in period $t-1$. The parameter α measures a ‘predisposition effect’: if α is positive (negative) a subject tends to choose alternative A (B), if profits are the same. The parameter β is the familiar ‘intensity of choice’ parameter, measuring how sensitive participants are with respect to differences in profits. We are going to assume that α and β are the same for each individual in a particular block, but may differ between blocks.⁵

⁵In order to check whether there is heterogeneity in the intensity of choice β across individuals, we also estimated the following model

$$\Pr(c_{i,t} = A | \mathcal{I}_{t-1}) = \frac{\exp U_{t-1}^A}{\exp U_{t-1}^A + \exp U_{t-1}^B},$$

where $U_{t-1}^A = \alpha + \beta(1 + \gamma_i) \pi_{t-1}^A$ and $U_{t-1}^B = \beta(1 + \gamma_i) \pi_{t-1}^B$.

Here the estimation includes the individual-specific parameter γ_i (where $i \in \{1, 2, \dots, I\}$, with I the total number of subjects in the treatment). If there is a lot of heterogeneity in the intensity of choice β between individuals, we should see that γ_i is significantly different from 0 at the 5%

The estimation results are presented in Table 7 in the Appendix. We see that the level of intensity of choice is different for different time series. As the payoff in all treatments is on average 5 points per period, the differences between the estimated intensities of choice can not be attributed to a scaling effect. We find the estimated β to be quite high in the *SI2* sessions, and low in the *WN2* sessions. This result has an intuitive explanation: choosing the alternative that generated a high profit in the previous period is sensible in particular when profits are highly and positively autocorrelated. The results suggest that the intensity of choice β is affected by the structure of the time series of profits, and subjects are successful in adjusting their reaction to past information in the right direction (e.g. their behavior is described by higher β when past information has more predictive power about the future). There is also a difference between the estimated β with experienced and unexperienced subjects, but the pattern is not clear: in two of the three cases unexperienced subjects are more sensitive to profit differentials than experienced traders, and in the other case they are less sensitive.

Figure 10 compares the probability of choosing alternative *A* using the estimated discrete choice model (3) with the share of subjects actually choosing alternative *A* in the experiment. At the aggregate level, this simple model already provides a very decent description of the experimental data. According to numbers reported in Table 7, the mean squared error of the model is highest in the *BH2* treatments, probably due to a larger variation in the dependent variable, and is very similar in the *WN2* and *BH2* treatments.

We are also interested in whether the model can make good out-of-sample predictions. Therefore we estimate the model using the experimental data from the first 20 periods (periods 1-20), and use the estimated coefficients to predict the data in the second set of 20 periods (periods 21-40). In Figure 11 we plot the actual shares of the last 20 periods from the experiment, together with the out-of-sample predicted probabilities (based upon the first 20 periods) and the in-sample predicted probabilities (based upon all 40 periods). This figure suggests that the out-of-sample predictions are very close to the in-sample predictions.

level for many subjects. The level of heterogeneity turns out to be mild, however. The number of values of γ_i that are significantly different from 0 at the 5% level is 16 out of 86 subjects (18%, which can be broken down as follows: 3 out of 16 unexperienced subjects for *WN2* in treatment 1, 6 out of 13 experienced subjects for *WN2* in treatment 6, 1 out of 15 unexperienced subjects for *BH2* in treatment 2, 2 out of 14 experienced subjects for *BH2* in treatment 5, 1 out of 14 unexperienced subjects for *SI2* in treatment 4, and 3 out of 14 experienced subjects for *SI2* session in treatment 3.)

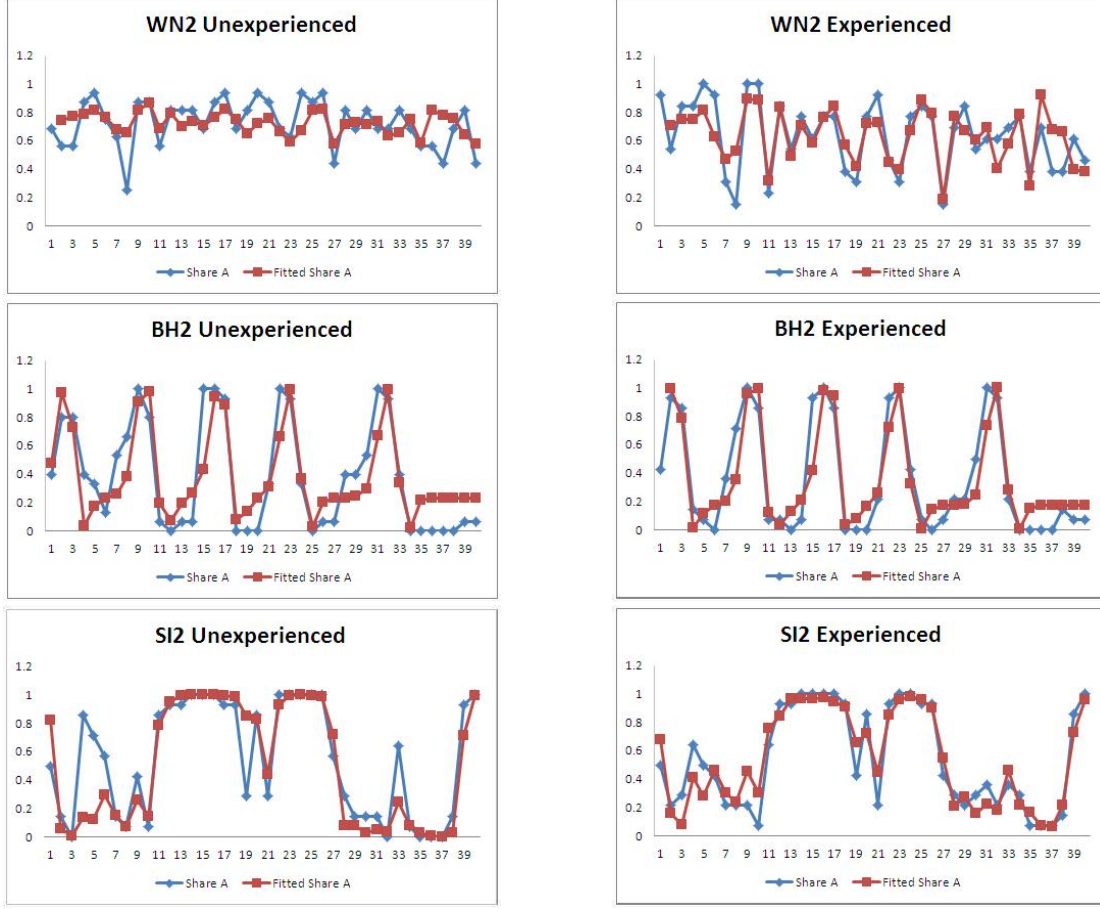


Figure 10: The actual and fitted time series by the simple 1 lag models. The diamonds are the experimental data, and the squares are the fitted fractions by the model.

This is confirmed by the mean squared errors for the in-sample and out-of-sample predictions presented in Table 5.

To check robustness of our results we estimated a more general form of the discrete choice model. This model, with l lags, is given by

$$\Pr(c_{i,t} = A | \mathcal{I}_{t-1}) = \frac{\exp U_{t-1}^A}{\exp U_{t-1}^A + \exp U_{t-1}^B},$$

$$\text{where } U_{t-1}^A = \beta_0^A + \beta_1^A \pi_{t-1}^A + \dots + \beta_l^A \pi_{t-l}^A,$$

$$\text{and } U_{t-1}^B = \beta_1^B \pi_{t-1}^B + \dots + \beta_l^B \pi_{t-l}^B.$$

Since the length of the “cycles” in the time series from *BH2* is typically fewer than 5 periods, we choose l to be maximally 5. We apply a likelihood ratio

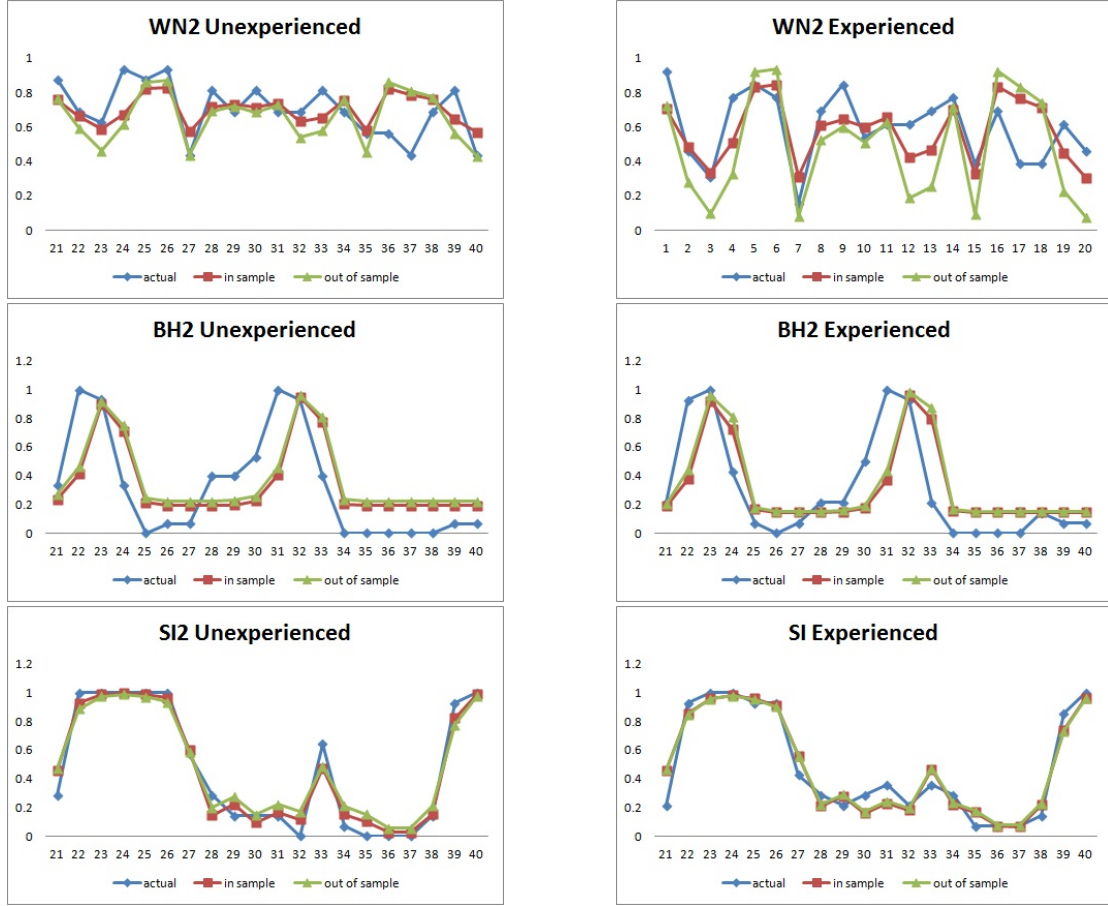


Figure 11: The actual and fitted time series by the simple 1 lag models in periods 21-40. The diamonds are the experimental data, the squares are the in sample prediction by the model using all data, and triangles are the out of sample prediction by the model using data in periods 1-20 only.

test iteratively to determine the number of necessary lags to include.⁶ We find that, when more lags can be included, the estimation results for the treatments with time series from *WN2* and *SI2* do not change very much. In fact, for the treatment with unexperienced subjects facing time series from *WN2* and for the treatment with experienced subjects facing time series from *SI2*, the estimated model does not incorporate more lags. For the treatment with experienced subjects facing time series from *WN2* and the treatment with unexperienced subjects facing time series from *SI2* the estimated model only involves one additional lag. For both treatments with time series from *BH2*, the estimated model incorporates up to the 4th lag. The estimation results are shown in Table 8 in the Appendix. Moreover, we observe some negative coefficients for higher order lags in the return,

⁶That is, we include new lags until the test does not reject the null hypothesis that the model with more lags is not an improvement over the model with fewer lags.

Time Series	Group	MSE In Sample	MSE Out of Sample
WN2	Unexperienced	0.0208	0.0291
	Experienced	0.0311	0.0796
BH2	Unexperienced	0.0732	0.0740
	Experienced	0.0688	0.0698
SI2	Unexperienced	0.0069	0.0109
	Experienced	0.0087	0.0089

Table 5: The mean squared error, for the last 20 periods, of the in-sample and out-of-sample predictions by the binary choice model (3).

especially for the treatments with time series from *BH2*. Figure 5 shows that the autocorrelation pattern in the profit difference in the time series from *BH2* is negative from the second to the fifth lag. The negative estimates for β in the higher order lags suggest that subjects recognize some of the long term reversion pattern in the profit dynamics, and use this information when choosing between alternatives *A* and *B*.

Figure 12 shows the fitted group with models with one lag and models allowing for more lags. Models allowing for more lags provide generally about the same level of fit as the model with one lag for the treatments with time series from *WN2* and *SI2*, and substantially outperform the model with one lag for the treatment with time series from *BH2*. The difference of the level of fitness can also be seen from the mean squared error numbers in Table 6.

Time Series	Group	MSE One Lag	MSE More Lags
WN2	Unexperienced	0.0210	0.0210
	Experienced	0.0299	0.0066
BH2	Unexperienced	0.0809	0.0254
	Experienced	0.0391	0.0167
SI2	Unexperienced	0.0219	0.0162
	Experienced	0.0140	0.0140

Table 6: The mean squared error of the in sample and out of sample prediction of the logit model with 1 lag and model that allows for more lags.

These findings suggest that the structure in the profit time series are an important determinant of the choice behavior of experimental subjects. When there is either almost no structure, or very much (autocorrelation) structure, subjects condition their choice only upon the most recent information. For the time series

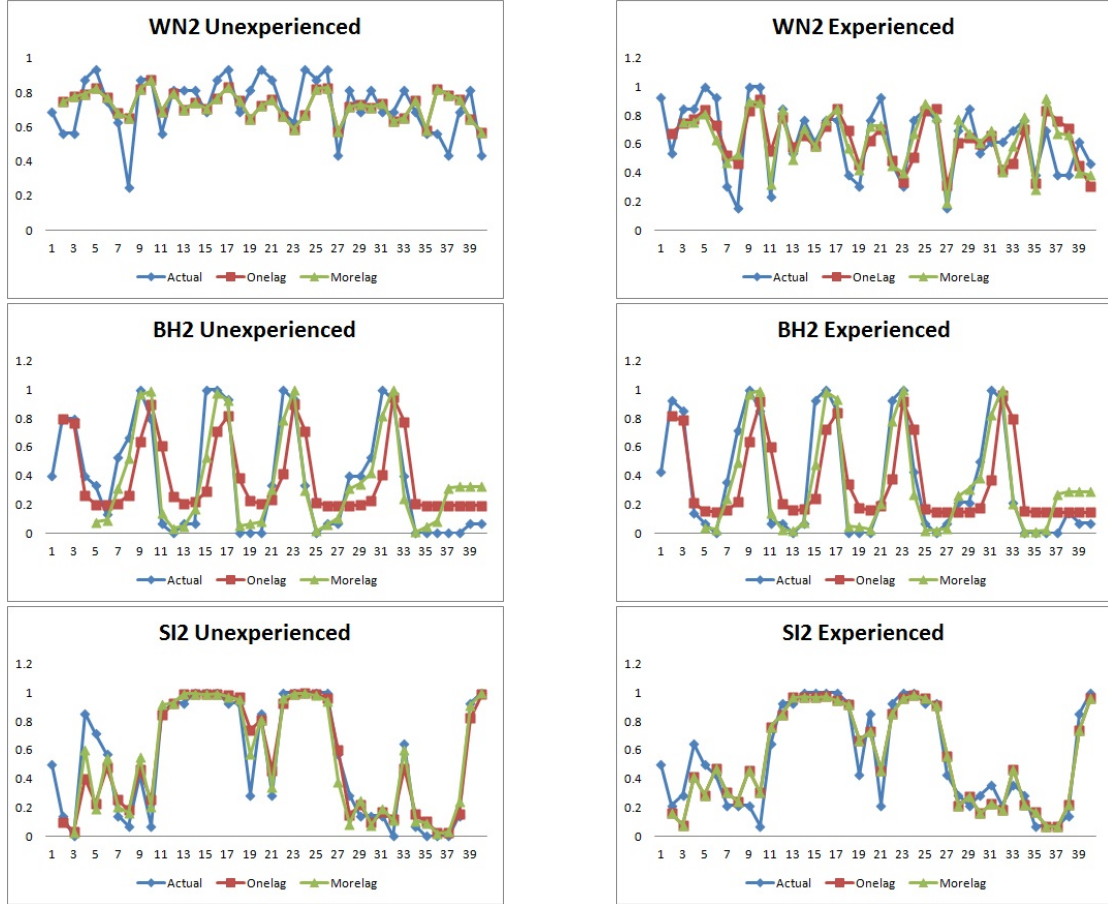


Figure 12: The actual and fitted time series by the simple 1 lag models and models that allow more lags. The diamonds are the experimental data, the squares are the fitted fractions by the 1 lag model, and triangles are models allowing for more lags.

from *WN2* there is not much information in either more recent or less recent data, and for the highly positively autocorrelated data, the most recent information already contains a lot of the information about the upcoming profit realizations. When there is a cyclical pattern in the data (as is the case with the time series from *BH2*), subjects need to use more past information because this information is helpful in predicting the reversal point of the current trend.

4.2 Multiple Choice Model

For sessions with three or four alternatives (that is, treatments with time series from *BH3*, *BH4*, *SI3* and *SI4* time series), we estimate a discrete choice model of the following form:

$$\Pr(c_{i,t} = A | \mathcal{I}_{t-1}) = \frac{\exp U_{t-1}^A}{\exp U_{t-1}^A + \exp U_{t-1}^B + \exp U_{t-1}^C},$$

when there are three alternatives, and

$$\Pr(c_{i,t} = A | \mathcal{I}_{t-1}) = \frac{\exp U_{t-1}^A}{\exp U_{t-1}^A + \exp U_{t-1}^B + \exp U_{t-1}^C + \exp U_{t-1}^D},$$

when there are four alternatives (the probabilities of choosing the other alternatives then follow directly from the estimated model).

We normalize the performance of the last alternative, and for the other performance measures we consider

$$U_{t-1}^K = \begin{cases} \alpha + \beta \pi_{t-1}^A & \text{if } K = A \\ \beta \pi_{t-1}^K & \text{if } K \neq A \end{cases}.$$

The result of the aggregate estimation can be found Table 9. The estimated β is very small or even insignificant for the treatments with time series from *BH3*, which indeed have very little structure, and relatively large for the treatments with time series from *SI4*. These estimation results confirm that the estimated value of β is to a large extent determined by the predictability of the profit time series.

Figure 13 plots the fractions simulated from the estimated models against the experimental data. The simulation reproduces the pattern in the data quite well.

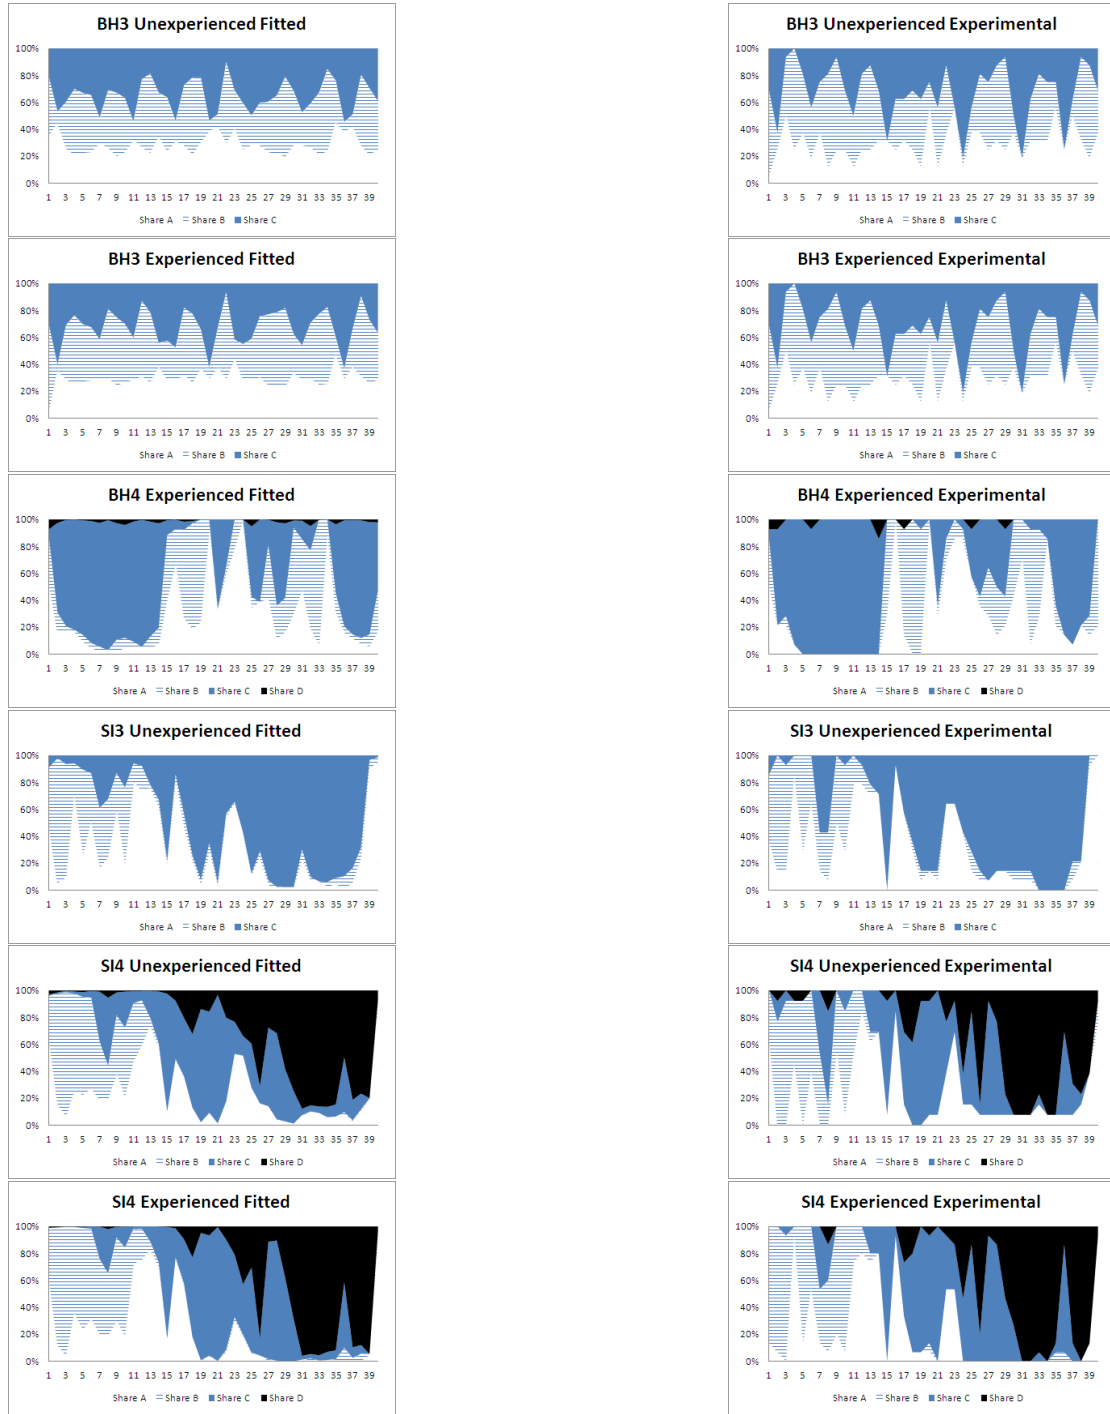


Figure 13: Simulated shares of choice by multiple choice model and the experimental data. The simulated shares are shown in the left panel, and actual experimental data are shown in the right panel.

5 Concluding remarks

We run an experiment where subjects have to make a discrete choice, in a number of subsequent periods, between several profitable investment alternatives (e.g. financial funds). Our aim is to investigate to what extent differences in profitability drives decisions of the subject, to identify what are the driving forces that make them switch between these alternatives and to analyze whether the well known discrete choice model (Brock and Hommes (1997, 1998)), which serves an important role in the literature on heterogeneous agents models, provides a good description of the switching decisions of the subject in the experiment.

The experimental data shows that human subjects switch often when they are faced with dynamic discrete choice between different investment alternatives. Moreover, this switching behavior is typically to a large extent driven by past performance of the different alternatives. When the time series of profits of the different alternatives exhibits a certain structure, such as autocorrelation and cyclical patterns, subjects are able to exploit this structure quite well. When there is limited structure in the time series of profits, e.g. when the time series are chaotic or random, and highly unpredictable, the typically tend to use a heuristic decision rule such as “choosing last period’s most profitable alternative”. Finally, we do not find any effect of experience on the quality or characteristics of the decisions made by the subjects.

The behavior of the subjects can be described by discrete choice models, such as those used in Brock and Hommes (1997, 1998), quite well. A simple discrete choice model, where decisions are only conditioned upon last period’s performance of the different alternatives, already captures the patterns in the data to a large extent, in particular when there is limited structure in the time series of profits. However, when there is a cyclical pattern or a high level of autocorrelation in the time series of profits, more lags need to be included into the model. This suggests that subjects pay attention to apparent structures in the past data. An important implication for this is that the discrete choice model, in the form in which it is used in a large part of the literature on heterogeneous agents models, may only be valid when the time series of the payoffs generated by different alternatives or heuristics exhibits limited structure.

There are several interesting possible extensions for future work in this direction. For example, one could think of an experiment where subjects are only informed about the payoff of the alternative they actually chose, in order to in-

investigate how they deal with the trade-off between exploration and exploitation. Another extension would be to subtract from the subject's earnings a small fee every time they switch between alternatives to investigate how large payoff differentials have to be in order for participants to change alternatives).

By considering exogenous payoff histories in the experiment we abstract from any influence the subject's decision may have on the actual outcomes. While this might be justified for some market environments, it does not hold in general. In fact, a key difference between natural and social sciences is that the latter often correspond to feedback systems: past market behaviour determines individual beliefs and choices, which in turn determine current aggregate market outcomes. An investor, for example, buys a stock that he expects to go up in the future based upon the past development of the price. If many speculators have the same beliefs and invest in this stock, demand for the stock goes up, as will the price, leading to an increase in the profits of the speculators investing in that stock. Therefore it might be worthwhile to consider a laboratory experiment in which this feedback is explicitly incorporated. Payoffs of the subjects will then also depend upon the decisions the other subjects in the experiment make.

References

- [1] Anderson, L., and C. A. Holt (1997): “Information Cascades in the Laboratory”, *American Economic Review*, **87**, 847-862.
- [2] Anufriev, M., T. Bao, A. Sutan and J. Tuinstra (2012): “A Simple Experiment on Fee Structure and Mutual Fund Choice”, *Mimeo*, University of Amsterdam.
- [3] Anufriev, M. and C. H. Hommes (2012a): “Evolution of Market Heuristics”, *Knowledge Engineering Review*, **27**, 255-271.
- [4] Anufriev, M. and C.H. Hommes (2012b): “Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments”, *American Economic Journal: Microeconomics*, **4**, 35-64.
- [5] Alevy, J.E., M.S. Haigh and J. A. List (2007): “Information Cascades: Evidence from a Field Experiment with Financial Market Professionals”, *Journal of Finance*, **62**, 151-180.
- [6] Audretsch, D. B., and G. Stadmann (2005): “Biases in FX-forecasts: Evidence from Panel Data”, *Global Finance Journal*, **16**, 2005, 99-111.
- [7] Bereby-Meyer, Y. and I. Erev (1998): “On Learning To Become a Successful Loser: A Comparison of Alternative Abstractions of Learning Processes in the Loss Domain”, *Journal of Mathematical Psychology*, **42**, 266-286.
- [8] Bossaerts, P., C. Plott, and W. Zame (2007): “Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiments”, *Econometrica*, **75**, 993-1038.
- [9] Boswijk, H.P., C.H. Hommes and S. Manzan (2007): “Behavioral Heterogeneity in Stock Prices”, *Journal of Economic Dynamics and Control*, **31**, 1938-1970.
- [10] Branch, W.A. (2004): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations”, *Economic Journal*, **114**, 592-621.
- [11] Brock, W. A. and C. H. Hommes (1997): “A Rational Route to Randomness”, *Econometrica*, **65**, 1059-1095.

- [12] Brock, W.A. and C.H. Hommes (1998): “Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model”, *Journal of Economics Dynamics and Control*, **22**, 1235-1274.
- [13] Camerer, C. F. and T.H. Ho (1998): “EWA Learning in Normal-form Games: Probability Rules, Heterogeneity and Time Variation”, *Journal of Mathematical Psychology*, **42**, 305-326.
- [14] Camerer, C., and T. H. Ho (1999): “Experience-Weighted Attraction Learning in Normal Form Games”, *Econometrica*, **67**, 827-874.
- [15] Cornea, A., C.H. Hommes and D. Massaro (2012): “Behavioral Heterogeneity in U.S. Inflation Dynamics”, CeNDEF Working paper 12-03 University of Amsterdam.
- [16] Drehmann, M., J. Oechssler and A. Roeder (2005): “Herding and Contrarian Behavior in Financial Markets: An Internet Experiment”, *American Economic Review*, **95**, 1403-1426.
- [17] Erev, I., Y. Bereby-Meyer and A.E. Roth (1999): “The Effect of Adding a Constant to All Payoffs: Experimental Investigation, and Implications for Reinforcement Learning Models”, *Journal of Economic Behavior & Organization*, **39**, 111-128.
- [18] Frijns, B., T. Lehnert and R. Zwinkels (2010): “Behavioral Heterogeneity in the Option Market”, *Journal of Economic Dynamics & Control*, **34**, 2273-2287.
- [19] Franke, R., and F. Westerhoff (2011): “Estimation of a Structural Stochastic Volatility Model of Asset Pricing”, *Computational Economics*, **38**, 53-83.
- [20] Goldbaum, D., and B. Mizrach (2008): “Estimating the intensity of choice in a dynamic mutual fund allocation decision”, *Journal of Economic Dynamics & Control*, **32**, 3866-3876.
- [21] Heemeijer, P., C.H. Hommes, J. Sonnemans and J. Tuinstra (2009): “Price Stability and Volatility in Markets with Positive and Negative Expectations Feedback: An Experimental Investigation”, *Journal of Economic Dynamics & Control*, **33**, 1052-1072.
- [22] Hommes, C.H., J. Sonnemans, J. Tuinstra and H. van de Velden (2005): “Coordination of Expectations in Asset Pricing Experiment”, *Review of Financial Studies*, **18**, 955-980.

- [23] Jain, C.J., and J.S. Wu (2000): “Truth in Mutual Fund Advertising: Evidence on Future Performance and Fund Flows”, *Journal of Finance*, **55**, 937-958.
- [24] de Jong, E., W. Verschoor and R. Zwinkels (2009): “Behavioural Heterogeneity and Shift-Contagion: Evidence from the Asian Crisis”, *Journal of Economic Dynamics & Control*, **33**, 1929-1944.
- [25] Jongen, R., W. Verschoor, C. Wolff, and R. Zwinkels (2012): “Explaining Dispersion in the Foreign Exchange Market: A Heterogeneous Agent Approach”, *Journal of Economic Dynamics & Control*, **36**, 719-735.
- [26] Kouwenberg, R., and R. Zwinkels (2010): “Chasing Trends in the U.S. Housing Market”, Erasmus University Rotterdam working paper.
- [27] Marcet, A. and J.P. Nicolini (2003): “Recurrent Hyperinflations and Learning”, *American Economic Review*, **93**, 1476-1498.
- [28] Patel, J., R. Zeckhauser and D. Hendrick (1994): “Investment Flows and Performance: Evidence from Mutual Funds, Cross-Border Investments and New Issues”, in: *Japan, Europe and International Financial Markets: Analytical and Empirical Perspectives*, edited by Sato, R., Levich, R., and R. Ramachandran, Cambridge University Press.
- [29] Sirri, E., and P. Tufano (1998): “Costly Search and Mutual Fund Flows”, *Journal of Finance*, **53**, 1589-1622.
- [30] Zheng, L. (1999): “Is Money Smart? A Study of Mutual Fund Investors’ Fund Selection Ability”, *Journal of Finance*, **54**, 901-933.

A Appendix

A.1 Experimental Instructions

In this experiment you will observe the time series of the financial returns of 2, 3 or 4 investment funds: fund A, B (or fund A, B, C, fund A, B, C, D) for 100 consecutive periods. You will have to make an investment decision by choosing one of these funds. Your payoff in the experiment will depend on the return of the fund chosen as your investment decision.

The actual returns of the funds will always be positive. The information you have when making your forecasts and investment decision consists of the actual returns of the funds in the recent past. After each 50 periods of predictions, we will show you the returns of new funds, which might not have the same pattern as the former ones.

Earnings

You may earn points for every period of the experiment. The earned points will be transformed to the payoff in Euros at the end of the experiment. **The number of points you get is the return from the fund you have chosen this period.**

At the end of the experiment you are paid 1 euro for each 20 points you earned during the experiment. In addition you will get a show-up fee of 5 euros. As an example, if for one period you choose between 3 funds, and your choice and the actual returns are as in the table below:

The Fund You choose	Profit of Fund A	Profit of Fund B	Profit of Fund C	Your Payoff
A	3	1.3	4.6	3

Then: you will get 3 points for this period, which is the actual return on fund A.

You earn 1 euro for each 20 points you make till the end of the experiment.

A.2 Estimation Result of Discrete Choice Models

A.2.1 Binary Choice

Phase	Type	α	β	log-likelihood	McFadden R-squared	MSE
T1_40	WN2	0.901 (0.090)	0.139 (0.033)	-367.402	0.024	0.0210
T6_80	WN2	0.430 (0.098)	0.293 (0.039)	-307.493	0.098	0.0299
T2_40	BH2	-0.197 (0.103)	0.378 (0.034)	-315.741	0.212	0.0809
T5_80	BH2	-0.314 (0.113)	0.438 (0.038)	-267.038	0.271	0.0391
T4_40	SI2	0.573 (0.144)	4.504 (0.374)	-200.018	0.482	0.0219
T3_80	SI2	0.381 (0.119)	3.293 (0.271)	-246.644	0.361	0.0140

Table 7: Aggregate estimation of binary choice models with one lag.

Phase	Type	β_0^A	β_1^A	β_2^A	β_3^A	β_4^A	log-likelihood	McFadden R^2	MSE
T1_40	WN2	0.901 (0.090)	0.139 (0.033)				-367.402	0.024	0.0210
T6_80	WN2	0.554 (0.107)	0.313 (0.041)	-0.153 (0.040)			-286.185	0.120	0.0066
T2_40	BH2	-0.271 (0.232)	0.892 (0.123)	-0.591 (0.097)	0.072 (0.101)	0.234 (0.065)	-202.661	0.428	0.0254
T5_80	BH2	-0.572 (0.289)	0.982 (0.138)	-0.671 (0.113)	0.199 (129)	-0.411 (0.098)	-161.371	0.498	0.0167
T4_40	SI2	0.591 (0.158)	5.830 (0.535)	-1.337 (365)			-175.433	0.521	0.0162
T3_80	SI2	0.381 (0.119)	3.293 (0.271)				-246.644	0.361	0.0140

Table 8: Aggregate estimation of binary choice models allowing for more lags.

A.2.2 Multiple Choice

Treatment		Coefficient	Std. Err.	T Statistic	p-value
BH3-Uexp	π_{t-1}^A	0.0056	0.0423	0.13	0.895
	π_{t-1}^B	0.3692	0.0165	2.24	0.025
	constant	-0.1163	0.2051	-0.57	0.571
	α	-0.255	0.0935	-2.73	0.006
	β	0.035	0.0164	2.13	0.033
BH3-Exp	π_{t-1}^A				
	π_{t-1}^B	0.0965	0.019	5.08	0
	constant	0.1295	0.0902	1.43	0.151
	α	0.3295	0.355	0.93	0.353
	β	-0.098	0.0955	-1.03	0.305
BH4-Exp	π_{t-1}^A	1.6978	0.1244	13.65	0
	π_{t-1}^B	1.6293	0.137	11.89	0
	π_{t-1}^C				
	constant	1.5952	0.1032	15.46	0
	α	1.5961	0.1033	15.44	0
	β	1.6734	0.1174	14.26	0
SI3-Uexp	π_{t-1}^A	0.4273	0.0528	8.09	0
	π_{t-1}^B	-0.6777	0.0238	-2.85	0.004
	constant	-2.4296	0.3165	-7.68	0
	α	-0.0725	0.1041	0.49	0.486
	β	0.0075	0.0196	0.38	0.705
SI4-Uexp	π_{t-1}^A	2.2317	0.1783	12.52	0
	π_{t-1}^B	2.5075	0.1848	13.57	0
	π_{t-1}^C	3.8102	0.3175	12	0
	constant	-0.4416	0.136	-3.25	0.001
	α	-0.1305	0.1085	-1.2	0.229
	β	2.5152	0.1549	16.23	0
SI4-Exp	π_{t-1}^A	4.0784	0.2657	15.35	0
	π_{t-1}^B	4.0972	0.2699	15.18	0
	π_{t-1}^C	4.7236	0.3586	13.17	0
	constant	-0.0739	0.1306	-0.57	0.571
	α	0.0500	0.1129	0.44	0.658
	β	4.1987	0.2439	17.21	0

Table 9: Aggregate estimation of multiple choice models.