Fee structure, return chasing and mutual fund choice: an experiment*

Mikhail Anufriev[†] Te Bao[‡] Angela Sutan[§] Jan Tuinstra[¶]

First Draft: 22 March 2012 This draft: 25 April 2015

Abstract

We present an experiment that investigates the effect of the fee structure and past returns on mutual fund choice. We find that subjects pay too little attention to the (periodic and small) operation expenses fee, but that the more salient front-end load is used as a commitment device and leads to lock-in into one of the funds. In addition we find that, even when subjects know that future returns are independent of past returns, these past returns are an important determinant of subjects' investment choices.

Keywords: Mutual fund choice, fee structure, experimental economics, return chasing

JEL Code: C91, G02, G11.

*We benefitted from stimulating discussions with John Duffy, Jacob Goeree, David Goldbaum, Cars Hommes, Alan Kirman, Marco van der Leij, Francois Koulischer, Malwina Mejer, Andreas Ortmann, Joep Sonnemans, John Wooders, Nan Yang and participants of the CEF 2012 Conference in Prague, the Experimental Finance conference 2013 in Tilburg and seminars at the Université libre de Bruxelles, the University of Amsterdam and the University of New South Wales. We are grateful to the financial support from the EU 7th framework collaborative project "POLHIA" Grant No. 225048 and the Research Priority Area Behavioral Economics of the University of Amsterdam. Mikhail Anufriev and Jan Tuinstra acknowledge financial support from the Australian Research Council through Discovery Project DP140103501.

[†]Economics Discipline Group, University of Technology Sydney, P.O. Box 123, Broadway, NSW2007, Sydney, Australia. E-mail: Mikhail.Anufriev@uts.edu.au

[‡]Department of Economics, Econometrics and Finance, University of Groningen, P.O. Box, 9700 AV, The Netherlands. E-mail: T.Bao@rug.nl

[§]Burgundy School of Business, 29, rue Sambin-21000 Dijon, France. E-mail: Angela.Sutan@escdijon.eu

[¶]Amsterdam School of Economics and CeNDEF, University of Amsterdam, Valckenierstraat 65-67, 1018XE, Amsterdam, The Netherlands. E-mail: J.Tuinstra@uva.nl

1 Introduction

Mutual funds are important investment vehicles: according to Investment Company Institute (2014) the value of assets invested in mutual funds worldwide was about 30 trillion US dollars in 2013, which is more than one third of world GDP. Moreover, mutual funds are particularly important for household finance. A total of 46.3% of US households own mutual funds, and the median value of mutual fund assets owned per household is $100,000$ US dollars. Given this large size of the mutual fund industry it is evident that understanding how investors choose between mutual funds is crucial for evaluating the performance of financial markets and might help in constructing appropriate regulatory policies.

This paper presents a laboratory experiment on sequential individual decision making under uncertainty that is aimed at shedding some light on mutual fund choice. In particular, we focus on two specific determinants of this decision problem. First, we study the role that the structure of the fee, charged by mutual funds, plays. More precisely, we consider two types of fees that are commonly used in practice: a *front-end load*, which is a fixed commission that has to be paid when an investor purchases shares of the fund, and an *operation expenses fee* or *management fee* that represents the costs for operating the fund and that, as opposed to the front-end load, needs to be paid by the shareholder periodically. Second, we investigate whether past returns of the mutual funds affect investors behavior, even under circumstances where it should be clear to these investors that they do not convey any additional information about future returns.

Much of the earlier empirical research in this area uses data from actual mutual funds, which introduces several endogeneity problems. The advantage of our laboratory experiment is that we have full control over the mutual fund return and fee structures. This allows us to test the effects of these two aspects more directly. In our experiment in each period subjects choose between investing their wealth in one of two experimental funds (A and B). The subjects know the return generating processes and fee structure of both funds, and observe past returns of both funds as well. However, they are explicitly informed that neither the past returns nor their own actions affect future returns. We construct three treatments, with no fees charged by fund A and the fee structure for fund B different across treatments. In particular, in treatment **N** fund B charges no fee, in treatment **O** it charges an operations expenses fee and in treatment **F** a front-end load, respectively Moreover, our experiment is designed in such a way that, although the decision problem is framed differently – with the fee in treatment **F** much more salient than that in treatment **O** – subjects in the three different treatments face essentially the same choice, since investing in fund B gives the same expected return (net of fees) in all three treatments. This expected return is higher than the expected return of fund A (which is also the same across treatments), implying that the optimal decision is to invest in fund B in every period. In particular, neither the fee structure, nor the past returns should have a significant effect upon investment choices of subjects.

However, we find that both the fee structure as well as past returns actually do have a substantial impact upon investment decisions. First, subjects choose fund B more frequently in treatment **O** than in treatment **N**. This suggests that subjects are driven more by gross returns (which are higher for fund B in treatment **O** than in treatment **N**) than by net returns and that they tend to ignore the, not very salient, operation expenses fee. Second, although the difference in average investments in fund B between treatments **F** and **O** is not significant, there is much more variation in individual investment decisions in the former treatment. In particular, the fraction of subjects that invest in fund B in (almost) every period is much higher in treatment **F**. These subjects understand they should pay the front-end load not more than once, and therefore get locked-in to the more profitable fund B. On the other hand, some subjects in treatment **F** switch between funds A and B more often and consequently end up with relatively low payoffs. Finally, for all treatments we find that past realized returns do have a substantial effect on investment decisions.

The impact of the fee structure on mutual fund choice has also been studied empirically. Barber, Odean, and Zheng (2005), for example, find that the flow of money into mutual funds is negatively correlated with the front-end load, but not correlated with the operation expenses fee. They argue that people pay more attention to the front-end load fee because it is more salient and transparent. Khorana and Servaes (2012), however, find that fund families that charge a front-end load have a significantly larger market share than funds charging operation expenses, which they explain by the fact that, much like as in our experiment, the front-end load may serve as a commitment device for investors, and thereby reduces search costs. Related to this, Chordia (1996) shows that funds may use front-end and back-end loads to discourage switching to another fund.[1] To our knowledge, only limited experimental research on the relation between fee structure and fund choice exists. An early contribution is Wilcox (2003), who shows that subjects pay too much attention to past performance and to the front-end load, and too little to the operation expenses fee – which is consistent with our findings. His experimental design differs from ours since he considers a fund that

[1]A separate literature focuses on whether the size of the fee is justified by the performance of the fund. Remarkably, Carhart (1997) and Gil-Bazo and Ruiz-Verdú (2009) find that the fees charged by funds are *negatively* related to returns. Khorana, Servaes, and Tufano (2009) provides a comprehensive investigation on mutual fund fees and finds a large dispersion of fees between different countries, which is difficult to explain by the difference in returns to the investors.

charges both a front-end load and an operation expenses fee. Moreover, his experiment conducts a one period cross-sectional comparison, while ours uses a multi-period setting to study investment over a longer horizon. Ehm and Weber (2013) employ a large scale survey (without payments to the subjects) to investigate how people choose between two hypothetical funds that only differ in the fee structure: one fund charges a performance fee (meaning that the fund charges a fraction of the gains, instead of a fraction of the total asset value) whereas the other charges an operation expenses fee. They find that people have a higher propensity to choose funds that charge a performance fee, which they explain by loss aversion.

There is substantial empirical evidence that investors on actual financial markets base their investment decisions on past performance of the funds, see e.g. Sirri and Tufano (1998), even though past performance may have limited predictive power on the future returns of the funds.[2] In particular, Choi, Laibson, and Madrian (2010) find that people rely strongly on the annualized past return of funds in making fund selection decisions even when such information is irrelevant. They design a field experiment where subjects choose between four funds that are based on the same index, and will therefore generate the same returns. In the manuals of the funds, however, they have different annualized past return rate records due to differences in the launching time, and the funds with a higher past return record charges a higher fee. Although people should ignore the past return information and select the fund with the lowest fee, many of them fail to do so. In an earlier paper by Choi, Laibson, Madrian, and Metrick (2009), the authors find that people's investment choices can be described by reinforcement learning: investors who experience rewarding outcomes from 401(k) savings tend to increase their savings level more than they optimally should. Our work is also related to Bloomfield and Hales (2002), who ask subjects to predict the next step of an earnings time series that follows a random walk. They find the subjects do not take the random walk as random, but divide the time series into "trend" and "mean reverting" regimes, and try to use the frequency of past earnings reversals to predict the likelihood of a future earnings reversal.

The remainder of the paper is organized as follows. Section 2 describes our experimental design. The experimental results are presented in Section 3 and

[2]There is a debate about whether there is persistency in the performance of mutual funds, for example because the fund manager is capable of consistently selecting high performing stocks. The general view is that such a persistency does not exist and that, after controlling for risk and trading costs, the typical manager is unable to consistently generate excess returns (see Carhart (1997)). In addition, Jain and Wu (2000) find that funds that advertise higher past returns attract more money, but do not perform significantly better than other funds in the periods following the advertisement. Hendricks, Patel, and Zeckhauser (1993) and Zheng (1999) find that there may be a "hot hand effect" in fund performance in the short run, but that in the long run there is no significant difference between funds that performed well in the recent past and other funds.

Section 4 concludes. Appendices A and B contain the experimental instructions.

2 Experimental Design

2.1 Summary Information

The experiment took place on December 8 and 9, 2011 at LESSAC, the experimental laboratory of the Burgundy School of Business in Dijon (France). In total 76 subjects participated in three treatments, with 22 subjects in treatment **N**, 19 in treatment **O** and 35 in treatment **F**.[3] All subjects were first year master students at the Burgundy School of Business, with no prior experience with laboratory experiments on a related topic. These students had two years of training in economics, statistics and mathematics before passing the exam to enter the school, and they took many other courses on business economics after entering the school.[4] Subjects could choose to have instructions in English or in French. The duration of the typical session was one and a half hours.

2.2 Task Design

The experiment is an individual choice experiment and divided in three blocks of fifteen periods each. At the beginning of each block the subject is given an initial wealth of $M_0 = 1000$ points. In each period t of that block the subject has to decide where to invest its accumulated wealth M_t: either it invests all of his/her wealth in fund A, or all in fund B, or he or she does not invest at all (note that the subject is not allowed to divide his or her wealth more evenly between the different funds). The wealth in the subsequent period of that block, M_{t+1}, is determined by the (exogenously given) stochastic return of the chosen fund (with $M_{t+1} = M_t$ if the subject decided not to invest his/her wealth in either fund A or fund B). The starting wealth is reset to $M_0 = 1000$ at the beginning of the second and third block.

At the end of the experiment, the subjects are paid according to their final wealth from one of the three blocks, where each block has the same probability of being chosen.[5] Appendix A provides the experimental instructions for this

[3]The treatments are discussed in detail in Section 2.4. The difference in the number of subjects per treatment is due to variations in the show-up rate.

[4]The students tend to be quite good. From a population of about 4000 students that participate in an entrance exam, the Burgundy School of Business has the right to select about 150, with grades between 14/20 and 17/20. Students with a grade higher than 17/20 go to HEC Paris.

[5]The experimenter throws a dice separately for each subject. The subject is then paid according to his/her final wealth in the first (second, third) block if the dice shows 1 or 2 (3 or 4, 5 or 6).

experiment for one of the treatments. Before the subjects start the experiment, they have to answer several control questions on paper in order to make sure that they understand the experiment. We start the experiment only when all subjects have answered all control questions correctly. The control questions and answers can be found in Appendix B.

2.3 Mutual Funds: Returns and Fees

Consider the (open-end) mutual fund X with a price per share of $P_{X,t}$ at time t. This price evolves according to $P_{X,t} = (1 + g_X + \epsilon_{X,t})\, P_{X,t-1}$, where $g_X > 0$ is a positive *growth constant* and $\{\epsilon_{X,t}\}$ is a white noise process, which can take on only two values: $\varepsilon > 0$ or $-\varepsilon$, with equal probability. The (gross) return of fund X then is given by

$$R_{X,t} = \frac{P_{X,t}}{P_{X,t-1}} = 1 + g_X + \epsilon_{X,t} = \begin{cases} 1 + g_X + \varepsilon & \text{with probability } \frac{1}{2} \\ 1 + g_X - \varepsilon & \text{with probability } \frac{1}{2} \end{cases} .$$

Because $\epsilon_{X,t}$ equals zero in expectation, the *expected* one-period return at the beginning of period t (that is, before $P_{X,t}$ is known) is given by $\mathrm{E}_t\left[R_{X,t}\right] = 1 + g_X$. More generally – and for now abstracting from any fees to be paid – the τ-period expected return at time t of investing one unit of money is given by

$$\mathrm{E}_t\left[R^{(\tau)}_{X,t}\right] = \mathrm{E}_t\left[\frac{P_{X,t-1+\tau}}{P_{X,t-1}}\right] = \sum_{s=0}^{\tau} \binom{\tau}{s} \frac{1}{2^\tau} \left(1 + g_X + \varepsilon\right)^s \left(1 + g_X - \varepsilon\right)^{\tau - s} =$$
$$= \left(1 + g_X\right)^\tau .$$

For the experimental design we consider two funds, $X = A, B$, with growth constants $g_A = a$ and $g_B = b$ respectively. The random components $\epsilon_{A,t}$ and $\epsilon_{B,t}$ are independent but identically distributed (in particular, the absolute size of the random components is equal to ε for both funds). Moreover, we impose that $\varepsilon < a < b$. The first inequality means that, independent of the realizations of $\epsilon_{A,t}$ and $\epsilon_{B,t}$, the prices of shares of both funds are monotonically increasing over time. The second inequality implies that *expected* τ-period returns for fund B are higher than for fund A (for any $\tau \geq 1$). In addition, by requiring that $a + \varepsilon > b - \varepsilon$ or, equivalently, that $\varepsilon > \frac{1}{2}(b - a)$, there is a positive probability that the *realized* return of fund A is higher than that of fund B.[6] A rational investor that knows the data generating mechanism (as is the case in our experiment) will always choose investing in fund B over investing in fund A, since past realized returns do not

[6] It would happen in the period, when the price of fund B incurred a negative shock, and the price of fund A incurred a positive shock.

convey any additional information about future returns and expected returns are always higher for fund B than for fund A. However, investors that do respond to past realized returns might now and then switch to fund A.

We consider two types of fees: an *operations expenses fee* and a *front-end load.* The first type of fee, sometimes also referred to as a *management fee*, is a periodic payment that represents the costs for running the mutual fund and providing service to its shareholders. It corresponds to a fraction γ_X of the investment to be paid each period as a fee. That is, the operation expenses fee in period t is $\gamma_X P_{X,t-1}$ per share. The (one-period) return for fund X, net of this fee, then becomes $R_{\mathbf{O},X,t} = 1 + g_X - \gamma_X + \epsilon_{X,t} = 1 + g'_X + \epsilon_{X,t}$, with $g'_X = g_X - \gamma_X$. The expected τ-period return is

$$\mathrm{E}_t\left[R^{(\tau)}_{\mathbf{O},X,t}\right] = (1 + g_X - \gamma_X)^\tau \ .$$

Alternatively, the investor may be charged with a purchase commission, or a so-called *front-end load.* That is, the investor pays a fixed percentage F_X of his investment M_t as a commission when he invests in fund X. With the remainder of his investment, $(1 - F_X)\,M_t$, shares of the mutual fund are purchased. The front-end load only has to be paid upon purchasing the shares. However, if the investor withdraws his money from the fund and wants to re-invest in the fund at a later stage he has to pay the front-end load again. The expected τ-period return from investing in mutual fund X at time t follows as

$$\mathrm{E}_t\left[R^{(\tau)}_{\mathbf{F},X,t}\right] = (1 - F_X)\,\mathrm{E}_t\left[R^{(\tau)}_{X,t}\right] = (1 - F_X)\,(1 + g_X)^\tau \ .$$

2.4 Treatments

The experimental design features three treatments: **N**, **O** and **F**. Subjects only participate in one of the treatments. In each of these treatments subjects could increase their wealth by investing in one of the two funds, A or B, as described in Section 2.2. No fees are required for investing in fund A, but for two of the three treatments a fee is charged when investing in fund B. Table 1 summarizes the design. In addition, we take $\varepsilon = 0.02$.

In the *baseline treatment*, **treatment N**, none of the two funds requires a fee and we set the growth constants to $a = 3\%$ and $b = 4\%$. The optimal decision is to invest in fund B in every period (although the realized return of investing in fund A will be higher than that of investing in fund B in periods t when $\epsilon_{A,t} = \varepsilon$ and $\epsilon_{B,t} = -\varepsilon$). Investing in fund B in every period gives an expected (net) return of about 73%, whereas the net expected return of only investing in fund A is around

Treatment	Fund A		Fund B			
	g_A	$\mathrm{E}_1\left[R_{A,1}^{(14)}\right]-1$	g_B	γ_B	F_B	$\mathrm{E}_1\left[R_{B,1}^{(14)}\right]-1$
N	3%	51.26%	4%	–	–	73.17%
O	3%	51.26%	5%	1%	–	73.17%
F	3%	51.26%	5%	–	13%	72.25%

Table 1: Experimental design.

51%.[7]

For the *operating expenses treatment*, **treatment O**, the growth constants are equal to $a = 3\%$ and $b = 5\%$, but there is an operations expenses fee for fund B equal to $\gamma = 1\%$. Effectively, therefore, expected (and realized) returns to investing in fund B are exactly the same in treatments **O** and **N**, and higher than the expected returns of investing in fund A.

In the third and final treatment, **treatment F**, we impose a *front-end load* on fund B. We take $a = 3\%$ and $b = 5\%$ again and choose a front-end load of $F = 13\%$. For this value of the front-end load the expected return of investing in mutual fund B from the beginning of the block is (roughly) equal to the expected return of investing in fund B for the other two treatments.[8] Note however that, where for treatments **N** and **O** it is always optimal to switch from investing in fund A to investing in fund B, for treatment **F** this is only worthwhile if enough periods remain to 'earn back' the front-end load. In particular, switching from fund A to fund B later than period 7 decreases expected returns.[9]

For each treatment the chosen parameters (a, b, γ and F) are the same in each block, and for each block we generate time series of prices for fund A and fund B, respectively. The only difference between blocks in the same treatment is that different seeds for generating the white noise process are used. Moreover, we use the same realization of the shocks $\epsilon_{A,t}$ and $\epsilon_{B,t}$ in the three treatments. Therefore

[7]Note that, although subjects need to make an investment decision for 15 consecutive periods there will only be 14 return rates. It can be easily checked that investing in fund B for all periods gives at least a net return of 32.0% (when all price shocks are negative) and at most a net return of 126.1% (when all these price shocks are positive). The corresponding numbers for fund A are 15.0% and 98.0%, respectively. Moreover, the return on always investing in fund A is going to be higher than the return on always investing in fund B only if fund A experiences at least four more positive price shocks than fund A does.

[8]To be precise: for the expected returns of fund B to be *exactly* the same in all three treatments, constant F has to satisfy $(1-F)\,(1+b)^T = (1+a)^T$. This equation, for $a = 0.03$, $b = 0.05$ and $T = 14$, gives $F^* \approx 0.1254$. We selected $F = 0.13$, because this is the closest integer (in percentage points) to F^*. Note that the difference between expected returns for fund B in the different treatments is very small.

[9]It can be easily checked that $0.87 \cdot 1.05^t \leq 1.03^t$ for all $t \geq t^* \approx 7.2414$, and choosing fund B (and paying the front-end load) can only be profitable (in expectation) when at least 8 periods remain.

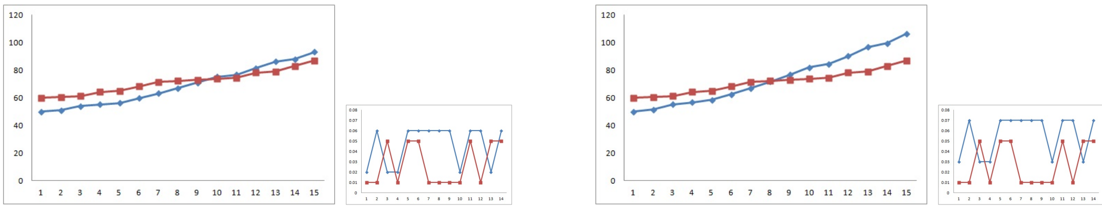

(a) **Block 1.**

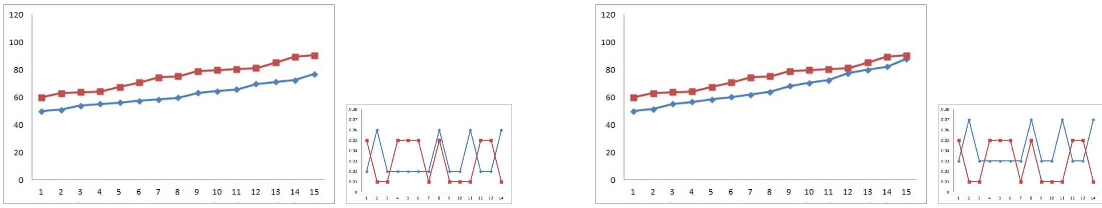

(b) **Block 2.**

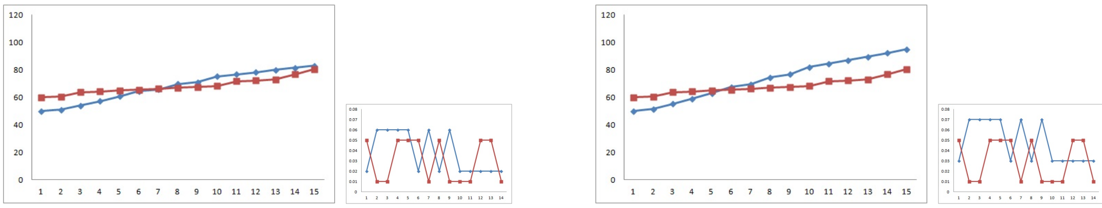

(c) **Block 3.**

Figure 1: Prices and returns in treatments **N** (left) and **O** and **F** (right). Prices and returns for fund A are shown by squares. Prices and returns for fund B are shown by circles. The expected return for fund A is 3% in all treatments, and the expected net return for fund B when a participant always chooses B is 4%.

any difference we observe between treatments can be attributed to differences in the fee structure. Figure 1 shows the generated time series of prices and returns for the three different blocks.[10]

Each subject has full information about the price generating mechanisms. As a given block of the experiment evolves, subjects are shown a table and a graph with the time series of past prices ($P_{A,t}$ and $P_{B,t}$). The subjects are also shown the value of their portfolio in each period. Figure 2 provides an example of a typical experimental screen.

2.5 Hypotheses

We designed our experiment in such a way that, although the decision problem is framed differently, subjects face essentially the same choice. In all three treatments fund B gives higher expected payoffs than fund A. In addition, the price generating mechanism is such that past performance does not influence future

[10]Realized returns of investing for 15 periods in fund A (fund B) are 45.1% (86.2%) in the first block, 50.9% (53.9%) in the second block and 34.3% (66.2%) in the third block, respectively.

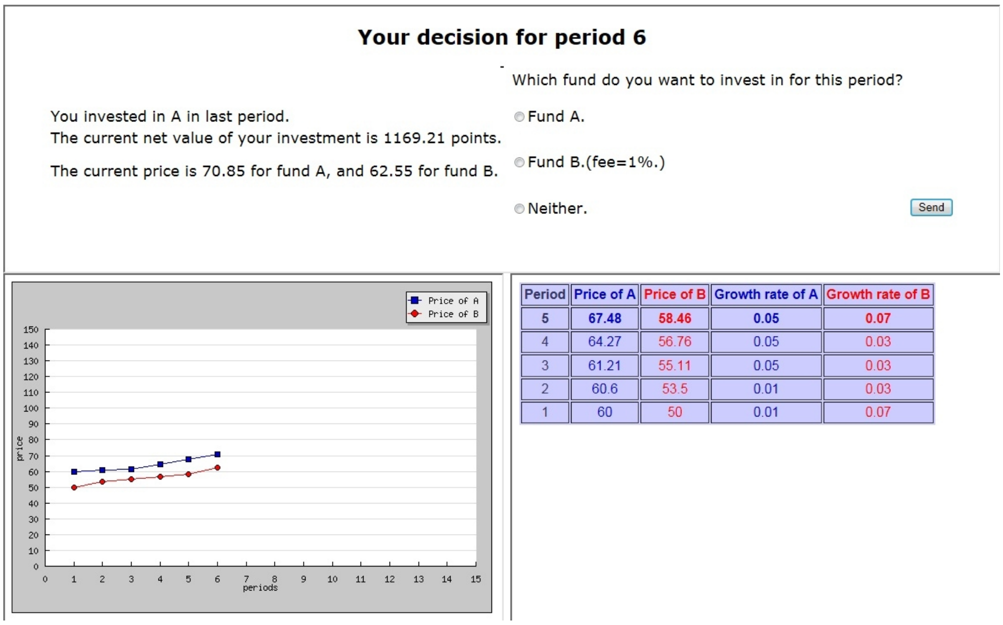

Figure 2: A typical experimental screen. The subjects make a choice between investing in fund A, B or neither of them in the upper part of the page. They can refer to the past prices of the funds in the graph in lower left part, and the table in the lower right part. The prices of fund A are shown by squares, and the prices of fund B are shown by diamonds.

performance of the funds, which is known by the subjects. Therefore one would predict that neither the fee structure, nor past realized returns of the funds will have a significant impact upon the investment choices of the subjects. This leads to the following set of hypotheses, which we will test in Section 3.

First, there should not be a significant difference between treatments **N** and **O**. Although the choice problem is framed differently in the two treatments, expected and realized returns on both funds are the same in these treatments.

Hypothesis 1. *There is no significant difference in subjects' behavior between treatments **O** and **N**.*

If we do find a statistical significant difference between the treatments, this can be attributed to the way the returns are framed in the two different treatment (either as an expected return of 4%, or as an expected return of 5% minus a fee of 1%).

Similarly, we do not expect a difference between treatments **O** and **F**, since – again – the expected returns to choosing fund B in every period is about the same for treatment **O** as it is for treatment **F**.

Hypothesis 2. *There is no significant difference in subjects' behavior between treatments **F** and **O**.*

Contrary to the comparison between treatments **O** and **N**, there are explanations other than framing for a possible significant difference between treatments **F** and **O**. First, for treatments **O** and **N** it is straightforward to understand which of the two funds generates a higher expected return – since 4% is clearly higher than 3%. However, in treatment **F** the comparison between the two funds requires a non-trivial computation, which may lead to uncertainty with the subject about how to evaluate fund B. Moreover, the front-end load in treatment **F** is much higher, and therefore much more salient than the operating expenses in treatment **O**. These two effects might explain *underinvestment* in fund B in treatment **F**. On the other hand, the fact that the front-end load has to be paid every time that the subject starts to invest in fund B (and only then) implies that switching back and forth between fund A and fund B is much more costly in treatment **F** than it is in treatment **O**. In that sense, the front-end load may serve as a commitment device and force the subject to exert more effort into thinking about the investment decision at the start of the experiment, or at the start of a new block. This might increase the fraction of subjects choosing fund B in treatment **F**.

Our last hypothesis focuses on a so-called 'return chasing behavior'.

Hypothesis 3. *For treatments **N** and **O** the fraction of choices for fund B in any particular period is uncorrelated with the realized returns of funds A and B in the previous period.*

If this hypothesis is rejected, subjects are triggered by realized returns, although these realized returns do not convey additional information.

3 Experimental Results

3.1 Investment Decisions

Table 2 reports, for each of the three treatments and averaged over periods within a block, the fraction of subjects choosing fund A, fund B or none of the two, respectively.

From Table 2 we see that it is quite rare, in particular for treatments **N** and **O**, that a subject does not in invest in either of the two funds, as was to be expected because at least fund A always generates a strictly positive return. More surprising is the finding that in all three treatments the fraction of choices for fund A is substantial and ranges from about 20% (block 3 in treatment **O**) to even slightly more than 50% (block 2 in treatment **O**). In addition, although the fraction of choices for A (in treatments **N** and **O**) is lower in block 3 than it is in block 1 – suggesting that with experience subjects learn to make better decisions

Treatment	Periods	Fraction of Choosing		
		A	B	**Neither**
N	Period 1	45.45%	54.55%	0.00%
	Block 1	41.52%	58.48%	0.00%
	Block 2	48.79%	50.91%	0.30%
	Block 3	27.27%	72.73%	0.00%
	Average	39.19%	60.71%	0.10%
O	Period 1	68.42%	31.58%	0.00%
	Block 1	32.63%	67.02%	0.35%
	Block 2	50.88%	48.77%	0.35%
	Block 3	20.35%	79.65%	0.00%
	Average	34.62%	65.15%	0.23%
F	Period 1	37.14%	62.86%	0.00%
	Block 1	24.38%	74.48%	1.14%
	Block 2	37.90%	60.76%	1.33%
	Block 3	24.76%	74.86%	0.38%
	Average	29.01%	70.03%	0.95%

Table 2: Fraction of participants/times choosing different options in each treatment.

– the fraction of choices for fund A is highest in block 2, for all treatments. A possible explanation for this lies in the fact that the difference between the realized returns of the two funds is quite small in block 2. In particular, in block 2 fund A experiences three more positive price shocks than fund B. Moreover, in block 2 it happens more often (six times) that the realized return of fund B is less than that of fund A (which happens two and three times in blocks 1 and 3, respectively). Although, these realized returns should not have an effect on investment decisions, apparently they do. We will get back to this in Section 3.4.

Now let us compare treatments **N** and **O**. Both over all blocks and separately for blocks 1 and 3, there is a higher fraction of choices for fund B in treatment **O** than in treatment **N**. For block 2 it is the other way around, although the difference is quite small for that block. The difference between the treatments is statistically significant at the 5% level according to the Mann–Whitney–Wilcoxon test. This gives us our first result.

Result 1. *We reject Hypothesis 1 and find that in treatment **O** there is a significant higher fraction of choices for fund B than in treatment **N**.*

This result suggests that subjects pay close attention to the (gross) expected return of 5%, but (partially) ignore the operations expenses fee of 1% that they should subtract from this expected return.

Secondly, we can compare treatments **F** and **O**. From Table 2 we see that the fraction of choices for fund B is higher for treatment **F** than for treatment **O** in the first two blocks, but lower in the third block. Although, averaging over all three blocks, the fraction of choices for fund B is higher for treatment **F**, the difference is not statistically significant at the 5% level (Mann-Whitney-Wilcoxon test). We, therefore, find that making fee much more salient has a limited influence on *aggregate* investment behavior. In our next step we investigate if it has an effect on *individual* investment behavior. Table 2 shows that the fraction of choices for both fund A and fund B are substantial. The question arises whether this is because most subjects switch regularly between the two funds, or because some subjects almost always choose fund A, whereas others typically choose fund B, and whether this depends upon the treatment. Figure 3 shows a histogram of the frequencies of individual choices for fund B, for each of the treatments. The histograms for treatments **N** and **O** are quite similar. All subjects from treatment **O** and 90% of the subjects from treatment **N** choose fund B at least 40% of the time. For both treatments most subjects choose fund B around 60% of the time, with only about 10% of the subjects in each treatment choosing fund B more than 90% of the time. In other words: there is some heterogeneity in individual choices, but this heterogeneity is mild. In particular, there are not many subjects that (almost) always choose fund B.

For treatment **F** the histogram looks completely different, with about 40% of the subjects choosing fund B for more than 90% of the time, and also a substantial number of subjects that rarely invest in fund B. Given the substantial cost of switching, represented by the front-end load, this makes intuitive sense. Assuming that the vast majority of the subjects in treatment **F** do not pay the front-end load more than once this suggests that most subjects choose fund B in one of the first couple of periods, and stick with that decision, and that there is a limited number of subjects choosing fund B for a small number of periods. Clearly, although on the aggregate level choice behavior in treatment **F** and treatment **O** is not significantly different, on the individual level there are substantial differences. A Kolmogorov-Smirnov test, at the 5% level, cannot reject the hypothesis that the distribution functions of individual frequencies of choices of fund B are the same for treatments **N** and **O**, but rejects the hypothesis that the distribution functions are the same for treatments **N** and **F** and the hypothesis that they are the same for treatments **O** and **F**.

Based upon the above discussion, which suggests that on the aggregate level there is no difference between treatments **O** and **F**, but on the individual level there clearly is, we formulate the following:

Result 2. *We reject Hypothesis 2: there is a significant difference, at the individual level, of behavior of subjects in treatments* ***O*** *and* ***F****.*

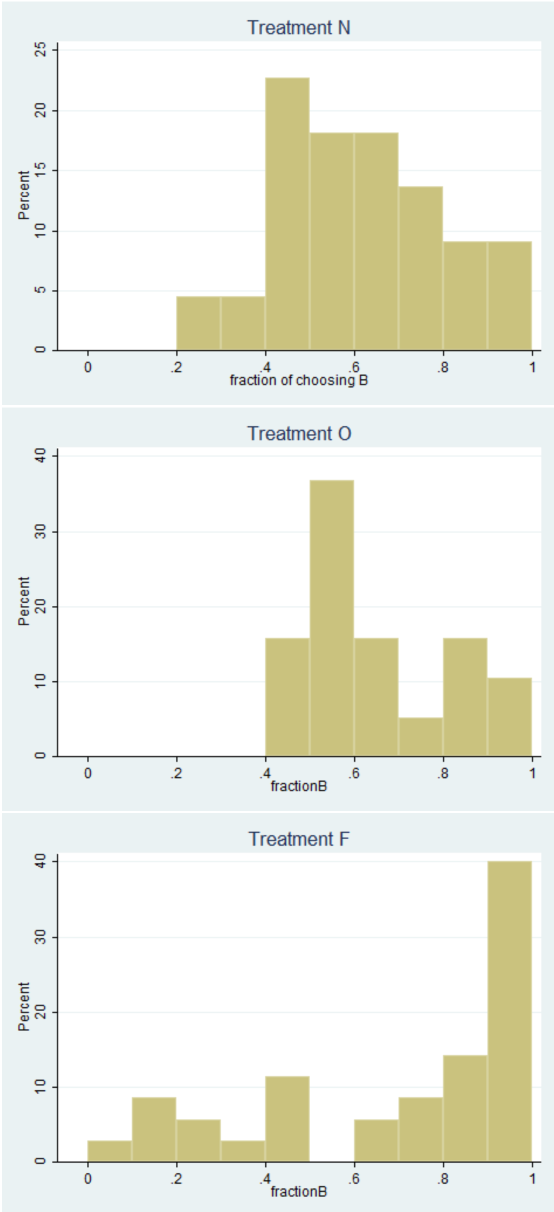

Figure 3: The histogram of individual choices of fund B. The horizontal axis represents the percentage of choice of fund B by each subject, and the vertical axis represents the percentage of subjects who choose fund B at this frequency.

3.2 Lock-in and switching between funds

We propose two explanations for the substantial difference between individual choices in treatment **F** and the other two treatments, illustrated in Figure 3. First, in treatment **F** most subjects understand that they should not pay the front-end load more than once, and moreover, that they need to stay with fund B for enough periods to at least recover the front-end load. In contrast, in treatments **N** and **O** the cost of switching between funds is small, and therefore subjects may want to explore both funds in those treatments. Second, given that subjects understand

	Treatment		
Periods	**N**	**O**	**F**
Block 1	0.28	0.20	0.22
Block 2	0.32	0.30	0.18
Block 3	0.27	0.20	0.12
All Blocks	0.29	0.23	0.17

Table 3: Frequency of switching per subject per period in different treatments and blocks.

that exploration and mistakes are prohibitively costly in treatment **F**, they may think better about the optimal investment choice, and conclude, already at the start of the experiment, that they should choose fund B in every period.

To investigate this issue further we determine, for each treatment and each block, the average fraction of times subjects switch between funds. Table 3 shows the results. In general, most switches occur in treatment **N**, and the fewest occur in treatment **F**, although the number of switches in treatment **F** is still quite high. In particular, the maximum number of times a subject in treatment F can switch, without paying the front-end load more than once, is two times, implying a frequency of 14.3%. Notably the frequency of switches in blocks 1 and 2 of treatment **F** are higher, implying that at least some subjects in these blocks paid the front-end load more than once, which clearly is suboptimal.[11] A Mann-Whitney-Wilcoxon test indicates that the difference between the frequency of switches across treatments is significant at the 5% level. Within each treatment, the difference between the number of switches across different blocks is not significant, except for treatment **F**, where the frequency of switching is significantly smaller in the third block compared to the first and second block. Fig. 4 shows the empirical cumulative distribution function of the number of switches by individual in each treatment. Consistent with the analysis above, the number of switches per individual is generally smallest in treatment **F**, and largest in treatment **N**.

We delve a bit deeper into this issue and investigate how often subjects in treatment **F** paid the front-end load, see Table 4. For reference we also include how often subjects in treatments **N** and **O** started with investing in fund B (although there was no explicit switching cost in these treatments for starting to invest in that fund, as there was in treatment **F**). Notice that there are 22, 19 and 35 subjects in the three treatments, meaning that we have observations from 66, 57 and 105 'individual blocks', respectively. Whereas the number of times subjects start to invest in fund B in treatments **N** and **O** is relatively evenly distributed

[11] In particular, the highest expected net return one can get by paying the front-end load twice is equal to 47.0%, which is lower than the expected return of investing in fund A in every period (which is 51.3%).

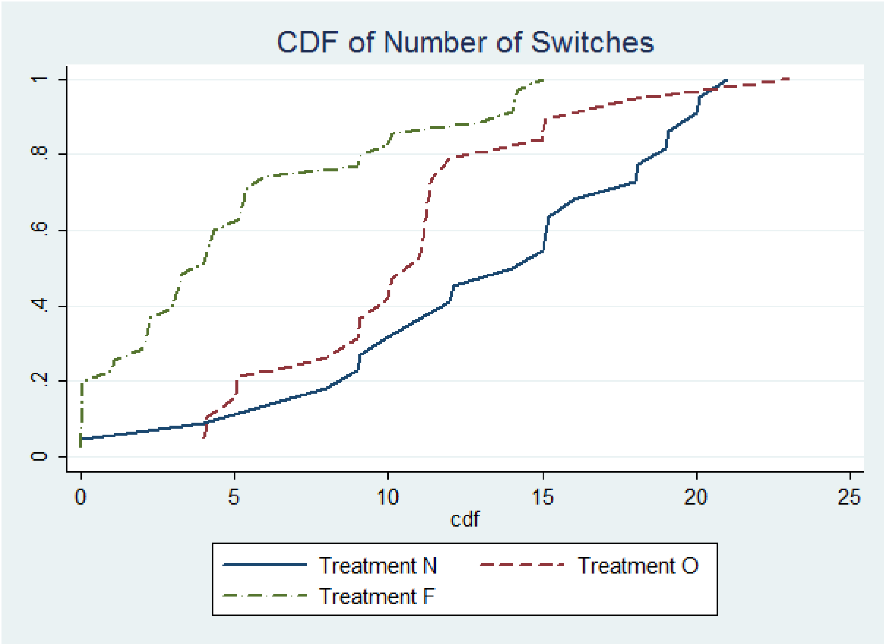

Figure 4: The empirical cumulative distribution function of total number of switches by each individual. The horizontal axis represents the number of switches, and the vertical axis is in terms of percentage. The solid line represents treatment **N**, the dashed line represents treatment **O**, and the dash-dotted line represents treatment **F**.

between 1, 2, 3 and 4 times, for a vast majority (70%) of the 105 individual blocks in treatment **F**, the subject only invests in fund B once, which is in accordance with the lock-in argument given above. Note that in almost one quarter of the blocks in treatment **F** the subject begins to invest in B at least twice, which is clearly suboptimal. Table 4 also presents, for treatment **F**, the results separately for each of the three blocks. From these last three rows it is clear that subjects *learn* to make better decisions, in particular from block 2 to block 3. In that last block there are only three out of 35 subjects that invest in fund B more than once. In addition (like in block 2) three subjects are deterred from investing in fund B at all, possibly due to a disappointing experience with that fund in one of the earlier blocks, or because of risk aversion and a misunderstanding or uncertainty about the size of the front-end load relative to the returns of the different funds.[12]

To repeat, a substantial number of subjects in treatment **F**, therefore, invest exactly once in fund B (from 70% of the subjects in all blocks, to 83% in block 3). A second question is *how long* these subjects invest in fund B. Expected returns are maximized when the subject invests in fund B for all 15 periods. Above we have already determined that they should invest for at least eight periods in

[12]The reason that there does not seem to be much learning between block 1 and block 2 in treatment **F** might be that in block 2 fund A is performing relatively well and outperforms B in 6 periods – also see the discussion following Table 2.

Treatment	Block	Periods of choosing B after investing in A					
		0	1	2	3	4	5
N	All	0 (0%)	12 (18%)	17 (25%)	24 (36%)	10 (15%)	3 (5%)
O	All	0 (0%)	21 (37%)	18 (32%)	10 (18%)	7 (12%)	1 (2%)
F	All	6 (6%)	73 (70%)	14 (13%)	7 (7%)	4 (4%)	0 (0%)
F	Block 1	0	22	4	5	3	0
	Block 2	3	22	8	1	1	0
	Block 3	3	29	2	1	0	0

Table 4: The number of times subjects started with fund B or switched from fund A to fund B (and, for treatment **F**, have to pay the front-end load) in each treatment. Numbers in brackets show the frequency for the corresponding number among all 'individual blocks' in a given treatment.

fund B to expect to recover the front-end load. It turns out that of the 73 of single investments in fund B in treatment **F**, 43 investments (i.e., 59% of the 73 investments, and 41% of all individual blocks) lasted for the full block of 15 periods and 53 investments (i.e., 73%) lasted for at least 12 periods. Of the 73 single investments only 10 (14%) lasted shorter than eight periods (and would therefore not correspond to *a priori* profitable investments). As a comparison: in treatment **N** only in six of the 66 blocks (9%) and in treatment **O** only in five of the 57 blocks (also 9%), the subject invested in fund B for all 15 periods.

Summarizing: the front-end load locks a substantial fraction of the subjects into the choice for fund B, which in the end might give higher payoffs for those subjects (see the discussion on earnings in the next Section 3.3). A reason for subjects to switch relatively often in treatments **N** and **O** may also be that, even if a subject understands that choosing B in every period is the optimal decision, he/she is still curious about what happens if A is chosen instead, which is not very costly in these treatments. This is related to the result by Blume and Ortmann (2007), who find that subjects may feel curious about alternative actions, and deviate from the efficient equilibrium even after they have played it for a long time.

3.3 Earnings

The analysis from Section 3.2 suggests that the variation of subjects' performance in treatment **F** will be much higher than in the other two treatments. Those subjects in treatment **F** that understand that they should pay the front-end load not more than once will often get locked into fund B, resulting in high payoffs. On the other hand, it will be difficult for subjects that pay the front-end load more

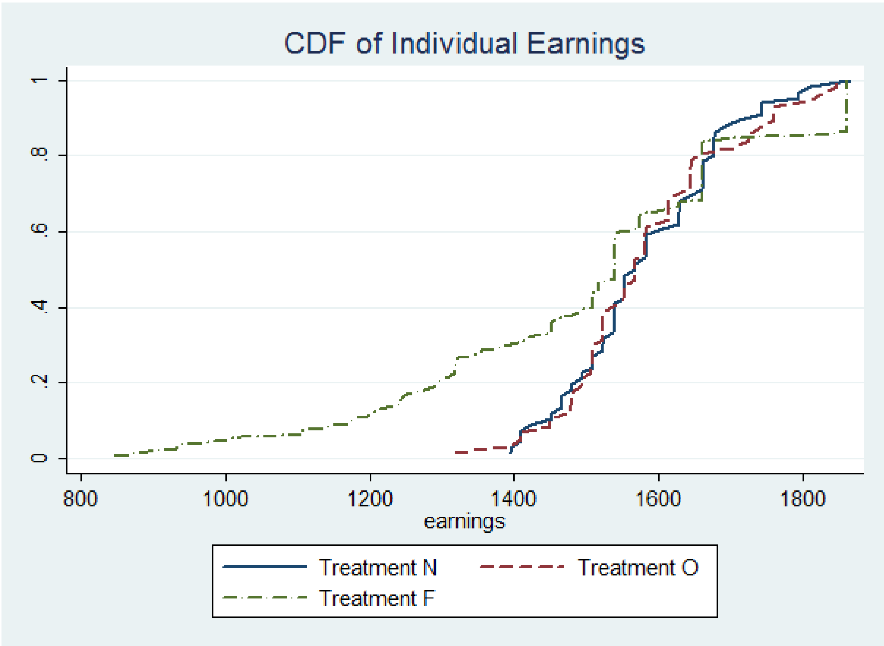

Figure 5: The empirical cumulative distribution function of individual earnings in different treatments, the earnings are in terms of points. Treatment **N** is shown by the solid line, treatment **O** is shown by the dashed line, and treatment **F** is shown by the dash-dotted line.

than once to obtain a good return on their investments.[13]

Figure 5 confirms this conjecture. It shows the empirical cumulative distribution function of individual earnings in each block in the three different treatments. There is clearly much more variation in earnings in treatment **F** than in treatments **N** and **O**, with both the share of low and high payoffs being higher in treatment **F** than in the other two treatments.[14] A Kolmogorov-Smirnov test indeed shows that the distributions for treatments **N** and **O** are not significantly different, but that the distribution for treatment **F** is significantly different both from that of treatment **N** as well as from that of treatment **O**.

We will now go into a bit more detail about the earnings. First, we describe in Table 5, earnings for the different funds in the different blocks under different

[13]In fact, the highest expected payoff for paying the front-end load fee twice – by investing for only one period in fund A, in between two longer investment runs in fund B – equals 1470 points, which is lower than the expected number of points (1512) from only investing in fund A. Moreover, the expected number of points when going back and forth between fund A and B as much as possible (meaning that the front-end load is paid seven times) is around 650 points, whereas in the other two treatments the payoffs can never be lower than 1000 points (which happens when a subject decides never to invest in any of the two available funds).

[14]On the one hand, both in treatment **N** and treatment **O** in only two cases individual earnings in a block are smaller than 1400 points (corresponding to 3.0% and 3.5% of the 66 and 57 observations, respectively), whereas in treatment **F** there are 31 such cases (29.5% of 105 observations). On the other hand, in treatment **F** there are 16 cases (15.2%) where individual earnings in a block are higher than 1800 points, whereas these numbers for treatments **N** and **O** are much smaller, i.e., 2 (3.0%) and 4 (7.0%) observations, respectively.

Block	Fund A (ex ante)	Fund B (ex ante)	Fund A (ex post)	Fund B (ex post)	Min.	Max.	Return chasing
1	1512.6	1731.7	1451.1	1862.3	1369.4	1976.7	1829.6
2	1512.6	1731.7	1508.6	1539.0	1267.8	1831.3	1495.0
3	1512.6	1731.7	1342.7	1662.0	1230.8	1813.0	1694.1

Table 5: Payoffs for each block of 15 periods for different scenarios (for treatments **N** and **O**). The first four columns show the expected and realized number of points for always choosing fund A or fund B, respectively. The next two columns show the minimum and maximum number of points that can be earned by always choosing the worst or best fund, and the last column gives the number of points for a subject that always chooses to invest in the fund that had the highest return in the previous period.

scenarios. The first four columns give the (*ex ante*) expected number of points of always investing in fund A or always investing in fund B, as well as the (*ex post*) realized number of points of always investing in fund A or B.[15] Obviously, the expected number of experimental points are the same for each block, but the realized number of points are different, with the differences between funds larger than expected in blocks 1 and 3 and (much) smaller than expected in block 2.[16] We have also added the minimum and maximum number of points a subject earns when always choosing (by sheer good or bad luck) that fund that has the lowest, respectively highest realized return in that period.[17]

Finally, the last column of Table 5 represents the payoffs for the hypothetical case where a subject is "chasing returns", that is, when the subject invests in that fund that had the highest realized return in the previous period. Clearly, because returns are uncorrelated (and subjects are informed about this) and fund B has higher expected returns than fund A, such an investment strategy will not be optimal. However, in any period there is a 25% probability that the return on fund A is higher than that on fund B and, although this provides no information about future return differences, it may induce subjects to switch to (or remain with) fund A. The last column in Table 5 reveals that the difference between

[15] The payoffs in Table 5 reflect those in treatments **N** and **O**: due to the different fee structure the payoffs in treatment **F** for investing only in fund B are marginally different.

[16] In blocks 1 and 3 the number of positive shocks for fund B is higher than for fund A (9 versus 6, and 6 versus 4, respectively), whereas in block 2 it is the other way around, with 7 positive shocks for fund A and only 4 for fund B.

[17] Note that, as discussed above, the minimum possible earnings in treatment **F** are much smaller, because a subject may decide to pay the front-end load more than once. In fact for 5 out of the 105 blocks in treatment **F** (i.e., 4.8%) payoffs are lower than 1000 points (with the lowest being 846 points in block 2) and 21 out of 105 (20.0%) lead to payoffs lower than the minimum payoffs possible in the relevant block, whereas this – by construction – can only happen in one of the other two treatments if a subject chooses not to invest at all for at least one period.

always investing in fund B and chasing returns is limited in blocks 1 and 3. This is due to the fact that there are only a few periods in these blocks in which fund A does better than fund B to begin with.[18] However, the difference in block 2 is much more pronounced. Remarkably, in that case return chasing leads to fewer points than choosing the inferior fund A in every period.

Comparing this with the experimental data we find that in blocks 1 and 3 of treatments **N** and **O** (because there is an explicit cost of switching from fund A to fund B in treatment **F**, it makes less sense to investigate return chasing in that treatment) there is only one participant earning less in a block than what he/she would earn by only investing in fund A. In block 2 of treatments **N** and **O**, however, more than half of the subjects (10 out of 22 and 11 out of 18, respectively) earn less than what they would earn by only investing in fund A – this suggest that return chasing may help in describing subject behavior. We will get back to this point in Section 3.4.

Table 6 shows for each treatment and each block the average and median number of points obtained by the subjects. Because, as can be seen from Table 5, realized payoffs are different for the different blocks, we also added an efficiency measure (the numbers between brackets), defined as the ratio of the (average or median) net return obtained by subjects over the net return obtained by always investing in fund B (the *ex ante* optimal decision), in order to facilitate comparison between the blocks.

From Table 6 we see that, although people choose fund B with a higher frequency in treatment **F**, both the average and median earnings (apart from the median in block 3) in this treatment are still the lowest. The earnings in treatment **F** are significantly lower than in the other two treatments in the second block at the 5% level according to a Mann-Whitney-Wilcoxon test, but there are no significant differences in earnings between treatments in the other blocks. The difference between the median earnings is smaller than the difference in average earnings across treatments. This suggests that the low average earnings in treatment **F** are mainly caused by some participants who often choose A or pay the front-end load too often. To investigate this for each block of treatment **F** we separated the subjects who paid the front-end load not more than once from those that paid it at least twice. Average and median earnings for these two separate groups are given in the last two columns of Table 6, with $\mathbf{F}^1$ (resp. $\mathbf{F}^2$) referring to the subjects paying the front-end load at most once (resp. at least twice). From

[18]Return of fund A is larger than the return of fund B twice in block 1, six times in block 2 and three times in block 3. In addition, a positive return difference between funds A and B is followed by a negative return difference for both cases in block 1, for three of the six cases in block 2 and for one of the two relevant cases in block 3 (the other time in block 3 where fund A had a higher return was in the last period).

Treatment	**N**	**O**	**F**	$\mathbf{F}^1$	$\mathbf{F}^2$
			Mean		
Block 1	1653.86	1682.81	1565.55	1781.37	1200.33
	(75.8%)	(79.2%)	(65.6%)	(90.6%)	(23.2%)
Block 2	1503.22	1493.52	1385.09	1464.53	1186.46
	(93.4%)	(91.6%)	(71.5%)	(86.2%)	(34.6%)
Block 3	1593.08	1579.89	1553.51	1584.74	1220.33
	(89.6%)	(87.6%)	(83.6%)	(88.3%)	(33.3%)
All Blocks	4750.17	4756.22	4504.15	4804.35	3591.93
	(84.8%)	(85.1%)	(73.0%)	(87.5%)	(28.7%)
			Median		
Block 1	1637.80	1708.93	1574.69	1862.29	1244.01
	(74.0%)	(82.2%)	(66.7%)	(100.0%)	(28.3%)
Block 2	1516.51	1493.53	1481.18	1539.28	1230.40
	(95.8%)	(91.6%)	(89.8%)	(100.1%)	(42.8%)
Block 3	1606.76	1567.78	1661.15	1661.15	1175.41
	(91.7%)	(85.8%)	(99.9%)	(99.9%)	(26.5%)
All Blocks	4715.58	4767.36	4659.05	–	–
	(83.1%)	(85.7%)	(80.4%)		

Table 6: Average individual earnings in different treatments and blocks. The earnings are in terms of points. Column $\mathbf{F}^1$ ($\mathbf{F}^2$) refers to the average/median over the subject in treatment **F** that pay the front-end load at most once (at least twice). The numbers in brackets are the ratio of the (average or median) realized net eturn over the realized net return when choosing fund B in every period.

these two columns it is clear that in particular those subjects in treatment **F** that understand that the front-end load should not be paid more than once perform quite well (for each block the median earnings of these subjects gives the same payoffs as investing in fund B in every period).

3.4 Return Chasing

The discussion in Section 3.3 indicates that subjects' choice behavior – in particular the fact that they regularly choose fund A, which is suboptimal – may be partially explained by "return chasing". In this section we investigate this issue in a bit more detail.

To start with, Fig. 6 plots the time series of the actual fraction of subjects choosing B against the binary variable that indicates whether the realized return of fund B was larger than that of fund A in the previous period.

There is a clear pattern that the fraction of subjects choosing fund B increases when fund B generated a higher return in the previous periods. Table 7 illustrates a similar effect. It represents (for each treatment and block) the fraction of choices

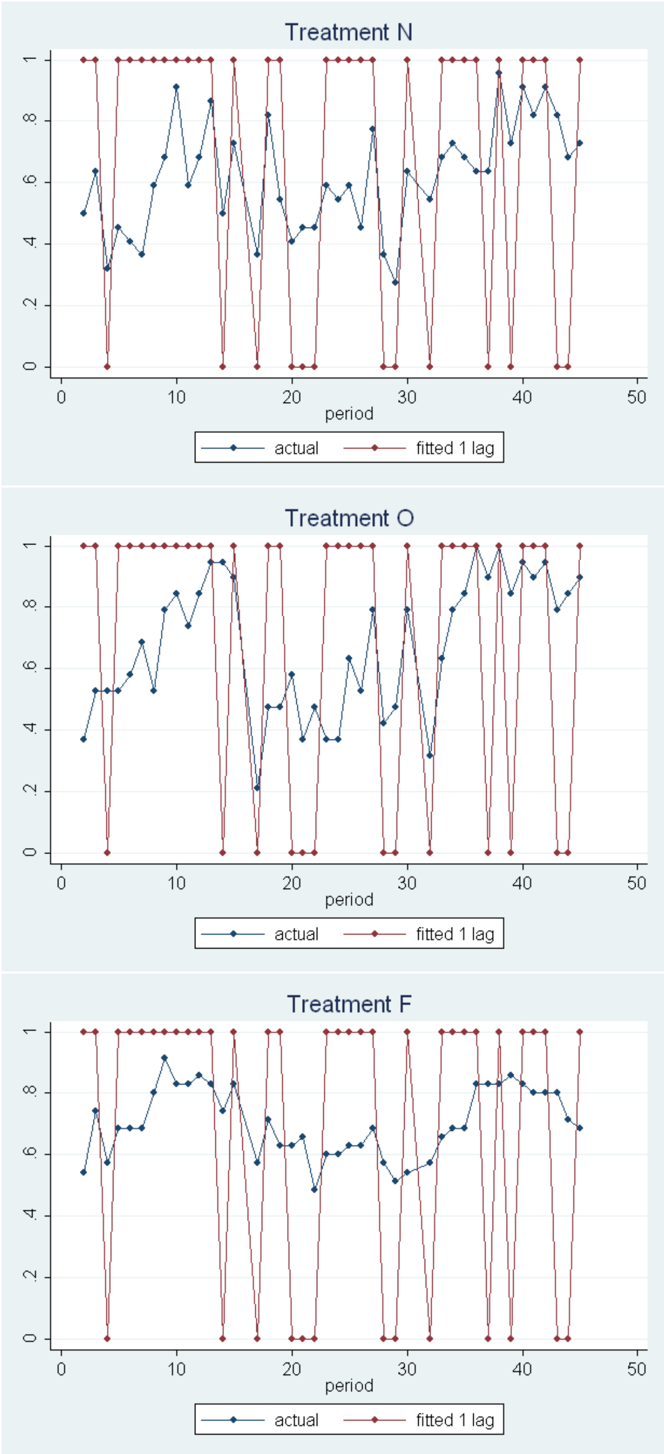

Figure 6: The time series of actual fraction choice of B against the binary variable that indicates whether the realized return of B is larger than A in the last period.

for fund A in periods that are immediately preceded by a period in which the return for fund A was higher than that of fund B, with the numbers in brackets indicating (for that treatment and block) the fraction of choices for fund A in all periods

Treatment	N	O	F
Block 1	59.09% (41.23%)	26.32% (30.08%)	32.86% (23.88%)
Block 2	61.36% (47.73%)	57.90% (50.00%)	41.43% (39.39%)
Block 3	25.33% (25.76%)	12.28% (16.92%)	26.67% (24.29%)
All Blocks	51.24% (38.10%)	39.71% (32.33%)	35.84% (29.18%)

Table 7: The fraction of choices for fund A in periods immediately following a period in which fund A had a higher realized return than fund B. The fraction in between brackets refer to the (unconditional) fraction of choices for fund A in that block.

(except the first period in that block).[19] Table 7 shows that for each treatment the fraction of choices for fund A increases after a positive return difference – this also hold for most of the blocks, with the effect most apparent in the first two blocks of treatment **N** and the second block of treatment **O**.[20] From Figure 6 and Table 7 we conclude that return chasing partially explains the relative high number of choices for fund A in the experiment.

Obviously, return chasing in treatment **F** is less likely than it is in the other two treatments, because switching (back) from fund A to fund B entails a switching cost (the front-end load), which inhibits regular switching between the two funds. In this respect also note that, from the fraction of choices for fund B, as represented in Fig. 6 has an inverse-U shape, in particular for blocks 1 and 3. This can be explained by the fact that the front-end load makes investments more persistent: most participants invest for a (large) number of consecutive periods in fund B, and choose fund A either for the first couple of periods of the block, or for the last couple of periods of the block (or both).

Our results thus far suggests that a model that explains aggregate choices on the basis of past returns may go some way in explaining the results for treatments **N** and **O**. One such model is the so-called discrete choice model (see, e.g., Brock and Hommes (1997)) where the fraction of the population choosing fund B is

[19]We excluded the first period in each block for these fractions in order to facilitate comparisons: in the first period subjects did not observe a realized return yet. This also explains that these fractions are slightly different from those in Table 2, with the numbers in Table 7 typically lower since the number of choices for fund A is relatively high in the first period of each block.

[20]Remarkably, in blocks 1 and 3 of treatment **O** the fraction of subjects choosing A after a positive return difference is *smaller* than the average fraction choosing fund A in those blocks. This can be partially explained by the fact that in the second periods of those blocks (like in the first periods) a relatively large number of subjects choose fund A (without the first return difference being positive): respectively 12 and 13 of the 19 subjects in this treatment. Ignoring this first period in which subject can chase returns the fraction in brackets become 28.34% (instead of 30.08%) and 12.96% (instead of 16.92%), which are both much closer to the average fraction choosing fund A.

Treatment	N	O	F
β_0	0.3798	0.5689	0.7956
p-value	(0.0000)	(0.0000)	(0.0000)
β_1	0.1423	0.0975	0.0628
p-value	(0.0000)	(0.0000)	(0.0000)
McFadden's R^2	0.0327	0.0145	0.0070

Table 8: Estimated coefficients of the logit model in the three different treatments.

modelled as follows

$$n_{B,t} = \frac{\exp\left[\beta_0 + \beta_1\left(r_{B,t-1} - r_{A,t-1}\right)\right]}{1 + \exp\left[\beta_0 + \beta_1\left(r_{B,t-1} - r_{A,t-1}\right)\right]}. \tag{1}$$

Here $r_{A,t-1}$ and $r_{B,t-1}$ refer to the return in percentage points of funds A and B in period $t-1$, respectively. The coefficient β_0 represents a *pre-disposition effect* for choosing fund B and the coefficient β_1, called often *intensity of choice*, measures the sensitivity with respect to past return differences. Given that the optimal decision is to choose fund B in every period we expect β_0 to be positive. A positive value of β_1 implies that subjects are indeed chasing returns. Note that optimal behavior (choosing fund B independent of past returns) corresponds to a (very) high value of β_0 and $\beta_1 = 0$. On the other hand, the case $\beta_0 = 0$ would suggest that subjects only use past returns to decide in which fund to invest, and do not use the information they have about expected returns.

We have estimated equation (1) on the (aggregate) data from all treatments.[21] Table 8 shows the estimation results.

Consistent with our earlier findings the estimated values of β_0 and β_1 are positive for all treatments. That is, subjects in both treatments take into account expected returns for the funds (represented by a positive value of β_0), as well as past return differences (represented by a positive value of β_1). The estimated models predict a higher fraction of optimal choices and less return chasing in treatment **F**, as expected. Moreover, the fraction of optimal choices is higher in treatment **O** than in treatment **N** and return chasing is lower (which is consistent with the descriptive statistics in Tables 2 and 7). In fact, the estimated models predict that the fraction of choices for fund A directly following a period in which fund A had a higher return is equal to 51.17% and 43.13% and 34.01% (which seems to be consistent with the fractions 51.24%, 39.71% and 35.84% from Table 7), and that the overall fraction of choices for fund A is equal to 37.14% and 33.83% and

[21] There is a growing literature that focuses on fitting the discrete choice model to experimental or empirical data, see, e.g., Boswijk, Hommes, and Manzan (2007), Bao, Hommes, Sonnemans, and Tuinstra (2012), Anufriev and Hommes (2012), Anufriev, Hommes, and Philipse (2013) and Anufriev, Bao, and Tuinstra (2015).

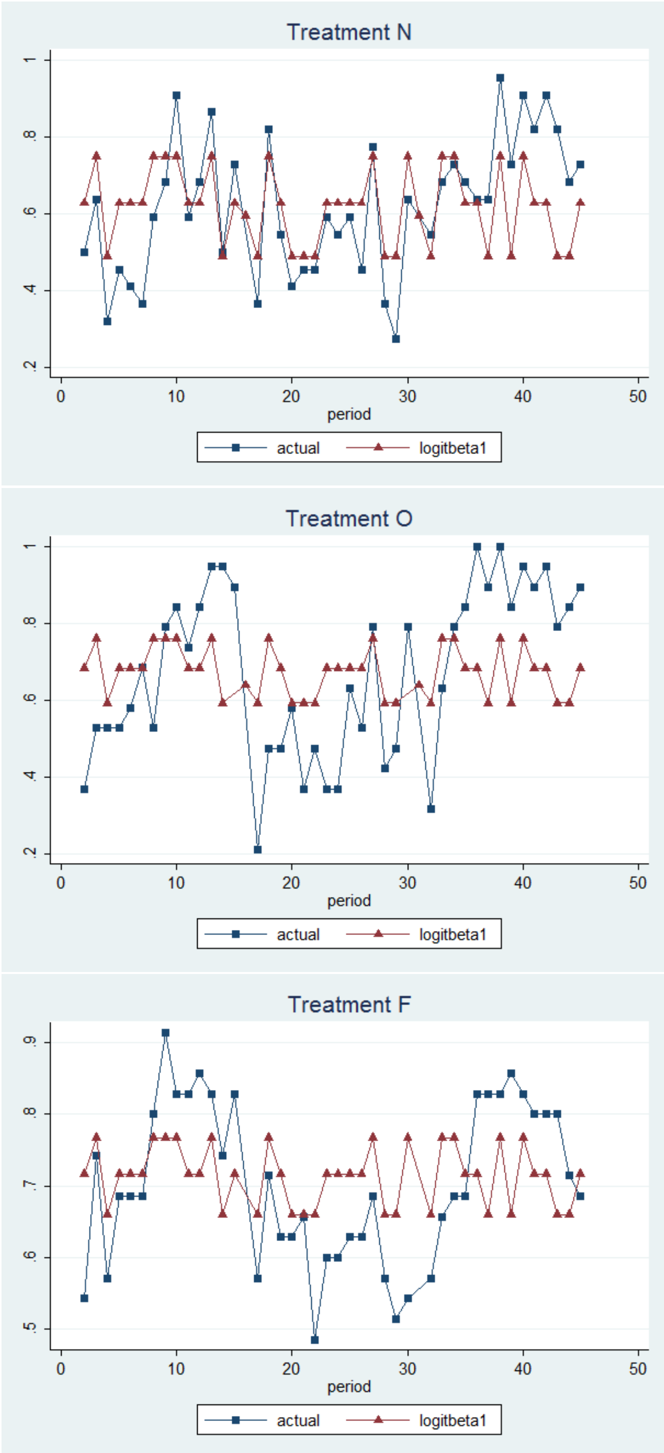

Figure 7: The actual and fitted fractions of subjects who choose B. The squares corresponds to the experimental data, and the triangles to the fitted time series by model (1).

28.25%, respectively (again similar to the fractions of 38.10%, 32.33% and 29.18% from Table 7). More generally, Figure 7 shows that the estimated models capture the choice dynamics in the experimental data quite well.

Based on the findings in this section we have the following result.

Result 3. *We reject Hypothesis 3. Fund choices are, to a substantial extent, explained by past realized returns.*

4 Conclusion

The experiment presented in the paper was aimed at investigating the effect of the fee structure and past returns on mutual fund choice. Subjects have to choose between two experimental funds, where the expected return of fund B is higher than that of fund A. Moreover, expected returns are independent of past returns and subjects know this. We impose different fee structures for fund B in the three treatments, but in such a way that expected returns (over the course of a long-run investment and net of fees) for fund B are the same in each treatment. Our prediction therefore is that investment behavior is the same in the different treatments and, in addition, does not depend upon past returns.

This prediction turns out to be incorrect. There are substantial differences between treatments. On the one hand, subjects pay too little information to the operation expenses fee. On the other hand, the front-end load acts as a commitment device for many subjects and locks them into permanently choosing fund B. Furthermore, in particular for treatments **N** and **O**, we find that subjects behavior can be, to a substantial extent, explained by past returns.

Since our subjects had to repeatedly make investment decisions in a stationary environment, our findings suggest that bounded rationality in mutual fund choice, as for example also found in Choi, Laibson, and Madrian (2010), can not be easily mitigated by experience and learning. It highlights the desirability for regulatory authorities and other government agencies to exert effort in enhancing the transparency of the mutual fund industry, and the level of financial literacy in society.

We find that although a front-end load fee is more salient than an operation fee, it is not more discouraging to investors *per se*. In fact, it is used a commitment device and leads to a lock-in into the more profitable fund. Obviously, such a lock-in does not necessarily lead to higher payoffs. In fact, our findings raise the question whether we would have a similar lock-in if (net of the front-end load) fund B generates *lower* expected returns than fund A. Another interesting direction for future research is to investigate what happens if, after a number of periods, fund A becomes (much) more attractive than fund B and subjects should switch to fund A. In this way we can test whether the sunk-cost fallacy, see Friedman, Pommerenke, Lukose, Milam, and Huberman (2007) and the status quo bias, see Brown and Kagel (2009), play a role in the choice between mutual funds.

We are aware that several authors documented that subjects in financial market experiments (or even simpler experiments) lack game form recognition, see e.g. Chou, McConnell, Nagel, and Plott (2009). These authors show that, to a surprising degree, subjects seem to have little understanding of the experimental environment in which they participate. This has also been underlined by Kirchler, Huber, and Stöckl (2012), who show that running an experiment with a different context ("stocks of a depletable gold mine" instead of "stocks") reduces confusion and thereby significantly reduces mispricing and overvaluation. In our experiment, however, control questions and after-experimental questionnaires suggest that our subjects fully understood the experiment.

References

Anufriev, M., T. Bao, and J. Tuinstra (2015): "Microfoundations for Switching Behaviour in Heterogeneous Agents Models: An Experiment," Mimeo.

Anufriev, M., and C. Hommes (2012): "Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments," *American Economic Journal: Microeconomics*, 4(4), 35 – 64.

Anufriev, M., C. H. Hommes, and R. H. Philipse (2013): "Evolutionary selection of expectations in positive and negative feedback markets," *Journal of Evolutionary Economics*, 23(3), 663–688.

Bao, T., C. Hommes, J. Sonnemans, and J. Tuinstra (2012): "Individual expectations, limited rationality and aggregate outcomes," *Journal of Economic Dynamics and Control*, 36(8), 1101–1120.

Barber, B. M., T. Odean, and L. Zheng (2005): "Out of Sight, Out of Mind: The Effects of Expenses on Mutual Fund Flows," *The Journal of Business*, 78(6), 2095–2120.

Bloomfield, R., and J. Hales (2002): "Predicting the next step of a random walk: experimental evidence of regime-shifting beliefs," *Journal of Financial Economics*, 65(3), 397–414.

Blume, A., and A. Ortmann (2007): "The effects of costless pre-play communication: Experimental evidence from games with Pareto-ranked equilibria," *Journal of Economic Theory*, 132(1), 274–290.

Boswijk, P., C. Hommes, and S. Manzan (2007): "Behavioral Heterogeneity in Stock Prices," *Journal of Economic Dynamics and Control*, 31(6), 1938–1970.

Brock, W., and C. Hommes (1997): "A rational route to randomness," *Econometrica*, 65(5), 1059–1095.

Brown, A. L., and J. H. Kagel (2009): "Behavior in a simplified stock market: the status quo bias, the disposition effect and the ostrich effect," *Annals of Finance*, 5(1), 1–14.

Carhart, M. M. (1997): "On persistence in mutual fund performance," *The Journal of Finance*, 52(1), 57–82.

Choi, J., D. Laibson, B. Madrian, and A. Metrick (2009): "Reinforcement Learning and Savings Behavior," *Journal of Finance*, 64(6), 2515–2534.

Choi, J. J., D. Laibson, and B. C. Madrian (2010): "Why does the law of one price fail? An experiment on index mutual funds," *Review of Financial Studies*, 23(4), 1405–1432.

Chordia, T. (1996): "The structure of mutual fund charges," *Journal of Financial Economics*, 41(1), 3–39.

Chou, E., M. McConnell, R. Nagel, and C. R. Plott (2009): "The control of game form recognition in experiments: understanding dominant strategy failures in a simple two person guessing game," *Experimental Economics*, 12(2), 159–179.

Ehm, C., and M. Weber (2013): "When risk and return are not enough: The role of loss aversion in private investors' choice of mutual fund fee structures," Discussion paper, University of Mannheim, Available at `https://ub-madoc.bib.uni-mannheim.de/34213/1/SSRN-id2252646.pdf`.

Friedman, D., K. Pommerenke, R. Lukose, G. Milam, and B. A. Huberman (2007): "Searching for the sunk cost fallacy," *Experimental Economics*, 10(1), 79–104.

Gil-Bazo, J., and P. Ruiz-Verdú (2009): "The relation between price and performance in the mutual fund industry," *The Journal of Finance*, 64(5), 2153–2183.

Hendricks, D., J. Patel, and R. Zeckhauser (1993): "Hot hands in mutual funds: Short-run persistence of relative performance, 1974–1988," *The Journal of Finance*, 48(1), 93–130.

Investment Company Institute (2014): *2014 Investment Company Fact Book: A Review of Trends and Activities in the U.S. Investment Company Industry*. Available at: `http://www.ici.org/pdf/2014_factbook.pdf`.

Jain, P., and J. Wu (2000): "Truth in Mutual Fund Advertising: Evidence on Future Performance and Fund Flows," *Journal of Finance*, 55(2), 937–958.

Khorana, A., and H. Servaes (2012): "What drives market share in the mutual fund industry?," *Review of Finance*, 16(1), 81–113.

Khorana, A., H. Servaes, and P. Tufano (2009): "Mutual fund fees around the world," *Review of Financial Studies*, 22(3), 1279–1310.

Kirchler, M., J. Huber, and T. Stöckl (2012): "Thar she bursts: Reducing confusion reduces bubbles," *The American Economic Review*, 102(2), 865–883.

Sirri, E., and P. Tufano (1998): "Costly Search and Mutual Fund Flows," *Journal of Finance*, 53, 1589–1622.

Wilcox, R. T. (2003): "Bargain Hunting or Star Gazing? Investors Preferences for Stock Mutual Funds*," *The Journal of Business*, 76(4), 645–663.

Zheng, L. (1999): "Is Money Smart? A Study of Mutual Fund Investors' Fund Selection Ability," *Journal of Finance*, 54(3), 901–933.

APPENDIX

A An Example of Translated Experimental Instructions (Treatment N)

General information. In this experiment you are asked to make subsequent investment decisions. You will start with 1000 points which you can invest. In every subsequent period you will have the possibility to reinvest your accumulated points. In every period you can only invest all of your points in fund A, all of your points in fund B, or invest in neither of the two funds. Your earnings from the experiment will depend upon how well your investments will do.

The funds and their prices. The price of fund A is $P_A(t)$ in period t, and the price of fund B is $P_B(t)$ in period t. Over time prices of the funds grow in the following way. The price of fund A in period $t+1$ is equal to $(1+g_A)$ times the price of fund A in period t, that is

$$P_A(t+1) = (1+g_A) \times P_A(t).$$

The growth rate g_A can only take one of two values. It is either equal to 0.05, or it is equal to 0.01. Both values are equally likely to occur (that is, both occur with the probability equal to 0.5). The history of values of g_A does not influence the probability of either value occurring.

Similarly, the price of fund B grows with growth rate g_B, which could either be 0.06 or 0.02. Again both values are equally likely to occur. The price of fund B in period t therefore is

$$P_B(t+1) = (1+g_B) \times P_B(t).$$

Prices of the two funds do not influence each other. Moreover, your decisions will not influence the price of the two funds.

Example: Suppose the price of fund A is equal to 50 in period 1, and the growth rate in period 1 is equal to 0.05. In that case we have $P_A(2) = 1.05 \times 50 = 52.5$. If the growth rate in period 2 is given by 0.01 then the price in period 3 will be given by $P_A(3) = 52.5 \times 1.01 = 53.03$, and so on.

Investing. If you invest your points in one of the two funds, the number of points you have will grow. For example, suppose you invest your 1000 points in fund A in period 1, when the price of fund A is $P_A(1) = 50$, and you keep your points in fund A until period 6. By then the price of fund A has grown to, for example, $P_A(6) = 60$. Then your points will have increased up to

$$1000 \times \frac{60}{50} = 1200.$$

If you then decide to invest these points in fund B for the next two periods and, for example $P_B(6) = 56$, $P_B(7) = 59$ and $P_B(8) = 61$, your total number of points at the end of period 8 will be equal to

$$1200 \times \frac{61}{56} = 1307.14.$$

by period 10. Note that your points will remain constant in the periods in which you invest in none of the two funds.

Your task. The experiment consists of three parts of 15 periods. In each part you start out with 1000 points, and you can increase the number of points by investing in the two funds A and B. In every period you have three options. Either to invest all of your points in fund A, or to invest all of your points in fund B, or to invest your points in neither fund. You are allowed to switch between funds as often as you want to, but you do not have to.

After the first 15 periods are finished, the experiment will be restarted. Your initial points will be reset to 1000 points and the prices of funds A and B will be reset to their initial values again. The values that the growth rates of the two prices can take are the same again (0.05 and 0.01 for fund A, with equal probability, and 0.06 and 0.02 for fund B, also with equal probability). Because these values are random, the actual growth rates in this second part of 15 periods will, most likely, be different for the actual growth rates in the first part of 15 periods.

After the second part of 15 periods, the experiment will be restarted in the same way as described above for another 15 periods.

Information. The information that you have at the beginning of time t, when you have to make your investment decision for period t, consists of the current prices, all past prices and all past growth rates of both funds. The current prices are shown in the top part of the computer screen. Both past prices and past growth rates are shown in a table on the computer screen. The prices of the funds are also shown in a graph on the screen. Moreover, we show your total accumulated (from the beginning of the current part) number of points in the top part of the computer screen.

Earnings. After the experiment you are paid out according to only one of the three parts. For which part you are paid is determined randomly, and with equal probability. You will be paid for the total number of points, and for each point you will receive 1 euro cent. For example, suppose in the first part your initial number of points increased from 1000 to 1800 points, in the second part your number of points increased from 1000 to 1400 points, and in the final part your number of points increased from 1000 to 1600 points. Then you will earn 18 euros if you are paid according to the first part, and 14 euros if you are paid according to the second part and 16 euros if you are paid according to the final part.

B Control Questions with Answers

B.1 Treatment N

1. Suppose that in the current period the price of fund A is 70 and the price of fund B is 74.1. You have 700 points in the current period and you choose to invest in fund A. Suppose that in the next period the price of funds A and B turned out to be 73.5 and 76.3, respectively. How many points do you have at the beginning of the next period? [735]

2. You have 1100 points at the beginning of the current period and want to invest in fund B. Would your investment decision from the previous period (i.e., in which fund you invested previously) matter for the number of points you will earn? [No]

B.2 Treatment O

1. Suppose that in the current period the price of fund A is 70 and the price of fund B is 74.1. You have 700 points in the current period and you choose to invest in fund A. Suppose that in the next period the price of funds A and B turned out to be 73.5 and 76.3, respectively. How many points do you have at the beginning of the next period? [735]

2. Suppose you have 600 points and you invest your points in fund B whose price in the current period is 57. Fund B charges a fee of 1%. How much fee would you pay for this period? [6]

3. You have 1100 points at the beginning of the current period and want to invest in fund B. Would your investment decision from the previous period (i.e., in which fund you invested previously) matter for the number of points you will earn? [No]

B.3 Treatment F

1. Suppose that in the current period the price of fund A is 70 and the price of fund B is 74.1. You have 700 points in the current period and you choose to invest in fund A for which there is no fee. Suppose that in the next period the price of funds A and B turned out to be 73.5 and 76.3, respectively. How many points do you have at the beginning of the next period? [735]

2. Suppose you invested in fund A in the last period, you have 1000 points at the beginning of this period and want to invest in fund B in this period. Fund B charges a fee of 13%. How much fee would you pay? [130]

3. Recall that fund A does not charge a fee, and fund B charges a fee of 13%. You have 1100 points at the beginning of the current period and want to invest in fund B. Would your investment decision from the previous period (i.e., in which fund you invested previously) matter for the number of points you will earn? [Yes]