

The Follow-on Purchase and Repurchase Behavior of Individual Investors: An Experimental Investigation *

Martin Weber, and Frank Welfens¹

This version: October 2007

Abstract: We analyze two recently documented follow-on purchase and repurchase patterns experimentally: Individual investors' preference for purchasing additional shares of a stock that decreased rather than increased in value succeeding an initial purchase (pattern 1) and investors' tendency for purchasing stocks that they previously sold at a higher price (pattern 2). Similar to the field data study by Odean, Strahilevitz, and Barber (2004), subjects in our experiment are about 2.5 to 3 times as likely to purchase units of a single fictitious good if the price of the good declined following a purchase or sale in the previous period. As an assignment of choices clearly reduces the effect, we argue that investors are involved in counterfactual thinking: They refrain from purchasing additional shares or repurchasing shares at a higher price because doing so means admitting to their ex post wrong decision.

JEL code: C91, D14, D81, G11

* Financial support from the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 504, at the University of Mannheim is gratefully acknowledged. Helpful comments were obtained from participants at the 2005 conference of the German Association for Experimental Economics (GEW).

¹ Martin Weber is from the Lehrstuhl für Bankbetriebslehre, Universität Mannheim, L5, 2, 68131 Mannheim, Germany and CEPR, London. E-Mail: weber@bank.BWL.uni-mannheim.de. Frank Welfens is from the Lehrstuhl für Bankbetriebslehre, Universität Mannheim, L5, 2, 68131 Mannheim, Germany. E-Mail: welfens@bank.BWL.uni-mannheim.de.

1 Introduction

Individual investors do not tend to behave as normative theory would suggest. They do not hold well-diversified portfolios and do not always trade in a direction to increase their portfolio diversification. (See e.g. Blume and Friend 1975, Glaser 2003, and Goetzmann and Kumar 2005 for empirical evidence as well as Kroll, Levy, and Rapoport 1988 and Bossaerts, Plott, and Zame 2006 for experimental evidence.) Nevertheless, for some reason, investors trade excessively, thus reducing their portfolio returns (Odean 1999). Examining individual-level trading behavior reveals transaction patterns that cannot easily be explained by rational decision making. Examples for such patterns are attention-based buying on the purchase side (Barber and Odean 2006) and the disposition effect on the sales side (Odean 1998, Weber and Camerer 1998, Grinblatt and Keloharju 2001).²

Not only initial purchase and selling decisions seem to be biased. As recently shown empirically by Odean, Strahilevitz, and Barber (2004), individual investors also exhibit a tendency to base their follow-on purchases and repurchases on past experiences with the stock, which, from a normative point of view, should not affect (or, by changing total wealth only marginally affect) transactions today. Odean et al. (2004) document that stock purchases occur more often if the stock decreased in value subsequent to a prior purchase (pattern 1), or a prior sale (pattern 2), or if the stock was previously sold at a gain rather than at a loss (pattern 3).

From a normative point of view, follow-on purchases and repurchases should be of low interest, as investors have the opportunity to select their assets from a comprehensive

² See section 4.2 for a short literature overview on these biases.

universe of stocks, derivatives, bonds, and structured financial products. If investors were aware of all investment opportunities, they should only rarely consider purchasing additional shares of a stock they already own or stocks they previously sold. However, as important national and international indices are composed of only a small number of (blue chip) stocks, investors focusing on these stocks are naturally engaged with follow-on purchase and repurchase decisions over years. Investigating how investors react to these situations is thus an important task, if we want to come to a better understanding of general investment decision behavior and want to develop better market models.

We investigate transaction patterns 1 and 2 experimentally. While field data studies like that of Odean, Strahilevitz, and Barber (2004) are concerned with the population we want to study, i.e. individual investors, they naturally fail to reveal whether certain behavior is driven by biased expectations or biased preferences. Our experiment allows us to replicate the behavior described above in a more controllable environment and gives us the opportunity to test different explanations. We conduct an individual choice experiment with 145 students, most of them majoring in economics or business administration. The experiment consists of four successive periods in which our subjects can buy and sell units of a single good. The price of the good changes every period. Based on our experimental data we are able to replicate patterns 1 and 2: Our subjects prefer purchasing units of the single good if, subsequent to a prior purchase or sale, the good decreased rather than increased in value.

Odean, Strahilevitz, and Barber (2004) attribute these biased transaction patterns to the psychological concept of “counterfactual thinking”: Decision makers tend to create mental simulations of “what might have been” if they decided differently in the past or if certain events in the past did not happen or happened differently. If this mental simulation, the counterfactual, would have lead to a better outcome, decision makers shy away from actions that

would make their counterfactual more salient. For example, if an investor sold a stock for a low price and, succeeding the sale, the stock went up in value, he does not wish to repurchase the stock, because repurchasing would mean admitting that his prior decision to sell the stock was *ex post* a wrong one. Similarly, if the investor purchased a stock for a low price, he is likely to refrain from purchasing additional shares at a higher price, because this repurchase would confront him with a counterfactual in which he made a larger initial purchase.

Counterfactuals are reported to be affected by antecedent controllability (Kahneman and Miller 1986, Girotto, Legrenzi, and Rizzo 1991, Markman, Gavanski, Sherman, and McMullen 1995, Roese and Olson 1995a), which, in turn, depends on whether the decision maker decided freely or not in the past. Hence, we control for counterfactual thinking in two further treatments, in which we assign our subjects the number of units to purchase and sell in the first two periods of the experiment. An assignment of choices clearly reduces the effect: While subjects in the assigned-choice treatments purchase as often as in the standard treatment following a price decrease, succeeding a price increase, their purchase frequency is about two times as high if the prior purchase or selling decision was assigned. Consistent with the psychological literature (Roese and Olson 1995a and Roese 1997), we thus find that especially upward counterfactuals exhibit a strong impact on investment decisions.

The paper is structured as follows: In section 2, we first review related literature on individual investors' trading patterns, focusing on follow-on purchase and repurchase decisions. We then summarize possible explanations for investors' purchase behavior and derive our hypotheses. In section 3, we explain the experimental design and procedure. Section 4 presents our results, and section 5 draws conclusions. Translations of the experimental instructions and questionnaires can be found in the appendix.

2 Related Literature, Competing Explanations, and Hypotheses

Recent research documents that individual investors' purchasing and selling decisions strongly deviate from normative behavior. One of the best investigated biases in behavioral finance is the disposition effect. As first noted by Shefrin and Statman (1985), individual investors more readily sell stocks at a gain rather than at a loss. (See e.g. Odean 1998, Weber and Camerer 1998, and Grinblatt and Keloharju 2001 for evidence.) The most common explanation for this behavior is based on prospect theory's s-shaped value function (Kahneman and Tversky 1979, Tversky and Kahneman 1992) in combination with mental accounting (Thaler 1980, 1985), assuming some weighted average of the stock's past purchase prices as the reference point. On the other hand, when looking for a purchase opportunity, investment decisions are affected by individual investors' inability to follow all stocks in the market. Individual investors looking for a stock to purchase have to limit their choice set to only the small selection of stocks of which they are aware, e.g. stocks that recently garnered their attention. Lee (1992) and Hirshleifer, Myers, Myers, and Teoh (2003) find that individual investors are net buyers following both negative and positive earnings surprises. Barber and Odean (2006) provide further evidence of so-called attention-based buying using trading volume, extreme one day returns, and news as proxies for attention grabbing events.

Much less is known about individual investors' follow-on purchase and repurchase behavior. The only direct investigation comes from Odean, Strahilevitz, and Barber (2004). The authors analyze two U.S. data sets covering 66,465 individual investors with accounts at a large discount broker and 665,533 individual investors with accounts at a large retail broker. The data sets contain transactions from January 1991 through December 1996 and January 1997 through June 1999, respectively. Odean, Strahilevitz, and Barber discover three previously undocumented patterns in investors' purchasing decisions. Investors seem to prefer

- purchasing additional shares of stocks that decreased rather than increased in value subsequent to a prior purchase (pattern 1),
- repurchasing stocks that lost value rather than increased in value subsequent to a prior sale (pattern 2), and
- repurchasing stocks they previously sold at a profit rather than stocks they previously sold at a loss (pattern 3).

In our experiment, we investigate the first two patterns in detail.³ In the experiment we are able to control for alternative explanations that cannot be controlled for with field data. We also apply different treatment groups to shed some light on why people exhibit this behavior.

The follow-on purchase and repurchase patterns found in individual investors' trading behavior could be explained in different ways. Investors possessing superior private information could sell overpriced stocks and later repurchase these stocks at or below fundamental value, consistent with pattern 2. Alternatively, investors might believe in mean reverting stock prices and purchase stocks that recently decreased in value because they expect these stocks to rebound to the old price. We use the advantage of a controlled laboratory experiment to create conditions where neither of these explanations apply: Our experiment can be described as a simple multi-stage lottery framed as a single good market. Our subjects are not provided with superior information about this good, and, as the probabilities of the multi-stage lottery are constant over all periods, any belief in mean reversion becomes obviously irrational. The

³ We refrain from testing pattern 3, as subjects in our experiment rarely sell at a loss.

remaining behavioral explanations we want to discuss in the following are based on standard prospect theory and counterfactual thinking.

We start with pattern 1, i.e. investors' attitude for purchasing additional shares if the stock recently decreased in value. Prospect theory is consistent with this behavior as it predicts that the investor becomes risk seeking and more willing to escalate his commitment if he had previously incurred a loss. Purchasing additional shares of a stock at a lower price decreases the average purchase price, increases the probability of breaking even, and probably lowers the investor's reference point. To some extent, the same explanation applies to pattern 2, i.e. investors' preference for repurchasing stocks that decreased in value subsequent to a previous sale. If the investor did not sell all shares in the past, some (unrealized) gains and losses remain in his mental account. If the stock drops in value after the sale, the investor's mental account is more likely to contain an unrealized loss as opposed to if the stock had increased in value. Thus, investors should repurchase shares more often following price decreases. However, if an investor sold all shares in the past and closed his mental account, there should be no difference between the two cases.

Although patterns 1 and 2 are at least partly explained by standard prospect theory, Odean et al. (2004) attribute the above discussed transaction patterns to the psychological concept of "counterfactual thinking" (Kahneman and Tversky 1982, Kahneman and Miller 1986). The term counterfactual thinking refers to people's natural tendency to create mental simulations of "what might have been" had a certain event in the past not happened, or had they themselves made a different decision (see Kahneman 1995, Roese and Olson 1995b, and Roese 1997 for overviews of the literature). Counterfactual thinking is assumed to be important for everyday life and enables people to learn from their prior mistakes. If a person compares a situation he is actually in with some counterfactual, the comparison evokes feelings of

joy or regret, depending on whether the counterfactual would have led to a worse (downward counterfactual) or better (upward counterfactual) outcome. These feelings, in turn, interfere with subsequent behavior. For example, a decision maker might think counterfactually if he decided to participate in a lottery that led to a negative payout. A reasonable counterfactual would be to compare this loss situation to a mental simulation in which the decision maker refrained from playing the lottery. Note that counterfactual thinking and prospect theory are not necessarily conflicting explanations, as counterfactuals could also be interpreted as possible candidates for prospect theory's reference point (see Weber and Welfens 2007). However, counterfactual thinking emphasizes that not only real gains and losses but also hypothetical or missed gains and losses are important for the decision maker's subsequent behavior.

Imagine an investor who purchased a stock in the past and now has to decide on whether to purchase more shares of the same stock (the situation considered in pattern 1). The investor is likely to compare his current situation with the situation he could have been in had he made a larger initial purchase. If this comparison makes him feel good, i.e. if since the initial purchase the stock price has declined (downward counterfactual), he is more likely to purchase additional shares. If, instead, the stock price has increased (upward counterfactual) he is aware that he could have been better off by simply purchasing more shares at the time of the initial purchase. The investor may avoid further purchases to forget about this counterfactual.

Pattern 2 may be explained in a similar way: An investor who sold a stock in the past and now needs to decide on whether to repurchase this stock is likely to create a mental simulation in which he simply did not sell the stock. Consequently, he might compare his actual situation with two different counterfactuals: If, since selling the stock, the price has declined (downward counterfactual), reality prevails over the mental simulation; the investor is glad to have sold the stock, because he can now repurchase it at a lower price. He may think about the avoided loss as a gain. Instead, if the stock price has increased after the sale

the avoided loss as a gain. Instead, if the stock price has increased after the sale (upward counterfactual), the investor is sad about his ex post wrong decision. To avoid a feeling of regret, the investor tries to avoid paying attention to a stock that has increased in value since he sold it. He is less likely to repurchase the stock, because repurchasing makes his counterfactual more salient. Counterfactual thinking can thus explain pattern 2, even if the investor decided to sell all shares in the past.

In our experimental study we first try to replicate the findings of Odean, Strahilevitz, and Barber (2004) leading to the following hypotheses:

Hypothesis 1: (Pattern 1) Individual investors are more likely to purchase additional shares of a stock if this stock has lost value rather than increased in value since the initial purchase.

Hypothesis 2: (Pattern 2) Individual investors are more likely to repurchase a previously sold stock if this stock has lost value rather than increased in value since being sold.

In addition, we want to test whether patterns 1 and 2 are solely driven by changes in risk aversion following gains and losses, i.e. prospect theory, or whether counterfactual thinking (and thus hypothetical gains and losses) can add something to understanding this behavior. To test for counterfactual thinking, we apply the following mechanism: As known from psychological research, controllable and thus mentally mutable antecedents tend to create more counterfactual thoughts than uncontrollable antecedents. (See Kahneman and Miller 1986 for theory, Girotto, Legrenzi, and Rizzo 1991, Markman, Gavanski, Sherman, and McMullen 1995, and Roese and Olson 1995a for evidence and McMullen, Markman, and Gavanski 1995 for discussion.) Hence, decision makers that are personally responsible for an

ex-post wrong decision should feel more regret and be more affected in their follow up decisions. In our experiment, we vary antecedent controllability by assigning prior purchases and sales to two treatment groups. If the follow-on purchase and repurchase patterns described above are due to counterfactual thinking, we expect them to be more salient if the investor decided freely on the prior purchase or the prior sale. We expect them to be weaker if the decision to purchase or sell the stock was assigned by the experimenter. This leads to the following hypotheses:

Hypothesis 3: (Evidence of counterfactual thinking as an explanation for pattern 1) The follow-on purchase pattern described by hypothesis 1 is stronger if the investor decided freely on the initial purchase and weaker if the initial purchase was assigned.

Hypothesis 4: (Evidence of counterfactual thinking as an explanation for pattern 2) The repurchase pattern described by hypothesis 2 is stronger if the investor decided freely on the prior sale and weaker if the prior sale was assigned.

3 Experimental Design

Studying follow-on purchases and repurchases means investigating the second and third step of a decision sequence concerning the same stock. In our experiment, we capture these decision sequences with a four-period design.

3.1 Basic Design

We apply an experimental design similar to Weber and Zuchel (2005). Weber and Zuchel study how decision makers are affected by gains and losses under different choice

frames. They set up a three period individual choice experiment, in which subjects could buy and sell units of a fictitious asset (in an asset frame) or purchase lottery tickets (in a lottery frame). In the asset frame, prices of the fictitious asset change from period to period. Subjects need to decide how many units to purchase and sell in periods 1 and 2 while the purpose of period 3 is only to determine the final value of the asset. We stick to the asset frame and extend the experimental design to four periods so that our subjects now need to make three decisions in a row.

Our experiment therefore consists of three successive decision periods labeled period 1 to 3, with a fourth period determining final wealth and payout. In period 1, subjects are endowed with €12, which they can use to purchase up to 40 units of a single good⁴. Subjects retain all money not used for purchasing units. In periods 2 and 3 subjects can either purchase (additional) units of the single good or sell units they bought in previous periods, getting back the current price in Euros. They can also decide to do nothing and keep their number of units constant. The only constraints on our subjects' decisions are that their individual money account as well as their number of units held have to remain non-negative. In contrast to Weber and Zuchel (2005), our subjects do not receive additional money over periods and are not allowed to buy on credit. We eliminate both design properties to ensure that results are not driven by subjects' unwillingness to purchase on debt.

After each decision, the unit price changes: Starting at 30 cents in period 1, the price either increases by 12 cents or decreases by 10 cents from period to period. Price increases and decreases are equally likely and are independent of previous price changes. Our subjects

⁴ We use the German terms “Einheiten” (“units”) and “Gut” (“good”) rather than “shares” and “stock” to avoid framing effects.

are informed about these probabilities and corresponding future prices of the good. The event tree in figure 1 shows all possible price developments.

(please insert figure 1 about here)

From a normative point of view, our subjects are confronted with three successive lotteries in which they need to decide how much money to bet. Holding a number of x units of the good for one period means participating in a lottery which either results in a gain of 12 cents times x or a loss of 10 cents times x , with both outcomes being equally likely. The lottery leads to an expected payoff of x cents, but also increases in risk with the number of units held. While a risk neutral subject should purchase the maximum number of 40 units in period 1 and hold these units throughout the experiment, a risk averse subject might purchase fewer units. Ignoring income effects, subjects should not change their exposure to risk over periods and should neither purchase nor sell units of the single good in periods 2 and 3. Decreasing absolute risk aversion, on the other hand, makes a subject more willing to purchase additional units following a price increase and to sell units succeeding a price decrease – the opposite of what we expect by our hypotheses.⁵

We ask our subjects for their complete strategy using Selten's (1967) strategy method. This method, commonly used in game theory experiments, demands subjects to devise strategies for each of their opponents' possible moves. In our experiment we ask our subjects to decide how many units to buy or sell at each node of the first three periods of the event tree (nodes 1 to 7 in figure 1). Decisions in later periods are always based on those from prior pe-

⁵ Note that holding portfolio weights constant as proposed by portfolio theory in combination with constant relative risk aversion is not a normative strategy in our experiment. The substitution effect only arises if absolute price changes depend on current price levels, i.e. if relative price changes were constant. Instead, in our experiment, price changes are fixed in absolute terms. The amount of money a subject puts at risk is thus only a function of how many units the subject holds, not a function of current prices.

riods. For example: We first ask a subject how many units of the single good he wants to buy in period 1 (node 1). Afterwards (in period 2) we ask him how many units he wants to buy or sell if the unit price has either increased to 42 cents (node 2) or decreased to 20 cents (node 3). The number of units he can sell in period 2 is limited by the number of units purchased in period 1. The number of units he can purchase in period 2 is restricted by his money account. Figure 2 shows a screenshot from treatment 1.

(please insert figure 2 about here)

At the end of the experiment, one path in the event tree is chosen randomly by rolling a die for each price change. Even numbers mark price increases while odd numbers indicate price decreases. The subject's strategy for the resulting path is played and the total value of his monetary units plus the final value of his holdings of the single good is determined. The subject is paid this sum (at face value) in Euros.

To control for sequence effects, we vary the order in which we ask for transaction decisions 1 to 7. Sequence effects may emerge if a subject is faced with positive (or negative) feedback first. Consequently, half of our subjects pass through nodes 1 to 7 in ascending numerical order so that decisions after price increases always precede decisions after price decreases. For the other half, decisions after price increases succeed decisions after price decreases so that the order changes to 1, 3, 2, 7, 6, 5, 4.

3.2 Different Treatments

The experimental design so far allows us to test hypotheses 1 and 2. To test hypotheses 3 and 4, we have to change our design in two steps. In our standard treatment discussed above, our subjects decide freely in all periods. In a second treatment, which is run with another subject group, we eliminate choice in period 1. Rather than having our subjects decide

how many units to buy, the number of units is assigned by the experimenter. We assign our subjects a purchase of 20 units of the single good in period 1, which is the median number of units bought by subjects in period 1 of our standard treatment. For periods 2 and 3, subjects have to decide independently on their transactions. In a third treatment, which is run with yet another subject group, we assign not only the number of units in period 1 but also the number of units in period 2, and only ask our subjects for period 3 decisions. Again, all assignments are the median values from our standard treatment: Subjects are assigned to purchase 20 units in period 1, to sell 10 units in period 2 in the case of a price increase (node 2), and to buy 14 units in period 2 in the case of a price decrease (node 3). The assignment of choices allows us to test whether follow-on purchases and repurchases are driven by standard prospect theory or (an extended version including) counterfactual thinking.

3.3 Questionnaire

In the second part of the experiment, common to all treatments, we ask our subjects to complete a short questionnaire in which we address the investigated patterns and potential explanations quite directly. To make sure that our subjects are not influenced by their performance in the first part of the experiment, they have to answer the questionnaire after they have given their complete strategy and before the actual path in the event tree and their monetary payoff is determined. A translation of the questionnaire can be found in the appendix.

3.4 Procedure

The experiment was conducted in December 2004 at the University of Mannheim. Our subjects were 107 male and 38 female students, most of them majoring in economics or business administration. Our standard treatment was run with 80 subjects, while 33 and 32 subjects participated in the assigned-choice treatments, respectively. Half of the subjects in each

treatment was assigned to make their decisions in ascending numerical order, while the other half was assigned to the alternative sequence specified in section 4.3.1. Our subjects had an average age of 23 and have been studying at university for an average of 2.3 years.

After reading the instructions and before participating in the first part of the experiment, subjects had to answer a short questionnaire on the rules of the game, the price process, and probabilities involved to ensure that the procedure was fully understood. If a subject answered one of these questions incorrectly, which happened in only five cases, they were required to re-read the instructions and to answer the questionnaire a second time before starting the experiment. The experiment was conducted in a computer laboratory using a self-made software tool programmed in PHP. The average processing time was approximately 45 minutes. The average payoff was €12.60, ranging from €0 to €22. The entire experiment was conducted in German. A translation of the instructions for the individual choice part of our standard treatment (in ascending numerical order) can be found in the appendix.

4 Results

We expect our subjects to exhibit the follow-on purchase and repurchase patterns discussed in section 2. We test for these patterns by analyzing similar decision sequences within our experimental data set, i.e. purchases preceded by purchases and purchases preceded by sales.

4.1 Definition of Variables

For analyzing purchase behavior, it is not sufficient to compare the number of follow-on purchases and repurchases following price increases and decreases. We want to ensure that

results are not driven by monetary constraints, which could prohibit subjects from purchasing additional units at higher prices. Hence, similar to standard disposition effect measurements (see e.g. Odean 1998), we need to control for the number of purchase opportunities.

Our procedure is as follows: We analyze every node in the event tree separately and count purchases and purchase opportunities. A purchase opportunity is recorded if the subject has sufficient experimental money to buy at least one more unit of the single good.⁶ A purchase is counted whenever a subject decides to actually acquire units of the good, regardless of quantity. Based on the two measures and controlling for different path properties and treatments, we then compute follow-on purchase frequencies and repurchase frequencies succeeding price increases and decreases as the relation between corresponding purchases and purchase opportunities. For follow-on purchase frequencies, we count only those purchases and purchase opportunities that are preceded by another purchase in the previous node. Similarly, for repurchase frequencies, we focus on transaction decisions that are preceded by a sale in the previous period. These frequencies serve as estimates of our subjects' conditional purchase probabilities.

While hypotheses 1 and 2 are analyzed using our standard treatment only, we refer to treatments 2 and 3 to account for counterfactual thinking.

4.2 Follow-on Purchase Behavior

Hypothesis 1 expects individual investors to be more likely to purchase additional shares of a stock they already hold in their portfolios if the stock price declines rather than

⁶ Although subjects are allowed to purchase and sell half a good and smaller quantities, their transactions are always in whole numbers.

increases subsequent to a prior purchase. We test this hypothesis by comparing our subjects' purchase behavior between nodes 2 and 3 and nodes 6 and 7 in our standard treatment, i.e. treatment 1. Purchases and purchase opportunities at nodes 4 and 5 are not taken into account, to ensure that our results are comparable to the assigned-choice treatments 2 and 3.⁷ We calculate follow-on purchase frequencies separately for nodes 2 and 6 (following price increases) and nodes 3 and 7 (following price decreases) as discussed in section 4.1 and apply a simple binomial test on these measures. The first block of table 1 documents the results.

(please insert table 1 about here)

In line with hypothesis 1, subjects are over 2.5 times as likely (66.67 % vs. 26.23 %) to purchase additional units of the good if the price decreased rather than increased succeeding the initial purchase. The difference in follow-on purchase frequencies is highly significant ($p = 0.0000$). While this behavior is not consistent with standard expected utility theory in combination with constant or decreasing absolute risk aversion, it coincides with behavioral explanations like standard prospect theory or counterfactual thinking.

A potential shortcoming of our statistical test documented above is that we assume every purchase rather than every subject to constitute an independent observation. We perform a second, more robust test to verify the effect: First, we limit our data set to those subjects in treatment 1 that exhibit follow-on purchase opportunities both following price increases (at node 2, node 6, or both nodes) and succeeding price decreases (at node 3, node 7, or both nodes). In a second step, we determine for each subject whether he used at least one of

⁷ As subjects in treatment 3, node 2, were assigned to sell, they do not have follow-on purchase opportunities at nodes 4 and 5. Nodes 4 and 5 instead provide the opportunity to study repurchase behavior, i.e. purchases preceded by a sale in the previous period. See section 4.4.3.

his follow-on purchase opportunities after price increases. Conversely, we record whether the subject used at least one of his purchase opportunities following price decreases. We then perform a matched-pairs Wilcoxon test on these variables, expecting that those subjects with follow-on purchase opportunities both following price increases and decreases are more likely to use their opportunities in the second case. Again, the result is highly significant ($p = 0.0000$). 15 and 22 subjects use their purchase opportunities in neither or both cases, while 33 (4) subjects only purchase following price decreases (increases).

Further validation comes from our subjects' answers to question 2 of the final questionnaire: Subjects are told that they purchased one unit of good A and one unit of good B in the past and that since the purchase good A increased and good B decreased in value. They are asked to indicate whether they prefer to purchase an additional unit of good A or good B. While out of 145 subjects only 17 claim to be indifferent, the majority, i.e. 84 of the remaining 128 subjects, prefer to buy an additional unit of good B (65.63 %, $p = 0.0003$). Subjects thus not only act purely intuitively, but also behave consistently if the similarity between the two situations becomes salient.

We want to shed light on the above documented follow-on purchase pattern by investigating what kind of rationale may be behind our subjects' behavior. We do so by comparing transactions in our standard treatment with transaction behavior in the two assigned-choice treatments. While in the standard treatment our subjects have carte blanche at all nodes, decisions in period 1 of treatment 2 and in periods 1 and 2 of treatment 3 are predetermined by the experimenter. If our subjects commit some kind of counterfactual thought as discussed in section 2, we should expect them to be less susceptible to the follow-on purchase pattern if the prior purchase decision has been assigned. We therefore repeat the tests described above with

the two assigned-choice treatments based only on our subjects' decisions at nodes 2 and 3 of treatment 2 and 6 and 7 of treatment 3. The lower part of table 1 provides the results.

In the assigned-choice treatments, subsequent to a price increase (decrease) subjects use their follow-on purchase opportunities in 41.54 % (67.69 %) of all cases. As the difference between purchase frequencies is significant on a 1 percent level ($p = 0.0014$) we ascertain that the pattern still emerges if the prior purchase has been assigned. The effect, however, unfolds to a smaller extent. While following a price decrease purchase frequencies are almost identical across treatments 1, 2, and 3 ($p = 0.4434$), following a price increase subjects are much more likely to purchase additional shares if they are not responsible for the initial purchase decision ($p = 0.0160$). Consistent with the psychological literature (Roesse 1997 and Roesse and Olson 1995a), our subjects are more strongly affected if purchasing additional shares means dealing with an upward counterfactual like "I should have bought more shares in the first place. They are so expensive now!". Upward counterfactuals, in general, are documented to evoke counterfactual thinking more easily and to have a stronger impact on succeeding behavior than downward counterfactuals ("I am glad that I did not buy so many shares in the initial purchase. They are cheaper now!"). While the first case results in strong feelings of regret, which exacerbate if the subject actually decides to purchase more shares at a higher price, the second case evokes only weak pleasure, which influences subsequent behavior only marginally.

As a robustness check, in the final questionnaire we ask our subjects whether they would feel better in situation a or situation b: Subjects are informed that they purchased a good in the past. Today they recognize that since the purchase the price of the good has decreased. While in situation (a) the subject purchased the good although it was unnecessary, in situation (b) the subject bought the good because he urgently needed it. Out of 145 subjects

only 16 are indifferent to the two situations, while of the remaining 129 subjects who stated a preference, 97 feel better in situation b (75.19 %). Consistent with their behavior in the experiment and with counterfactual thinking, subjects seem to prefer not to be responsible for an ex post wrong decision ($p = 0.0000$).

4.3 *Repurchase Behavior*

Hypothesis 2 argues that individual investors are more likely to repurchase a stock if the stock decreased rather than increased in value subsequent to a prior sale. We stick to nodes 4 and 5 for investigating this behavioral pattern to be again comparable to the assigned-choice treatment 3.⁸ Table 2 provides the results.

(please insert table 2 about here)

The table reveals that, consistent with hypothesis 2, subjects are three times as likely (52.50 % vs. 17.50 %) to repurchase units of the good in period 3 if, since the sale in period 2, the price of the good has declined rather than increased. The effect is not only economically but also statistically highly significant ($p = 0.0005$).

In a second test we investigate if the same pattern can be observed on an intra-subject basis. We therefore check whether subjects with repurchase opportunities at both nodes 4 and 5 use their opportunities more often subsequent to a price decrease than following a price increase. A matched-pairs Wilcoxon test reveals robust results ($p = 0.0013$). While 17 and 5 subjects use their opportunities in no case or in both cases, 16 and 2 subjects repurchase only following price decreases or increases, respectively.

⁸ In treatment 3, subjects were assigned to sell at node 2 and to purchase at node 3. Hence, only nodes 4 and 5 provide repurchase opportunities in this treatment.

A final test of hypothesis 2 is based on question 3 of the final questionnaire. We ask our subjects whether they prefer to repurchase good A or good B. While both goods were sold in the past without realizing gains or losses, good A increased and good B decreased in value subsequent to the prior sale. Only 20 out of 145 subjects claim to be indifferent between the two repurchase opportunities while 106 of the remaining 125 subjects prefer to repurchase good B (84.80 %, $p = 0.0000$), confirming our previous results.

We perform the same tests for the assigned-choice treatment 3, i.e. we want to find out whether our subjects' preference for repurchasing units of a good that has declined rather than increased in value levels off if the prior sale in period 2 was assigned by the experimenter. Our tests are again based on our subjects' decisions in nodes 4 and 5. Results are reported in the bottom part of table 2. While subjects still prefer repurchasing following a price decrease, the effect loses its statistical significance ($p = 0.1587$). This loss in significance is mainly driven by our subjects' increased willingness to repurchase the good following a price increase if the prior selling decision was assigned. Repurchase frequencies following price increases are significantly different between treatments 1 and 3 ($p = 0.0074$), while they are almost equal succeeding price decreases ($p = 0.3755$). Similar to the follow-on purchase pattern analyzed in section 4.2, our subjects shy away from repurchasing a good they have sold for a lower price if they are personally responsible for the sale and thus confronted with a salient upward counterfactual ("Why did I sell this stock? It increased in value!").

A final test is based on our subjects' answers to question 5. Our subjects are told that they sold a good in the past and today realize that, since the sale, the price of the good has increased. They are asked to indicate whether they would feel better (a) if they sold the good in the past although they did not need to or (b) if they sold the good because they needed the money. Out of 145 subjects, 15 state indifference, while 102 of the remaining 130 subjects

claim to feel better in the second situation (78.46 %). As the frequency of answer b is highly significant ($p = 0.0000$), hypothesis 4 receives further support. Whether a good was sold freely or under compulsion influences our subjects' feelings and therefore also impacts their transactions. Our subjects seem to prefer not being responsible for an ex-post wrong decision and try to avoid the regret coming from an upward counterfactual thought by not considering repurchasing a good that went up in value.

5 Conclusion

Based on an individual choice experiment run with 145 students, we document two trading patterns concerning follow-on purchases and repurchases: Our subjects prefer purchasing units of a single good if it, subsequent to an initial purchase (pattern 1) or a prior sale (pattern 2), decreased rather than increased in value. Follow-on purchases and repurchases are 2.5 times and three times as likely following a purchase or sale at a higher price. Our experimental design ensures that these findings cannot simply be explained by belief in mean reversion or superior private information.

We instead argue that our subjects display characteristics of counterfactual thinking, a psychological concept that is documented to affect people's behavior in everyday life (Kahneman and Tversky 1982, Kahneman and Miller 1986). Counterfactual thinking can be described as people's natural tendency to create mental simulations of "what might have been" if they decided differently in the past or if certain events in the past did not happen or happened differently. Repurchasing a good following a prior sale is likely to evoke thoughts of having sold fewer units of the good or not having sold the good at all. Purchasing additional units of a good subsequent to a purchase in the previous period, correspondingly, causes sub-

jects to run a simulation in which they made a larger initial purchase. We believe that subjects compare reality with these counterfactuals. If this comparison makes the subject feel good, i.e. he would be worse off having previously made another decision, he is more likely to follow-on purchase or repurchase the good in order to make his downward counterfactual more salient. If the comparison makes him feel bad, e.g. because the good increased in value subsequent to the prior purchase or sale, he refrains from purchasing the good to forget about his ex-post wrong decision.

We control for this counterfactual story with two additional experimental treatments, in which we assign subjects to purchase and sell units of the good in previous periods. Assignment of choices is known to alter mental antecedent controllability, which, in turn, determines whether subjects are able to create salient counterfactual thoughts (Kahneman and Miller 1986, Girotto, Legrenzi, and Rizzo 1991, Markman, Gavanski, Sherman, and McMullen 1995, Roese and Olson 1995a). If they are not responsible for the prior purchasing or selling decision, subjects are less prone to the discussed follow-on purchase and repurchase patterns. The effect, however, is an asymmetric one: Following a purchase or sale at a higher price, subjects purchase additional units or repurchase units of the good in approximately 60 % of all cases - regardless of whether the prior purchase or sale was free or assigned. Conversely, following a price increase we obtain significant differences in follow-on purchase and repurchase frequencies with approximately 20 % in the free choice treatment and 40 % in the assigned-choice treatments. Subjects are thus less likely to purchase additional units or to repurchase units if they decided freely on the previous purchase or sale and are therefore responsible for an ex post wrong decision. This, in turn, provides evidence that counterfactual thinking affects investment decisions.

Although our results are derived from an experiment using students instead of real investors and conditional decisions based on a strategy method, we believe that our study contributes to understanding real investors' transaction behavior: Our students should know at least as much about normative theory and optimal investment decisions as most individual investors do. If we find that economics students are biased in their transactions by prior experiences and decisions, we expect to see the same or even exacerbated behavior when investigating less sophisticated decision makers. The strategy method, on the other hand, demands that our subjects make conditional decisions. Thinking "what if I incur a loss?" might result in decisions other than facing an unconditional real loss. If there were an impact, however, one would expect subjects to exhibit less emotion and show weaker reactions to hypothetical gains and losses than to real ones, thus weakening all effects and biases. Hence investigating these biases using a strategy method provides a conservative test. If we are able to find behaviorally motivated trading patterns, we can expect them to be even stronger in reality.

Further research is required to find out how and when investors think counterfactually and how exactly counterfactual thoughts interfere with subsequent investment decisions. While in our experiment, subjects, for the sake of simplicity, only traded in one good, one could also assume that alternative investment opportunities and their performance serve as comparisons for counterfactual thoughts. Understanding counterfactual thinking in the investment context might lead to better models of how investors form and update their individual reference points over time and could thereby enhance our understanding of individual biases and potential market outcomes.

Appendix: Instructions and Questionnaire

The following subsections contain translations of the German instructions for the standard treatment, i.e. treatment 1, as well as the final questionnaire. Instructions for treatments 2 and 3 differed only marginally from treatment 1.

Instructions for Treatment 1

You are participating in an experiment conducted by the Chair of Banking and Finance at the University of Mannheim. The experiment is funded by the Deutsche Forschungsgemeinschaft. It will take about 45 minutes, including time for reading instructions and completing a short questionnaire at the end. Please do not communicate with other participants during the experiment.

If you read the following instructions carefully, you will be able to earn a considerable amount of money. The exact amount depends on the decisions you make. You will be paid in cash immediately after the experiment. Please signal us if you have problems understanding these instructions.

In this experiment, you buy and sell units of a fictitious good in three successive periods. After each period, the price of the good will change. If it rises, you will make a profit. If it decreases, you will incur a loss.

Period 1

In Period 1, you receive an initial endowment of €12. With this money, you can buy up to 40 units of the fictitious good. Each unit of the good costs exactly 30 cents in period 1. After you have made your decision, the price of the good will change. It will either increase by 12 cents (to 42 cents) or decrease by 10 cents (to 20 cents). Both price changes are equally likely.

Period 2

In period 2, you can either sell units of the good at the new price, i.e. 42 cents in case of a price increase and 20 cents in case of a price decrease (as long as you have bought any units in period 1), buy additional units of the good (as long as you still have enough money), or do nothing. Subsequent to your decision, the price will once again rise by 12 cents or decrease by 10 cents. Both price changes are equally likely and independent of the previous price change.

Period 3

As in period 2, you can either buy or sell units of the good in period 3, or do nothing. Afterwards, the price of the good will change for the final time. Price increases (by 12 cents) and decreases (by 10 cents) are again equally likely and independent of all preceding price changes.

Therefore, over all periods, the price can change as shown in the attached figure.

(A figure similar to figure 1 was attached.)

The figures at the “nodes” of the “tree” represent probabilities. As price increases and price decreases are equally likely, the probabilities are always 50 %.

Whether the price of the good increases or decreases at the end of periods 1, 2, and 3 will be determined at the end of the experiment. For each price change a die is rolled.

As you do not know how prices will change during the experiment, we ask you to disclose your decisions for all possible price developments. You have to decide on how many units to buy or sell at the current price at each of the nodes in the attached figure (numbered 1 to 7 and marked in blue). Therefore, you have to make seven decisions.

Subsequent to the experiment, a die will be rolled three times to determine whether the price rises or decreases in periods 1, 2, and 3. If an even number is rolled, the price increases by 12 cents; if an odd number is rolled, the price decreases by 10 cents.

Example: You throw the die three times in a row and obtain 3, 6, and 4. This means the price decreases in period 1 and increases in periods 2 and 3. In this case, only your decisions at nodes 1, 3, and 6 are relevant for your financial payout (see the attached figure).

(A figure similar to figure 1 was attached; nodes 1, 3, and 6 were highlighted in red.)

Your financial payoff is equivalent to the total value of your goods at the end of the experiment plus your cash on hand. Your cash on hand is equivalent to the €12 which you received in period 1, minus your expenditure for buying units of the good, plus your revenues from selling units. You are paid your financial reward in cash at the end of the experiment.

Questionnaire

What did you think when making your decisions in the experiment?

(This was an open question.)

In the past, you purchased a unit of good A and a unit of good B. Since you purchased the goods, the price of good A has increased and the price of good B has decreased. Today, you purchase another unit of one of these goods. Of which of the two goods would you prefer to purchase an additional unit?

- I would prefer to purchase another unit of good A.
- I would prefer to purchase another unit of good B.
- I would be indifferent between purchasing another unit of good A or good B.

In the past, you sold two different goods, A and B. You made neither gains nor losses with these sales. Since you sold the goods, the price of good A has increased and the price of good B has decreased. Today, you repurchase one of these goods. Which of the two goods would you prefer to repurchase?

- I would prefer to repurchase good A.
- I would prefer to repurchase good B.
- I would be indifferent between repurchasing good A or good B.

In the past, you purchased a good. Today, you find out that, since the purchase, the price of the good has decreased. Please distinguish between two different situations:

- (a) You purchased the good although you did not need to. It was your free choice.
- (b) You purchased the good because you urgently needed it. You did not have any other choice.

In which situation would you feel better?

- I would feel better in (a).
- I would feel better in (b).
- I would feel the same in (a) and (b).

In the past, you sold a good. Today, you find out that, since the sale, the price of the good has increased. Please distinguish between two different situations:

- (a) You sold the good although you did not need to. It was your free choice.
- (b) You sold the good because you needed the money. You did not have any other choice.

In which situation would you feel better?

- I would feel better in (a).
- I would feel better in (b).
- I would feel the same in (a) and (b).

Figure 1: Event tree of price changes

The event tree shows all possible price developments over periods 1 to 4. In period 1, the single good starts at a price of 30 cents. At the end of each period, the price either increases by 12 cents or decreases by 10 cents. Price changes are equally likely and independent of previous price changes.

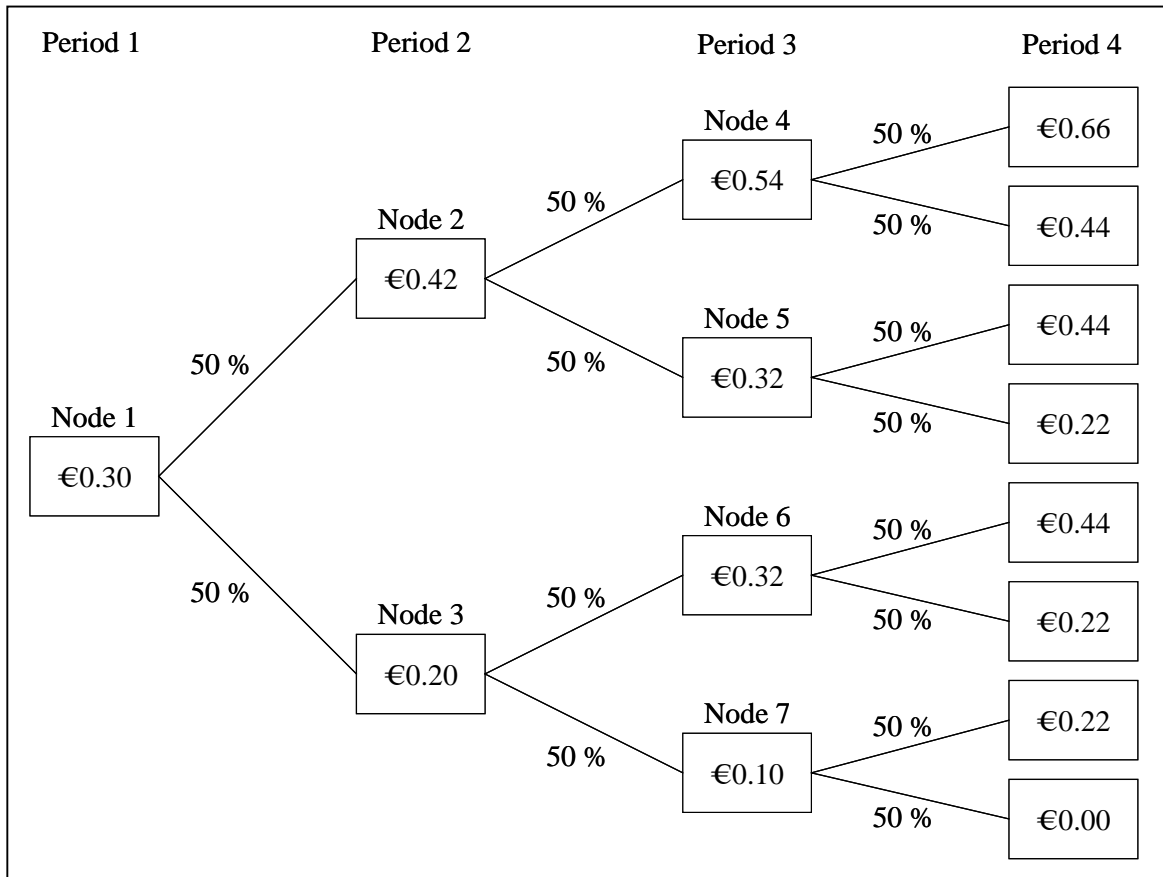


Figure 2: Computer screen

The figure shows a translation of the computer screen. The graph area displays the event tree, with the current node highlighted in red. On the left-hand side, subjects are informed about their current situation. Subjects could decide whether they want to purchase, to sell, or to do nothing by clicking on buttons below the graph area. The screen is taken from a subject that decided to sell and is now asked about the number of units.

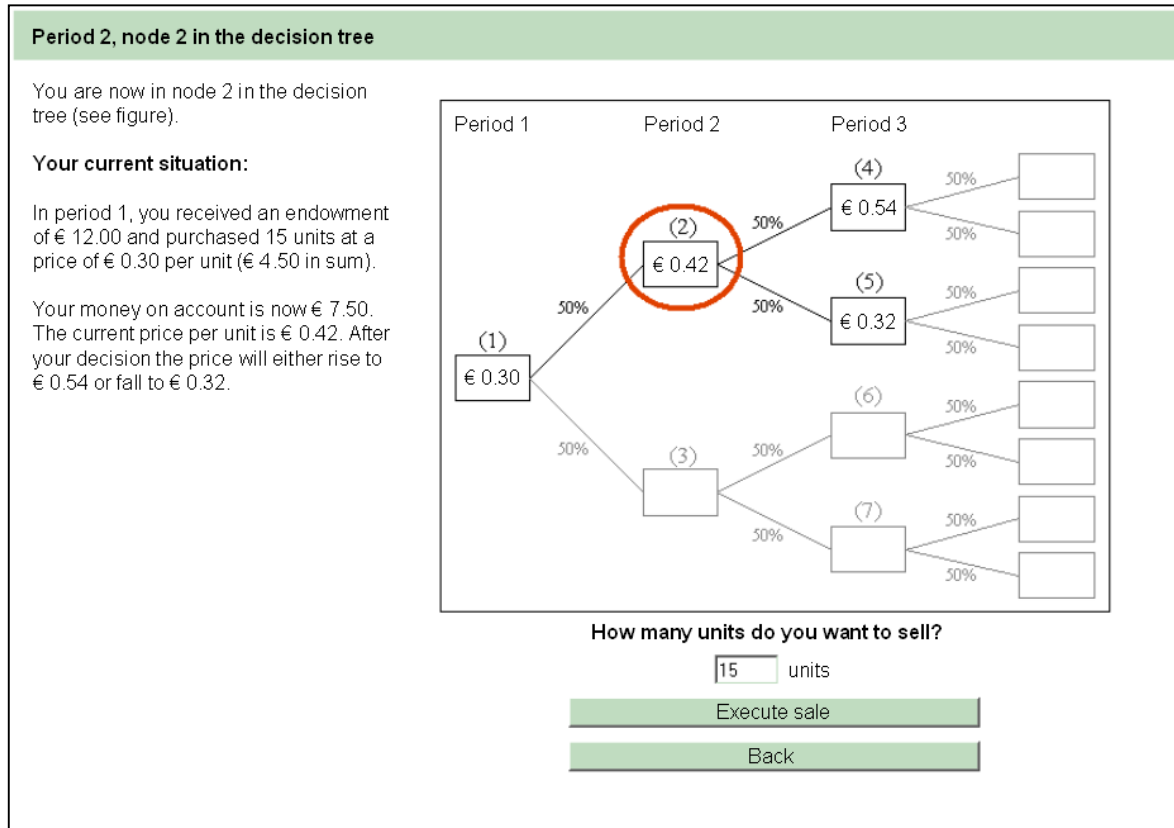


Table 1: Follow-on purchases and counterfactual thinking

The following table is based on our subjects' decisions in period 2, nodes 2 and 3, and period 3, nodes 6 and 7. It reports the number of purchase opportunities, the number of purchases, and resulting purchase frequencies following a price increase (nodes 2 and 6) or decrease (nodes 3 and 7). The first block is based on treatment 1 and documents statistics for all follow-on purchase opportunities, i.e. purchases with a purchase in the preceding period (nodes 1 and 3, respectively). The second block documents follow-on purchases and follow-on purchase opportunities in nodes 2 and 3 of treatment 2 and nodes 6 and 7 of treatment 3, i.e. following an assigned purchase in period 1, node 1, or period 2, node 3, respectively. P-values in the last column are based on binomial tests studying the difference between purchase frequencies following price increases and decreases.

		# of oppor- tunities	# of purchases	Purchase frequency	p
Treatment 1 with purchase in preceding period	following price increase	122	32	26.23 %	0.0000
	following price decrease	123	82	66.67 %	
Treatments 2 and 3 with assigned purchase in preceding period	following price increase	65	27	41.54 %	0.0014
	following price decrease	65	44	67.69 %	

Table 2: Repurchases and counterfactual thinking

The following table is based on our subjects' decisions in period 3, nodes 4 and 5 only. It reports the number of purchase opportunities, the number of purchases, and resulting purchase frequencies following a price increase (node 4) or decrease (node 5). The first block is based on treatment 1 and documents statistics for all repurchase opportunities, i.e. purchases with a sale in the preceding period (node 2). The second block documents repurchases and repurchase opportunities in nodes 4 and 5 of treatment 3, i.e. following an assigned purchase in period 1, node 1, and an assigned sale in period 2, node 2. P-values in the last column are based on binomial tests studying the difference between purchase frequencies following price increases and decreases.

		# of oppor- tunities	# of purchases	Purchase frequency	p
Treatment 1 with sale in preceding period	following price increase	40	7	17.50 %	0.0005
	following price decrease	40	21	52.50 %	
Treatment 3 with assigned sale in preceding period	following price increase	32	14	43.75 %	0.1587
	following price decrease	32	18	56.25 %	

References

- Barber, Brad M. and Terrance Odean (2006): All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *Review of Financial Studies*, forthcoming.
- Blume, Marshall E. and Irwin Friend (1975): The asset structure of individual portfolios and some implications for utility functions, *Journal of Finance*, Vol. 30, No. 2, 585–603.
- Bossaerts, Peter, Charles Plott, and William R. Zame (2006): Prices and allocations in financial markets: Theory and experiment, working paper.
- Giroto, Vittorio, Paolo Legrenzi, and Antonio Rizzo (1991): Event controllability in counterfactual thinking, *Acta Psychologica*, Vol. 78, No. 1–3, 111–133.
- Glaser, Markus (2003): Online broker investors: Demographic information, investment strategy, portfolio positions, and trading activity, working paper.
- Goetzmann, William N. and Alok Kumar (2005): Why do individual investors hold underdiversified portfolios?, working paper.
- Grinblatt, Mark and Matti Keloharju (2001): What makes investors trade?, *Journal of Finance*, Vol. 56, No. 2, 589–616.
- Hirshleifer, David, James N. Myers, Linda A. Myers, and Siew Hong Teoh (2003): Do individual investors drive post-earnings announcement drift? Direct evidence from personal trades, working paper.
- Kahneman, Daniel (1995): Varieties of counterfactual thinking, in: Neal J. Roese and James M. Olson (eds.): *What might have been: The social psychology of counterfactual thinking*, New Jersey, Lawrence Erlbaum Associates, 375–396.
- Kahneman, Daniel and Amos Tversky (1979): Prospect theory: An analysis of decision under risk, *Econometrica*, Vol. 47, No. 2, 263–292.

- Kahneman, Daniel and Amos Tversky (1982): The simulation heuristic, in: Daniel Kahneman, Paul Slovic, and Amos Tversky (eds.): *Judgment under uncertainty: Heuristics and biases*, New York, Cambridge University Press, 201–208.
- Kahneman, Daniel and Dale T. Miller (1986): Norm theory: Comparing reality to its alternatives, *Psychological Review*, Vol. 93, No. 2, 136–153.
- Kroll, Yoram, Haim Levy, and Amnon Rapoport (1988): Experimental tests of the separation theorem and the capital asset pricing model, *American Economic Review*, Vol. 78, No. 3, 500–519.
- Lee, Charles M. C. (1992): Earnings news and small traders: An intraday analysis, *Journal of Accounting and Economics*, Vol. 15, No. 2–3, 265–302.
- Markman, Keith D., Igor Gavanski, Steven J. Sherman, and Matthew N. McMullen (1995): The impact of perceived control on the imagination of better and worse possible worlds, *Personality and Social Psychology Bulletin*, Vol. 21, No. 6, 588–595.
- McMullen, Matthew N., Keith D. Markman, and Igor Gavanski (1995): Living in neither the best nor the worst of all possible worlds: Antecedents and consequences of upward and downward counterfactual thinking, in: Neal J. Roese and James M. Olson (eds.): *What might have been: The social psychology of counterfactual thinking*, New Jersey, Lawrence Erlbaum Associates, 133–167.
- Odean, Terrance (1998): Are investors reluctant to realize their losses?, *Journal of Finance*, Vol. 53, No. 5, 1775–1798.
- Odean, Terrance (1999): Do investors trade too much?, *American Economic Review*, Vol. 89, No. 5, 1279–1298.
- Odean, Terrance, Michal Strahilevitz, and Brad M. Barber (2004): Once burned, twice shy: How naïve learning and counterfactuals affect the repurchase of stocks previously sold, working paper.
- Roese, Neal J. (1997): Counterfactual thinking, *Psychological Bulletin*, Vol. 121, No. 1, 133–148.

- Roese, Neal J. and James M. Olson (1995a): Outcome controllability and counterfactual thinking, *Personality and Social Psychology Bulletin*, Vol. 21, No. 6, 620–628.
- Roese, Neal J. and James M. Olson (1995b): Counterfactual thinking: A critical overview, in: Neal J. Roese and James M. Olson (eds.): *What might have been: The social psychology of counterfactual thinking*, New Jersey, Lawrence Erlbaum Associates, 1–55.
- Selten, Reinhard (1967): Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopol-experiments, in: Heinz Sauer mann (ed.): *Beiträge zur experimentellen Wirtschaftsforschung*, Vol. 1, Berlin, J.C.B. Mohr (Siebeck), 136–168.
- Shefrin, Hersh and Meir Statman (1985): The disposition to sell winners too early and ride losers too long: Theory and Evidence, *Journal of Finance*, Vol. 40, No. 3, 777–790.
- Thaler, Richard H. (1980): Toward a positive theory of consumer choice, *Journal of Economic Behavior and Organization*, Vol. 1, No. 1, 39–60.
- Thaler, Richard H. (1985): Mental accounting and consumer choice, *Marketing Science*, Vol. 4, No. 3, 199–214.
- Tversky, Amos and Daniel Kahneman (1992): Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and Uncertainty*, Vol. 5, No. 4, 297–323.
- Weber, Martin and Colin F. Camerer (1998): The disposition effect in securities trading: An experimental analysis, *Journal of Economic Behavior and Organization*, Vol. 33, No. 2, 167–184.
- Weber, Martin and Heiko Zuchel (2005): How do prior outcomes affect risk attitude? Comparing escalation of commitment and the house money effect, *Decision Analysis*, Vol. 2, No. 1, 30–43.
- Weber, Martin and Frank Welfens (2007): Reference Point Formation over Time: A Weighting Function Approach, working paper.