

Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions

MARKKU KAUSTIA and SAMULI KNÜPFER*

ABSTRACT

We find a strong positive link between past IPO returns and future subscriptions at the investor level in Finland. Our setting allows us to trace this effect to the returns personally experienced by investors; the effect is not explained by patterns related to the IPO cycle, or wealth effects. This behavior is consistent with reinforcement learning, where personally experienced outcomes are overweighted compared to rational Bayesian learning. The results provide a microfoundation for the argument that investor sentiment drives IPO demand. The paper also contributes to understanding how popular investment styles develop, and has implications for the marketing of financial products.

WHAT CAN AN INVESTOR LEARN from personal experimentation with stocks? For experience goods such as wine, experimentation is essential: Product information and reviews can help to some extent, but ultimately one must consume the product to evaluate the utility it gives (see Nelson (1970)). Experimentation is also important when it is a prerequisite for acquiring information (e.g., Grossman, Kihlstrom, and Mirman (1977)). However, stocks lack the aforementioned characteristics, at least in traditional thinking. Therefore, whether one personally invests in a stock or merely observes the performance of a stock should not matter.

The impact of personally experienced investment outcomes on future actions tells us something important about how investors learn, as it represents a key distinction between fully rational models of learning and reinforcement learning. Reinforcement learning theory predicts that personally experienced outcomes have a greater effect on behavior than, say, just reading about the same information without personal involvement. In pure choice reinforcement,

*Markku Kaustia and Samuli Knüpfer are at Helsinki School of Economics, Department of Accounting and Finance. The paper benefited greatly from comments by Campbell Harvey (the editor) and the anonymous referee. We thank Malcolm Baker, Matti Keloharju, Mikko Leppämäki, Juhani Linnainmaa, Darius Palia, Vesa Puttonen, Elias Rantapuska, Sami Torstila, Dimitri Vayanos, Jeffrey Wurgler, and especially Jay Ritter for comments. We also thank the participants at the 2008 American Finance Association meeting in New Orleans, 2007 European Finance Association meeting in Ljubljana, and 2007 Financial Management Association European meeting in Barcelona, as well as at London School of Economics, London Business School, UCLA, University of Michigan, INSEAD, and New York University seminars. Both authors acknowledge financial support from the Academy of Finland and the Finnish Foundation for Advancement of Securities Markets. Knüpfer acknowledges as well support from the Helsinki School of Economics Foundation, Kemira Foundation, and Okobank Group Research Foundation.

only actual and directly experienced outcomes affect future choices, while in standard Bayesian belief-learning, actual, and forgone payoffs are equally important.

Theoretical research has paid much attention to reinforcement learning (e.g., Cross (1973), Arthur (1991), Ellison and Fudenberg (1993), Roth and Erev (1995)). Camerer and Ho (1999) develop a hybrid model that combines reinforcement learning and belief learning. This allows the influence of both actual and forgone payoffs, but possibly with different weights. In the field of financial economics, explanations for asset pricing anomalies have involved different forms of biased learning,¹ although there are also other theoretical explanations.² Evidence on whether, and how, investors learn has implications for this discussion.

Very little is known empirically about how investors learn. Recent research suggests that more experienced investors make fewer behavioral errors and use more sophisticated trading tactics (Feng and Seasholes (2005), Dhar and Zhu (2006), Korniotis and Kumar (2006), Seru, Shumway, and Stoffman (2006)). However, Linnainmaa (2006) argues that investors with less skill learn to exit the stock market. This could explain an *ex post* positive correlation between experience and sophisticated behavior.

In this paper we examine empirically what, if any, incremental forecasting power personally experienced investment outcomes have for future actions. The setting is unique: the Finnish initial public offering (IPO) market, with data on allocations of all individual investors for all 57 IPOs that occurred from January 1995 through December 2000. The actions are subscriptions to the offerings, and the outcomes are determined by IPO stock performance.

There are several reasons this setting is suitable for the research question at hand. First, it is an advantage to have data on all IPOs in the market and at the same time not have too many IPOs. Investor attention is likely to be much more evenly distributed over the 57 IPOs than over some broader set of investment opportunities, such as all common stocks. Further, IPOs typically receive a fair amount of publicity in the news media (Cook, Kieschnick, and Van Ness (2006)) and the investment outcome is thus easily observable at very

¹ Barberis, Shleifer, and Vishny (1998) model a representative investor who uses standard Bayesian updating to switch between two belief regimes for corporate earnings. The true regime is a random walk, but the agent is constrained to falsely believe earnings are either mean-reverting or trending. In Daniel, Hirshleifer, and Subrahmanyam (1998) the informed investor is Bayesian with respect to forecasts of asset value, but underweights signals that are contrary to his earlier private signals. Cecchetti, Lam, and Mark (2000) model investors who use simple learning rules implying excessive optimism and pessimism about the duration of business cycles. In Gervais and Odean (2001), the investor learns about his ability but underweights the role of chance in assessing performance. Brav and Heaton (2002) model the learning of one agent who considers only recent payoffs and ignores prior beliefs and another agent who overweights prior beliefs. Peng and Xiong (2006) model a representative investor with limited attention and overconfidence, but Bayesian updating in processing information.

² Campbell and Cochrane (1999) and Barberis, Huang, and Santos (2001) present models with nonstandard preferences that can account for these anomalies. Brav and Heaton (2002) show that the same aggregate outcome can result from biased learning or from rational structural uncertainty.

little cost. Second, IPOs are temporally discrete events: Although many IPOs can be offered simultaneously, often there is sufficient time between offerings for an investor to condition his/her future decisions on past outcomes. Third, a new IPO comes fresh to the market and no outside investors have had any previous experience with the stock. Fourth, short-term liquidity shocks also probably have a somewhat reduced impact on IPO participation relative to trading of listed stocks because of the fixed timing of the offering periods. Taken altogether, these characteristics contribute to a setting that is much closer to a natural experiment in sequential decision making than is usually the case in a study of stock market investments.

Our main contributions are an empirical test of the reinforcement learning hypothesis and an assessment of its strength in financial markets. This study constitutes a large-sample field application, while previous empirical evidence on reinforcement learning comes from experimental studies. In field applications, subjects' beliefs are generally unobservable, and thus it is not possible to test directly for optimal action conditional on the subjects' information sets. We are able to avoid this problem by focusing instead on the impact of personal experience, as it provides a testable distinction between reinforcement learning and fully rational learning.

We find that personal experience exerts an incrementally significant impact on future action. Controlling for IPO fixed effects, which encompass the effect of offer characteristics such as size and preoffer publicity, but also market returns and previous IPO returns, we find personally experienced returns to be an important determinant of future IPO subscriptions. The impact of personally experiencing good returns, rather than merely observing them, is consistent with reinforcement learning.

We also identify pairs of hot and cold IPOs with offer dates very close to each other and track the future subscription activity of investors who make their first IPO subscription in one of these offerings.³ More than twice as many investors participate in a subsequent offering if they first experience a hot offering rather than a cold offering. The effect continues to be strong for subsequent offerings as well. By the tenth offering, 65% of investors in the hot IPO group will have subscribed to another IPO, compared to only 39% in the cold IPO group. The importance of initial experience has long been recognized in marketing, consistent with what psychologists refer to as the *primacy effect* (Asch (1946); see also Bereby-Meyer and Roth (2006) for an analysis of related issues in strategic games). Our results can be viewed as an illustration of the primacy effect in the realm of finance and it may have implications for the marketing of financial products and services.

In addition to reinforcement learning, there are other theories that could also potentially explain our results. These include private information, portfolio rebalancing, wealth effects, liquidity constraints, investment banks' preferential treatment of customers, and factors related to the hot issue market

³ A cold IPO would have negative initial returns, while a hot one would have strictly positive initial returns.

from 1999 to 2000. However, these explanations do not stand up to further scrutiny.

There is a natural connection between our results and the literature on investor sentiment, which is suggested to be a determinant of IPO waves and pricing patterns, and various other anomalous asset pricing regularities.⁴ There is no exact definition of sentiment, which is variably referred to as, for example, the average bullishness of noise traders (De Long et al. (1990)), or the propensity to speculate (Baker and Wurgler (2006)). In typical models sentiment plays the role of an external factor. Empirically, it seems to be positively related to investor and consumer confidence, as well as past market returns.⁵

While the confidence of a given investor can also depend on factors beyond the individual, our results provide a natural way of interpreting sentiment through investor learning based on personal experience. Good experiences bolster confidence. When such experiences by many investors coincide, there may be market-level consequences, as shown theoretically in the context of IPOs by Derrien (2005) and Ljungqvist, Nanda, and Singh (2006), and also consistent with the empirical results of Kumar and Lee (2006).

Our results have further implications for understanding investor behavior, tastes, and cohort effects. According to anecdotal reports, even institutional investors want to “see for themselves” how a new asset class or a trading strategy might work for them, and they invest more money in it if they are satisfied with its performance. Particularly inexperienced money managers may be more prone to overweight recent events, as shown by Greenwood and Nagel (2008). Results on how investors learn provide a way to understand these effects better.

Fama and French (2007) assume that investors have different tastes for different kinds of stocks. Our results suggest that experience with particular stocks is a possible channel for forming these tastes. Cohort effects in asset market participation, such as the baby boom effect (see Bakshi and Chen (1994)), have large and long-lasting implications. Campbell (2006) proposes that intergenerational differences in how stock markets are viewed and in the use of different financial products can be driven by the different experiences of different generations. Our results are also consistent with this idea.⁶

In the remainder of the paper, Section I discusses reinforcement learning and the role of personal experiences. Section II describes the data and presents the main results, while Section III considers several alternative explanations and robustness checks. Section IV concludes.

⁴ See De Long et al. (1990) for a seminal sentiment model; Derrien (2005) and Ljungqvist, Nanda, and Singh (2006) for models of IPO markets and sentiment; Lee, Shleifer, and Thaler (1991), Rajan and Servaes (1997), Lowry (2003), and Cornelli, Goldreich, and Ljungqvist (2006) for empirical evidence on the role of sentiment in IPOs; Baker and Wurgler (2006) for the role of sentiment in the cross-section of stock returns; Fisher and Statman (2000) and Brown and Cliff (2004, 2005) for predicting market returns with sentiment; and Shleifer (2000), Hirshleifer (2001), and Baker and Wurgler (2007) for surveys.

⁵ See Fisher and Statman (2003), Brown and Cliff (2004), Lemmon and Portniaguina (2006), Baker and Wurgler (2007), Qiu and Welch (2006), and Ben-David, Graham, and Harvey (2007).

⁶ A recent paper by Malmendier and Nagel (2008) also finds support for this idea.

I. Reinforcement Learning and the Role of Personal Experience in Financial Decisions

Reinforcement learning is the leading theory of heuristic updating. It posits that decision makers tend to repeat actions that have produced favorable outcomes in the past. An example of a decision rule consistent with reinforcement learning is “win–stay, lose–switch,” which can be applied to sequential decisions when the outcome is binomial, or when the agent treats a more complex outcome simply as success or failure depending on the outcome exceeding his aspiration level. The theory of reinforcement learning is well established in the psychological literature, and dates back at least to the “law of effect” of Thorndike (1898). Reinforcement learning models in the economics literature include Cross (1973), Arthur (1991), and Roth and Erev (1995). Ellison and Fudenberg (1993) combine reinforcement learning with social interaction. The case-based decision theory of Gilboa and Schmeidler (1995) also bears some resemblance to these ideas. Camerer and Ho (1999) develop a hybrid model that combines reinforcement learning and belief learning, nesting each as special cases. Their model thus allows for the reinforcement of both actual payoffs and forgone payoffs. Theoretical research also shows that reinforcement learning often has long-term equilibrium implications (Einhorn (1980), Lettau and Uhlig (1999), Pouget (2007)).

The experimental literature provides support for the reinforcement hypothesis. Erev and Roth (1998) predict the behavior of subjects in 12 different experimental data sets involving strategic games with a unique equilibrium. They find that a very simple one-parameter reinforcement learning model outperforms equilibrium predictions for all values of its one parameter. Camerer and Ho (1999) test their hybrid model experimentally and find that subjects generally weight actual payoffs about twice as heavily as forgone payoffs. Charness and Levin (2005) set up an experiment so that Bayesian updating and reinforcement learning lead to different choices. They find that about 50% of all decisions violate Bayes’s rule.

Psychological and neurological mechanisms suggest a special role for personal experiences in learning. The cognitive processing of personal experiences can be inherently different and more thorough (Fazio, Zanna, and Cooper (1978)), which can make them more salient and available in decision making (Tversky and Kahneman (1973)). Kahneman, Wakker, and Sarin (1997) develop a theory of choice in which the utility of an outcome considered in decision making (“decision utility”) is allowed to be different from the utility of that outcome when experienced firsthand (“experienced utility”).⁷ Simonsohn et al. (2008) provide experimental evidence on the incremental impact of personal experience in the context of strategy games.

Recent neurological studies further support the special role of personal experience. Knutson and Peterson (2004) show that different physiological mechanisms and even different parts of the brain are involved when subjects

⁷ See also Kahneman (1994, 2000) for discussion and experimental evidence.

anticipate, rather than experience, monetary outcomes. Huettel, Mack, and McCarthy (2002) find that a specific part of the brain starts to anticipate another repetition of the same outcome in a sequence after it occurs only twice in a row, and another part of the brain reacts strongly if a repeating pattern is broken. This finding is consistent with the idea that even very limited personally experienced evidence can cause choice reinforcement.

There is also evidence that personally experienced payoffs could actually interfere with rational learning. Merlo and Schotter (1999) run an experiment to study how well subjects learn an optimal decision rule. Their results show that subjects who had monetary incentives in earlier rounds later performed significantly more poorly. Schotter and Merlo (2003) find that subjects who learn by making actual decisions perform significantly more poorly than subjects who learn by merely observing another person's decisions and their outcomes. Ariely et al. (2008) present experimental evidence that monetary rewards that are "too high" (on the order of average monthly consumption for subjects in rural areas of India) degrade performance in cognitive tasks. Finally, Charness and Levin (2005) find that when a prior outcome provides information along with a payoff, violations of Bayes's rule occur much more often than when the same information is provided without a payoff.

Based on the discussion above, the key prediction of reinforcement learning theory relevant for us is that personal experience can be effective in changing behavior. In our setting, this may be expressed in a hypothesis as follows: Personally experienced IPO returns are incrementally positive determinants of future subscriptions.

II. Data and Main Results

A. Data

To analyze the impact of personally experienced returns on future action, we construct a data set of retail investor subscriptions to initial public offerings and link them to past returns. The subscription decisions are derived from a nationwide share registry maintained by the Finnish Central Securities Depository (FSCD). This official registry includes shareholdings and transactions for all investors in Finland. The FSCD also records allocations to investor accounts in equity offerings. The data do not cover foreign IPOs. However, based on discussions with practitioners and the Finnish Financial Supervision Authority, foreign IPOs represent a negligible share in individual investors' portfolios.

Except for one IPO, all investors who had bid for shares received an allocation. This means that using allocations to measure participation does not produce any measurement error, except for that one single IPO.⁸ In other respects, the institutional environment in the Finnish IPO market during the sample period is quite similar to that in the United States. All issues are underwritten using the bookbuilding method, for example. One important difference is that investment banks are required to treat all retail investors equitably.

⁸ In the IPO of SSH Communications, a lottery was used to determine the recipients of the shares and about 71% of subscribers received an allocation.

Every allocation is tagged with a unique investor code, and all the future actions an investor takes can be tracked. Each investor code is also linked to investor characteristics, and we use investors' age and sex as control variables. The registry records all other stock market transactions, particularly stock market trades. We use these data for control variables such as portfolio value, and for constructing overall measures of investment performance. Further details on a subset of these data are provided in Grinblatt and Keloharju (2000).

The total sample over January 1995 through December 2000 includes 57 offerings. The next IPO after the sample period occurs in March 2002. Data on IPO characteristics are obtained from prospectuses and stock exchange releases. These characteristics include listing date, beginning of the retail subscription period, offer price, and gross proceeds. Stock and index return data are provided by the OMX Helsinki Exchanges.

Table I, Panel A, gives descriptive statistics of the sample IPOs. Average first-day underpricing is 22.3%, which is high by historical standards but lower than 32.7%, the figure for U.S. IPOs over the same period (Loughran and Ritter (2004)). The median return for the sample IPOs is 2.8%. Panel B shows averages and the number of IPOs for each year during the sample period. The number of IPOs and first-day returns are higher during 1999 and 2000, in line with the global market conditions at the time. Panel C summarizes investor characteristics for the sample of investors who make at least one IPO subscription during the sample period. The average investor makes 2.0 IPO subscriptions; the mean value of an allocation is about EUR 1,700. For the average investor, 33 IPOs occur after the first stock market investment; the median figure is 37. Average portfolio value is about EUR 65,000, but the median is much lower at about EUR 9,000. The distribution of trading activity is even more skewed. The average turnover (value of buys plus value of sells) is about EUR 286,000, and the median is about EUR 19,000. The average age of the investor is 46 years, and 66% are men.

Some of our empirical tests require the identification of the first IPO subscribed to by an investor. Because we do not have data on IPO subscriptions before the beginning of 1995, there is some measurement error as to whether a subscription is an investor's first. Only four IPOs took place in the 5 years before the beginning of our study period, however, and the offerings were quite small. Therefore, for most investors we are able to accurately measure the first IPO in 5 or more years.

We measure IPO stock performance by initial returns (first-day closing price/offer price -1). Of course, it is possible to define performance over some longer holding period, but we want to keep the holding period relatively short to be able to include the return from the latest IPO as we consider the impact on subscription. Returns over, for example, the first week of trading would allow this in most cases, but as the initial return is a permanent component of the stock price (at least over a horizon of a few months), it is very highly correlated with the returns over other short-term horizons. Initial returns are also highly visible and easily comparable.

An alternative way to think of IPO stock performance would be to calculate the returns actually realized by investors. Research on the disposition effect,

Table I
Descriptive Statistics

This table gives descriptive statistics of the IPOs and investors in the sample. Panel A shows statistics for Finnish IPOs from 1995 to 2000. Underpricing is calculated as the difference between the first-day closing price and offer price divided by the offer price. Number of investors equals the number of retail investors who were allocated shares in the offering and gross proceeds equals the number of allocated primary and secondary shares to retail and institutional investors times the offer price, before the use of any green shoe options. Panel B shows averages of the same variables and the number of IPOs for each year during the sample period. Panel C shows statistics for investors who participate in at least one IPO during the sample period. Allocation value equals the number of shares allocated to an investor times the subscription price. Portfolio value is calculated as the sum of the value of holdings in individual shares and number of shares as the number of individual shares in the portfolio. Portfolio characteristics are calculated for each investor in the beginning of the subscription period of each offering starting from the first IPO investment, averaged across offerings, and then across investors. Trading activity variables and age come from the beginning of the subscription period of the last IPO in the sample.

Panel A: IPOs, $N = 57$						
Variable	Mean	Median	SD	25% Percentile	75% Percentile	
Underpricing	22.27%	2.80%	53.51%	0.00%	20.43%	
Number of investors, Th.	7.78	1.97	15.37	1.16	5.24	
Gross proceeds, Mill. euros	86.29	27.18	200.03	10.21	0.08	
Panel B: IPOs by Year, $N = 57$						
	1995	1996	1997	1998	1999	2000
Average underpricing	0.77%	17.80%	14.41%	14.41%	35.83%	21.23%
Average number of investors, Th.	5.10	0.31	1.96	18.98	7.12	7.62
Average gross proceeds, Mill. euros	77.82	13.21	53.79	311.62	49.32	42.84
Number of IPOs	5	2	8	8	19	15
Panel C: Investors, $N = 183,919$						
	Mean	Median	SD	25% Percentile	75% Percentile	
No. of IPOs after the first subscription	32.88	37.00	13.32	18.00	38.00	
Number of IPO subscriptions	2.04	1.00	2.47	1.00	3.00	
Total value of allocations, Th. euros	3.86	1.35	66.38	0.19	4.64	
Value of allocation, Th. euros	1.72	1.13	33.16	0.52	2.64	
Portfolio value, Th. euros	64.67	9.20	1,096.12	2.72	29.10	
Number of shares in portfolio	3.60	2.57	3.65	1.18	4.67	
Value of stock market buys, Th. euros	133.29	7.57	1,976.98	0.50	37.32	
Value of stock market sells, Th. euros	152.34	11.89	2,054.75	2.34	49.61	
Age	46.34	47.00	17.84	33.00	59.00	
Gender dummy	0.66	1.00	0.47	0.00	1.00	

however, shows a strong tendency for investors to hold on to losers and to realize gains (see Odean (1998)). Using realized returns would thus produce a sample that would be self-selected according to selling, and hence the level of returns.

B. Impact of Prior Performance on Subscription

Activity: A Two-Period Analysis

Our investigation of the impact of prior performance uses three types of analyses. The first one divides the sample into two halves, and investigates how actions in the second half are affected by returns earned in the first period.

The first half of the sample includes 40 IPOs, the last occurring in October 22, 1999. The second half includes 17 offerings, the last in December 20, 2000. This division places an equal number of investor/offering pairs in each period. The median investor makes only one IPO subscription during the whole period, so a majority of investors are very passive. We therefore exclude investors who participate in fewer than three IPOs in the first half, because including them would significantly reduce the variation in the variable that measures personally experienced returns.

In Figure 1 investors are grouped into quintiles based on allocation-weighted initial returns. The columns correspond to the probabilities for ending up in a




Figure 1. Transition probabilities between previous and subsequent initial return quintiles. To obtain the figure, the sample period is divided into two halves. The first half has 40 IPOs, the last occurring in October 22, 1999; the second half has 17 offerings, the last in December 20, 2000. This split is determined by placing an equal number of investor/offering pairs in both periods. Investors who participate in at least three IPOs in the first half are included. Initial return is the average initial return from the IPOs in which the investor participates, weighted by the value of allocation.

particular return quintile during the second half, given the return quintile during the first half. Investors who make no subscriptions during the second half are placed in a category labeled “No subscriptions.” Conditional on at least one subscription made in the second half, returns in the second half do not seem to depend on the returns in the first half. There is, however, a strong relation between returns in the first half and the likelihood that an investor will subscribe to at least one IPO also in the second half. The likelihood of ending up in the category “no subscriptions” in the second half is about 20% for an investor in the highest return quintile during the first half, whereas this likelihood is about 50% in the group with the poorest performance. The proportions for the intermediate quintiles fall between these ends in a monotonic fashion.

To control for investor characteristics that affect subscription activity, we run logit regressions that explain the subscriptions in the latter period. The dependent variable in the regressions is assigned the value of one for an investor who participates in at least one IPO in the second half, and zero otherwise. The independent variable measuring past returns is the allocation-weighted initial return from the IPOs subscribed to in the first half.⁹ The activity levels of individual investors are highly persistent, that is, investors with only few subscriptions in the past tend to make only few in the future, and those with high past activity tend to be active in the future as well. To control for this tendency, we include a variable measuring an investor’s number of subscriptions in the first half. As other control variables we include quintile dummies for portfolio value and trading activity, as well as age, age squared, and a dummy for being male.

Table II presents the results of the subscription logit regressions. Column 1 does not include the trading activity or the portfolio value variable; columns 2 and 3 add either one of them; and column 4 includes both variables. In all the specifications, the percentage return on past IPO subscriptions of an investor is a strong determinant of future subscription activity. Past subscription activity is also a significant positive predictor of future subscription activity.

To assess the economic significance of our results, the table reports changes in subscription probability associated with changes in the explanatory variables. The probability changes are based on the regression in column 4 and are the products of the marginal effect of an explanatory variable and its standard deviation. The marginal effect is obtained by setting all the continuous explanatory variables in the model to their mean values, and the category dummies are set to indicate the omitted category. For the return variable, a change from its mean, 17.9%, by one *SD* (8.0 percentage points) implies a 7.7 percentage point increase in the probability of subscribing to future IPOs. This increase is considerable, given that the average investor’s probability of participating is 70%.

Figure 2 graphs the relation between past returns and future subscriptions by dividing investors into 25 categories based on five past return and five

⁹ Changing the definition of this variable to a simple average initial return without allocation weighting does not change the results materially.

Table II
Two-Period Analysis of the Relation between Past Returns
and Future IPO Subscriptions

The sample period is divided into two halves. The first half has 40 IPOs, the last occurring in October 22, 1999; the second half has 17 offerings, the last in December 20, 2000. This split is determined by placing an equal number of investor/offering pairs in both periods. We include investors who participate in at least three IPOs in the first half. Regressions are run using a logit model. The dependent variable is one if an investor participates in at least one IPO in the second half, zero otherwise. Past mean initial return is the average initial return from the IPOs in which the investor participates, weighted by the value of allocation. Unreported control variables are age, age squared, a dummy for being male, quintile dummies for portfolio value, and quintile dummies for value of trades (summed from the beginning of the sample period). Columns 1 through 4 report coefficient estimates of regressions. Column 5 reports the changes in probabilities associated with a one *SD* increase in an explanatory variable, evaluated at the mean of all other variables. Probability of reference investor is the model's predicted probability to participate, evaluated at the mean of all the explanatory variables; the male dummy set to zero; and the omitted quintile dummy set to zero for both portfolio value and value of trades. *t*-values (in parentheses below coefficients) are based on robust standard errors.

Specification	Logit (0 = No Future Subscriptions, 1 = At Least One Future Subscription)				Reference Probability and Probability Changes
	1	2	3	4	
Probability of reference investor					70.4%
Past mean initial return	5.41 (25.50)	4.67 (21.68)	4.67 (21.72)	4.65 (21.56)	7.7%
Number of past IPO subscriptions	0.55 (37.56)	0.47 (32.29)	0.47 (32.29)	0.47 (32.14)	24.0%
Portfolio value dummies	No	Yes	No	Yes	
Trading activity dummies	No	No	Yes	Yes	
Pseudo- R^2	0.21	0.24	0.24	0.24	
Number of observations	21,928	21,928	21,928	21,928	

subscription activity categories. Return categories are quintiles of allocation-weighted initial returns, and subscription groups are defined as one through five or more subscriptions.¹⁰ Every category (except one) is assigned a dummy variable, and the dummy variables are included as independent variables in a regression similar to that in Table II, column 4. The plot shows a marked increase in subscriptions as a function of the return experienced within all activity categories, with a somewhat stronger impact of returns for investors with more prior subscriptions.

Figure 2 also gives us a way to compare the extent of the return effect to the effect of past subscription activity. In most cases, higher past returns produces about the same effect on future activity as higher past activity. For example, consider an investor located in the middle quintile (the third) on returns, with

¹⁰ This specification also includes the passive investors with only a few subscriptions that were not included in the baseline results in Table II.




Figure 2. Likelihood of further participation by past subscription activity and past initial returns. This figure is based on a logit regression similar to that in Table II column 1, that is, the dependent variable is one if an investor participates in at least one IPO in the second half, zero otherwise (see more description in Table II header). Unlike in Table II, here previous IPO activity and initial return are divided into 25 (5×5) dummy variables each representing a pairing of one activity and one return interval. Initial return is the average initial return from the IPOs in which the investor participates, weighted by the value of allocation. Activity is divided into five intervals based on the number of IPO subscriptions, from one to five or more. Intervals for returns are based on quintiles formed on the distribution returns. The figure depicts the estimated coefficients for these dummy variables. The group with one subscription and the lowest return during the first half is the default, and is plotted as a zero.

three past subscriptions. Changing to the fourth quintile of returns (holding the number of past subscriptions constant) produces about the same effect as changing to four subscriptions (holding performance constant).

We examine other specifications in order to check the robustness of the results. First, we use an alternative way to divide the time period into two halves. Instead of an equal number of investor/offering pairs in both periods, we divide the sample by placing an equal number of offerings in both periods. This produces a much earlier cutoff date, which allows us to check possible effects related to the state of the IPO cycle. The results are very similar to our first set of results. We prefer the original analysis, however, as the alternative split produces a substantial imbalance in the number of observations—there are many more observations in the second half. In additional robustness checks, rather than defining the dependent variable by whether the investor makes subscriptions in the second period, we define it based on the difference in subscriptions and do not include the number of past subscriptions as an explanatory variable. The dependent variable in this specification is coded as one if the number of subscriptions increases from the first period, and zero otherwise. We also run an

ordered logit specification, where the dependent variable indicates a decline, no change, or an increase in subscriptions. These modifications do not markedly alter our conclusions.

C. Prior Performance and Timing of the Next Subscription

In this subsection we take the analysis to the level of investor-offering pairs, and analyze how an investor's future subscription after a particular offering is affected by her returns experienced up to that offering. Returns vary across offerings and investors, and this research design allows us to control for any offering-level effects.

We conduct this analysis because retail investor participation activity in an offering is affected by a number of factors that we cannot effectively control for in the two-period analysis. For example, investors have been shown to participate more after higher industry returns (Derrien (2005)) and to participate more in offerings with higher preoffer publicity (Cook et al. (2006)). Offering effects also control for effects related to offer size, distribution of the offering, and general market returns.

Even more important to our research question is the fact that the effects also include the impact of returns from recent IPOs in which the investor did not participate. Ibbotson and Jaffe (1975) and Lowry and Schwert (2002), among others, show that initial returns are autocorrelated, and this might induce investors to participate without regard to personal returns. Therefore, this test allows us to separate more cleanly an investor's personal return experience from information about the IPO cycle. This analysis also works to alleviate any possible remaining concern that the stage of the IPO cycle somehow drives the results through the placement of the cutoff point between the two periods in the previous section's analysis.

Our empirical specification involves a logit regression that analyzes an investor's decision to participate in the next offering, conditional on having participated in the current offering. The dependent variable takes the value of one if an investor participates in the next IPO, and zero otherwise. We limit the analysis to investors who have previously participated in at least three IPOs and calculate the independent variable, the average initial return from the IPOs subscribed to, over the past three offerings. As control variables we employ quintile dummies for portfolio value and trading activity, dummies for the past number of subscriptions (from 3 to 10 or more), age, age squared, a dummy for male investors, and a dummy variable for each offering.

The sequence of IPOs is determined by subscription periods and listing dates. In the beginning of the subscription period for each IPO, we designate the IPO with the most recent listing date as the prior IPO. Sometimes two or more IPOs have the same prior IPO. As an example, consider the case of three IPOs, labeled *A*, *B*, and *C*. IPO *A* is listed at $t = 0$. Subscription to *B* begins at $t = 1$ and *B* is listed at $t = 3$. Subscription to *C* starts at $t = 2$, that is, before *B* is listed. In this case *A* would be the prior IPO for both *B* and *C*. In such cases, that is, where two or more IPOs share a common prior IPO, we randomly select

one and exclude observations from the other simultaneous IPOs. We do this to avoid double counting of investor-offering observations, and believe this is the most conservative approach to dealing with the issue. As a result, the number of IPOs used in this analysis is reduced to 30.

Table III, Panel A, reports the results of the logit regression in column 1. A higher past mean initial return strongly predicts a higher probability of

Table III
Timing of Participation to Next IPO as a Function of Past Returns

This table gives results of investor-level regressions where the dependent variable is either an indicator variable representing the decision to subscribe to the next IPO (column 1 of Panel A) or the time to the next IPO subscription (column 2 of Panel A). For every investor participating in at least three IPOs during the sample period from 1995 to 2000, the time to the next IPO subscription is determined by recording how many IPOs elapsed until the investor participates in the next IPO, including the next IPO. The dependent variable is regressed on a set of explanatory variables using a logit model in column 1 and a duration model with an exponential duration distribution in column 2. The hazard ratio reported in column 2 is given by $\exp$ (coefficient obtained from of the duration regression). The independent variable is the average initial return from the previous three IPOs subscribed. Unreported control variables are age, age squared, a dummy for being male, quintile dummies for portfolio value, quintile dummies for value of trades (summed over the previous 120 trading days), dummies for the past number of subscriptions (from three to ten or more), and dummies for each offering. “Probability change” shows the changes in probabilities associated with a one *SD* increase in an explanatory variable, evaluated at the mean of all other variables. “Hazard ratio change” is the change in hazard ratio associated with a one *SD* increase in an explanatory variable. *t*-values (in parentheses below coefficients) are based on robust standard errors.

Panel A: Regression Results				
Specification	1		2	
	Decision to Subscribe to Next IPO		Time to Next Subscribed IPO	
Dependent Variable	Coefficient	Probability Change	Hazard Ratio	Hazard Ratio Change
Past mean initial return	0.47 (18.78)	6.90%	0.43 (24.32)	33.70%
Dummies for				
Offerings		Yes		Yes
Portfolio value quintiles		Yes		Yes
Trading activity quintiles		Yes		Yes
Subscription activity		Yes		Yes
Number of observations	110,501		115,040	
Panel B: Descriptive Statistics				
	All Observations		At Least Three Subscriptions	
Number of observations	321,394		115,040	
Number of failures	191,195		77,457	
Time to next IPO, mean	12.61		7.26	
Time to next IPO, median	11		3	

participating in the next IPO. The relation is both statistically and economically significant. The probability increase implied by the marginal effect for a one *SD* change in the return variable equals 6.9 percentage points, while the reference probability is 31.3%. The relative increase in probability is higher compared to the two-period analysis reported in Table II. The observation frequency is higher as well, as we consider each offering separately. This brings down the reference probability and amplifies the effect of return as the analysis is conditional on participating in the most recent IPO.

The logit specification ignores any information beyond participation in the next offering. We use this additional information by modeling the time to the next IPO as a duration model instead of the dichotomous variable measuring the decision to participate in the next offering. The analysis takes into account censoring at the end of our sample period. The independent variables are the same as in the logit model.

Table III, Panel B, presents descriptive statistics of the duration analysis. Duration is measured in units of IPOs. One unit of “time” is the period from one offering to the next, so one unit does not correspond to a constant amount of calendar time. The average number of offerings until the next subscription is 12.6 offerings for the whole sample and 7.3 for the sample of investors with at least three offerings.

Table III, Panel A, reports the results of the duration regression in column 2, estimated using an exponential duration distribution.¹¹ The past mean initial return is associated with an increase in the hazard rate, suggesting again that higher returns increase the investor probability of participating in future offerings. The effect is economically significant; a one *SD* increase in return (68.1%) implies about an $e^{0.43 \times 0.68} - 1 = 33.7$ percentage point higher rate of future participation.¹²

We also calculate standard errors that correct for clustering at the offering level. The *t*-value for the return coefficient equals 2.59 in the duration model, and 1.86 in the logit model. Note that these numbers effectively correspond to an analysis with 30 observations. Hence the *t*-values should be interpreted accordingly.

D. Natural Experiment

In a third test of the impact of personal experience, we run the following experiment. We identify six pairs of IPOs for which subscription periods overlap, and for which it happens to be the case that the initial return is positive for one of the IPOs and negative for the other one. The mean difference between

¹¹ This is a standard assumption, and we have no *ex ante* reason to assume a particular distribution. Our results are not very sensitive to the choice of the distribution, however. For example, using a Weibull distribution results in a coefficient value of 0.40 on the mean initial return variable.

¹² The *SD* of the personally experienced return (68.1%) is similar in magnitude to the *SD* of returns calculated over the IPOs. For comparison, the standard deviation of U.S. IPO returns was 55% over 1964 to 2004, and well over 100% from September 1998 to August 2000 (Lowry, Officer, and Schwert (2006)).

the initial returns within a pair of IPOs is 42 percentage points. We select investors who make their first IPO investment in one of these pairs, subscribing to either the positive- or negative-return IPO (but not to both). We then track the future subscriptions of these investors. We also single out another subsample of investors who enter the stock market by subscribing to one of these IPOs.

The benefit of this design is twofold. First, it allows for investigation of investors with identical IPO market histories, that is, those with no prior subscriptions, and, in the case of new investors, no stock market history whatsoever. This makes it close to a natural experiment. Second, it allows us to analyze the behavior of typical passive IPO investors who do not make many subscriptions (the most passive investors were excluded in the prior regression analyses for technical reasons, as explained).

Table IV reports the results. The first two columns show the proportion of first-time IPO investors participating in further IPOs, conditional on having

Table IV
Natural Experiment of Survival to Next IPO

This table gives the proportion of investors participating in a subsequent IPO conditional on having experienced either a negative or a positive initial return in their first IPO. The sample includes six pairs of IPOs. To qualify as a pair in this test, the two IPOs must have simultaneously running subscription periods, and it must be the case that one IPO has a positive initial return, while the other one has a negative initial return. The sample includes all such pairs. The “cumulative proportion of participating investors” is given by the number of investors who have made at least one further subscription since the initial IPO divided by all investors participating in the initial IPO, and it is calculated for ten subsequent offerings. “First IPO” includes investors who make their first IPO investment in either one of the IPOs in a matched pair. “Stock market entry via IPO” includes only investors who have made neither IPO nor other common stock investments before. The *z*-values are for a test of difference in proportions.

Number of Offering after the First IPO with Positive or Negative Initial Return	Cumulative Proportion of Participating Investors							
	First IPO				Stock Market Entry Via IPO			
	Pos. Initial Return	Neg. Initial Return	Difference Between Proportions	<i>z</i> - Value	Pos. Initial Return	Neg. Initial Return	Difference Between Proportions	<i>z</i> - Value
1 st	0.10	0.04	0.06	7.7	0.11	0.07	0.05	2.1
2 nd	0.20	0.07	0.13	12.6	0.24	0.09	0.14	5.0
3 rd	0.33	0.13	0.20	15.3	0.38	0.16	0.22	6.0
4 th	0.43	0.17	0.26	17.5	0.50	0.19	0.30	7.3
5 th	0.45	0.18	0.28	17.7	0.53	0.21	0.32	7.1
6 th	0.50	0.19	0.30	18.7	0.55	0.22	0.33	7.2
7 th	0.56	0.21	0.35	20.9	0.58	0.24	0.34	7.1
8 th	0.60	0.35	0.25	12.9	0.63	0.34	0.29	5.8
9 th	0.64	0.37	0.28	13.0	0.67	0.35	0.32	5.7
10 th	0.65	0.39	0.26	11.0	0.67	0.38	0.30	4.7
Number of investors at the beginning of 1st IPO	2,304	2,182			412	252		

experienced either a negative or a positive initial return in their first IPO. Of those investors who start with a positive initial-return IPO, 10.2% subscribe to the IPO that comes next to the market. Only 4.3% of the investors who experienced a negative initial return in their first IPO do so. A test for a difference in proportions shows high statistical significance with a z -value of 7.7.

Looking at further offerings and the cumulative proportions of investors participating, the differences continue to be great and statistically significant up to the tenth offering after the first IPO an investor subscribed to. By the tenth offering, 65% of investors in the positive initial IPO group have subscribed to another IPO, compared to only 39% in the negative initial IPO group.

The corresponding figures for investors who enter the market through their first IPO subscription are shown on the right-hand side of Table IV. This group of investors is a subset of the first-time IPO investors considered earlier; they had not made any kind of stock market investment before subscribing to the IPO. The results are very close to those generated with first-time IPO investors.

We also run a duration regression similar to that reported in Table III including all the control variables. We pool the six offerings and include a dummy to indicate the better IPO of each pair. The dummy enters with a coefficient of 0.78 and is highly significant (t -value 16.6). This implies that the hazard rate of participating after a positive initial IPO is more than twice as high ($e^{0.78} = 2.18$) than the hazard rate after a negative IPO.

III. Alternative Explanations

The results of our three analyses confirm that investors are significantly more likely to participate in future IPOs if they have personally experienced better returns. Other explanations besides reinforcement learning may also be consistent with our main results. We now consider these explanations.

A. Private Information

If uninformed investors participate only occasionally and the informed are more active (as in Benveniste and Spindt (1989)), one should observe a positive relation between investors' past returns and future subscriptions, similar to the reinforcement hypothesis.

The subjects of our study are individual investors, who quite likely are uninformed. However, we still allow for the possibility that there are differences among investors. To account for the effects of differentially informed investors, we divide investors into groups according to prior general stock market investment performance by the portfolio change measure developed in Grinblatt and Titman (1993). The portfolio change measure, r_p , is calculated as $r_p = \sum_{i=1}^N r_{i,t}(w_{i,t} - w_{i,t-L})$, where N is the number of stocks held, $r_{i,t}$ is the return from stock i during period t , $w_{i,t}$ is the weight of stock i in the portfolio during period t , and L is the rebalancing period. With this measure one calculates the performance accruing to a stock as a product of the stock's return and

Table V
Effect of Past Return for Different Investor Types

This table reports regressions similar to that in Table II column 4 except that the past return variable is interacted with two indicator variables. Investors are first divided into quintiles based on performance (measured as own benchmark adjusted returns, see text) and total portfolio return, both calculated for the first half of the sample. Investors with no trading activity during the first half are left out of the analysis. For every investor, indicator variables for performance and portfolio return are defined in the following way. If an investor belongs to the highest quintile, the indicator gets a value of one. If she belongs to the lowest quintile, the indicator takes a value of -1 . All other investors are given the value of zero. The t -values (in parentheses below coefficients) are based on robust standard errors.

Specification	Logit (0 = No Future Subscriptions, 1 = At Least One Future Subscription)		
	1	2	3
Past mean initial return	4.48 (19.79)	4.37 (19.23)	4.38 (19.25)
× Performance indicator	-0.22 (-0.64)		-0.18 (-0.52)
× Portfolio return indicator		-0.22 (-0.60)	-0.22 (-0.59)
Performance indicator	0.04 (0.55)		0.02 (0.35)
Portfolio return indicator		0.24 (3.43)	0.24 (3.41)
Number of past IPO subscriptions	0.45 (30.60)	0.46 (30.65)	0.46 (30.65)
Portfolio value dummies	Yes	Yes	Yes
Trading activity dummies	Yes	Yes	Yes
Pseudo- R^2	0.21	0.21	0.21
Number of observations	19,039	19,039	19,039

the change in the weight of that stock. Overall performance is then given by aggregating over all the stocks held by the investor. We calculate monthly returns using monthly rebalancing periods. The measure directly gives the value added by trading, and there is thus no need to specify an asset pricing model for risk adjustment. This is an improved and more direct measure of informativeness than proxies such as portfolio size used in the prior literature.

Table V presents results based on logit regressions similar to those in Table II, except that we include additional indicator variables, as well as their interactions with past returns.¹³ The performance indicator takes a value of one for investors in the highest overall performance quintile, and a value of minus one for those in the lowest quintile. All other investors get a value of zero. This variable, as well as its interaction with past returns, is insignificant. Hence,

¹³ In some situations a logit model can produce incorrect coefficient estimates or t -statistics for interaction variables. As an alternative to logit, we estimated a linear probability model with OLS and the results are the same.

subscription activity is heavily, and similarly, influenced by past IPO returns for both the best- and poorest performing investors. Individual investors in the poorest performance group are prime candidates to describe as uninformed and unsophisticated investors, on average. If any group of investors learns by personal reinforcement, we would expect it to be this group. The fact that we find similar effects also in the best performance group suggests an even greater role for reinforcement learning in investment decisions.

B. Portfolio Rebalancing

Changes in stock prices directly cause changes in allocation weights, which might prompt the need to rebalance for an investor taking an active view on sectors or styles. Chen, Ho, and Wu (2005) find that adding IPOs to a set of benchmark portfolios improves the investment opportunity set, which suggests that IPOs could be seen as a distinct asset class. Thus, bad performance leading to a lower amount allocated to IPO stocks should increase the probability of subscribing to the next IPO for an investor who has an asset allocation motive. At a minimum, rebalancing should not lead to a positive relation between personally experienced returns and future subscriptions. To the extent that rebalancing motives influence investors in the sample, the effect of reinforcement learning is understated.

C. Wealth Effects

High returns from IPOs increase an investor's wealth, which can lead to an increased appetite for risk. To account for such an effect, we use raw portfolio returns as a proxy for changes in wealth, and divide investors into quintiles on this basis. Column 2 of Table V shows that while the portfolio return variable is a positive and significant determinant of future subscription activity, the impact of personal IPO returns remains very strong. Furthermore, the interaction between IPO returns and the portfolio return indicator is close to zero and insignificant. This shows that the relation between personally experienced IPO returns and subscription activity is similar for investors with different changes in wealth.

D. Preferential Treatment

Investment banks may reward regular investors by favoring them in allocations of hot IPOs (e.g., Cornelli and Goldreich (2001), Loughran and Ritter (2002)). Learning about this kind of preferential treatment could also cause an investor to increase her subscription activity in response to high past returns. Preferential treatment is not allowed in Finland, however. The regulator requires equitable treatment of all retail investors in share issues. In practice, this means that investment banks apply a mechanical rationing rule to determine all retail allocations.

E. Liquidity Constraints

Stocks with good prior returns are known to have higher trading volume (Lakonishok and Smidt (1986), Statman, Thorley, and Vorkink (2006)). Investors who hold poorly performing IPO stocks could therefore be relatively more liquidity constrained. In addition to actual market liquidity, there could also be a personal liquidity effect due to reluctance to realize losses (Odean (1998); see also Kaustia (2004) for the impact on IPO stock trading).

We investigate liquidity effects by separating the observations into two categories: instances in which investors make some stock market purchases between their IPO subscriptions, and instances in which they do not. Making a purchase indicates that a liquidity constraint is not binding. We run regressions similar to those reported in Table III (logit of participating in the next offering, as well as time until the next participation) on the two samples. The return on past IPO subscriptions is a significant determinant of future subscription activity in both groups, and the relation is actually stronger for those who have made other purchases. This indicates that the results are not due to liquidity-constrained individuals.

F. Hot Issue Market from 1999 to 2000

In unreported analysis we show that the two-period results (Table II) continue to hold when the sample is divided into three periods, and the middle period (hot issue market) is excluded. Using this division, the mean initial return is 11.2% in the first period, 63.6% in the excluded middle period, and 0.0% in the third period. Furthermore, the analysis in Table III contains IPO fixed effects that pick up effects related to the IPO cycle.

We conclude that it is unlikely that private information, portfolio rebalancing, wealth effects, preferential treatment by investment banks, liquidity constraints, or the hot issue market from 1999 to 2000 play a major role in explaining the main findings we document.

IV. Conclusion

We investigate the link between individual investors' tendency to subscribe to initial public offerings and the returns on past IPOs. The results have implications for the IPO and asset pricing literature, particularly concerning investor sentiment. They are also related to a more general question about the role of reinforcement learning in economic decisions.

Using data on 183,000 retail investors in 57 Finnish IPOs occurring from 1995 to 2000, we find that an increase in the returns that an investor earns on past IPO investments has a positive impact on this investor's propensity to participate in future IPOs. Controlling for offering fixed effects, which encompass important offer characteristics but also the impact of market returns and previous IPO returns, we find personally experienced returns to be an important determinant of future IPO subscriptions. This finding is consistent with

reinforcement learning theory. That is, the individual repeats behavior that has produced good outcomes in the past and avoids behavior that has produced poor outcomes. Importantly, the individual overweights personal experience.

Although we see reinforcement learning as a natural explanation for our results, it is also possible to specify other forms of learning and information use that could produce a similar outcome. Suppose that there is a rational uninformed investor who is determined to learn about an asset class that she is not familiar with. In the spirit of bounded rationality, she would first work through the problem of deciding how to obtain information. Although information is publicly available, there might be costs to locate and process it. If the costs are sufficiently high, the investor might find it most efficient to obtain the information by first-hand experience. After subscribing to a series of offerings she would learn the return distribution and could then decide whether to continue investing. Yet constructing a reliable estimate of the return distribution takes many trials, so the dependence between past returns and future subscriptions would in this case be manifested over only a fairly long period. Furthermore, it seems that information acquisition costs would need to be quite high—perhaps implausibly high—for personal experimentation to be optimal.

Learning by reinforcement has several implications for the study of financial markets. Our results provide a way to interpret sentiment through investor learning from personal experience. Recent IPO models (Derrien (2005), Ljungqvist et al. (2006)), suggest that issuers and investment banks' regular customers benefit from the involvement of sentiment investors. Cook et al. (2006) suggest that investment bankers use marketing efforts to encourage retail investors' participation in IPOs. A natural experiment in which we consider investors subscribing to their first IPOs shows that initial personal experience has both immediate and longer term consequences for investor willingness to use the same financial product in the future. This fits the view of investors as consumers. Our results concerning the role of personal experience in investment decisions also help us understand how popular investment styles develop, how investors form their tastes, and how cohort effects in investment behavior arise.

REFERENCES

- Ariely, Dan, Uri Gneezy, George Loewenstein, and Nina Mazar, 2008, Large stakes and big mistakes, *Review of Economic Studies*, forthcoming.
- Arthur, Brian W., 1991, Designing economic agents that act like human agents: A behavioral approach to bounded rationality, *American Economic Review* 81, 353–359.
- Asch, Solomon E., 1946, Forming impressions of personality, *Journal of Abnormal and Social Psychology* 41, 258–290.
- Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *Journal of Finance* 61, 1645–1680.
- Baker, Malcolm, and Jeffrey Wurgler, 2007, Investor sentiment in the stock market, *Journal of Economic Perspectives* 21, 129–151.
- Bakshi, Gurdip S., and Zhiwu Chen, 1994, Baby boom, population aging, and capital markets, *Journal of Business* 67, 165–202.

- Barberis, Nicholas, Ming Huang, and Tano Santos, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics* 141, 1–53.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–343.
- Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey, 2007, Managerial overconfidence and corporate policies, NBER Working Paper 13711.
- Benveniste, Lawrence M., and Paul A. Spindt, 1989, How investment bankers determine the offer price and allocation of new issues, *Journal of Financial Economics* 24, 343–361.
- Bereby-Meyer, Yoella, and Alvin E. Roth, 2006, The speed of learning in noisy games: Partial reinforcement and the sustainability of cooperation, *American Economic Review* 96, 1029–1042.
- Brav, Alon, and J.B. Heaton, 2002, Competing theories of financial anomalies, *Review of Financial Studies* 15, 575–606.
- Brown, Gregory W., and Michael T. Cliff, 2004, Investor sentiment and the near-term stock market, *Journal of Empirical Finance* 11, 1–27.
- Brown, Gregory W., and Michael T. Cliff, 2005, Investor sentiment and asset valuation, *Journal of Business* 78, 405–440.
- Camerer, Colin F., and Teck-Hua Ho, 1999, Experience-weighted attraction learning in normal form games, *Econometrica* 67, 827–874.
- Campbell, John Y., 2006, Household finance, *Journal of Finance* 61, 1553–1604.
- Campbell, John Y., and John Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark, 2000, Asset pricing with distorted beliefs: Are equity returns too good to be true? *American Economic Review* 90, 787–805.
- Charness, Gary, and Dan Levin, 2005, When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect, *American Economic Review* 95, 1300–1309.
- Chen, Hsuan-Chi, Keng-Yu Ho, and Cheng-Huan Wu, 2005, Initial public offerings: An asset allocation perspective, Working paper, SSRN-594321.
- Cook, Douglas O., Robert Kieschnick, and Robert A. Van Ness, 2006, On the marketing of IPOs, *Journal of Financial Economics* 82, 35–61.
- Cornelli, Francesca, and David Goldreich, 2001, Bookbuilding and strategic allocation, *Journal of Finance* 56, 2337–2369.
- Cornelli, Francesca, David Goldreich, and Alexander Ljungqvist, 2006, Investor sentiment and pre-IPO markets, *Journal of Finance* 61, 1187–1216.
- Cross, John G., 1973, A stochastic learning model of economic behavior, *Quarterly Journal of Economics* 87, 239–266.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839–1886.
- De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldman, 1990, Noise trader risk in financial markets, *Journal of Political Economy* 98, 703–738.
- Derrien, François, 2005, IPO pricing in “hot” market conditions: Who leaves money on the table? *Journal of Finance* 60, 487–521.
- Dhar, Ravi, and Ning Zhu, 2006, Up close and personal? An individual level analysis of the disposition effect, *Management Science* 52, 726–740.
- Einhorn, Hillel J., 1980, Learning from experience and suboptimal rules in decision making, in T.S. Wallsten, ed.: *Cognitive Processes in Choice and Decision Behavior* (Lawrence Erlbaum Associates, Hillsdale, NJ).
- Ellison, Glenn, and Drew Fudenberg, 1993, Rules of thumb for social learning, *Journal of Political Economy* 101, 612–643.
- Erev, Ido, and Alvin E. Roth, 1998, Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria, *American Economic Review* 88, 848–881.
- Fama, Eugene F., and Kenneth R. French, 2007, Disagreement, tastes, and asset prices, *Journal of Financial Economics* 83, 667–689.

- Fazio, Russell H., Mark P. Zanna, and Joel Cooper, 1978, Direct experience and attitude behavior consistency: An information processing analysis, *Personality and Social Psychology Bulletin* 4, 48–51.
- Feng, Lei, and Mark Seasholes, 2005, Do investor sophistication and trading experience eliminate behavioral biases in financial markets? *Review of Finance* 9, 305–351.
- Fisher, Kenneth L., and Meir Statman, 2000, Investor sentiment and stock returns, *Financial Analysts Journal* 56, 16–23.
- Fisher, Kenneth L., and Meir Statman, 2003, Consumer confidence and stock returns, *Journal of Portfolio Management* 30, 115–127.
- Gervais, Simon, and Terrance Odean, 2001, Learning to be overconfident, *Review of Financial Studies* 14, 1–27.
- Gilboa, Itzhak, and David Schmeidler, 1995, Case-based decision theory, *Quarterly Journal of Economics* 110, 605–639.
- Greenwood, Robin, and Stefan Nagel, 2008, Inexperienced investors and bubbles, *Journal of Financial Economics*, forthcoming.
- Grinblatt, Mark, and Matti Keloharju, 2000, The investment behavior and performance of various investor types: A study of Finland's unique data set, *Journal of Financial Economics* 55, 43–67.
- Grinblatt, Mark, and Sheridan Titman, 1993, Performance measurement without benchmarks: An examination of mutual funds returns, *Journal of Business* 66, 47–68.
- Grossman, Sanford J., Richard E. Kihlstrom, and Leonard J. Mirman, 1977, A Bayesian approach to the production of information and learning by doing, *Review of Economic Studies* 144, 533–547.
- Hirshleifer, David, 2001, Investor psychology and asset pricing, *Journal of Finance* 56, 1533–1597.
- Huettel, Scott A., P.B. Mack, and Gregory McCarthy, 2002, Perceiving patterns in random series: Dynamic processing of sequence in prefrontal cortex, *Nature Neuroscience* 5, 485–490.
- Ibbotson, Roger G., and Jeffrey F. Jaffe, 1975, Hot issue markets, *Journal of Finance* 30, 1027–1042.
- Kahneman, Daniel, 1994, New challenges to the rationality assumption, *Journal of Institutional and Theoretical Economics* 150, 18–36.
- Kahneman, Daniel, 2000, Experienced utility and objective happiness: A moment-based approach, in Daniel Kahneman and Amos Tversky, eds.: *Choices, Values, and Frames* (Cambridge University Press and the Russell Sage Foundation, New York).
- Kahneman, Daniel, Peter P. Wakker, and Rakesh Sarin, 1997, Back to Bentham? Explorations of experienced utility, *Quarterly Journal of Economics* 112, 375–405.
- Kaustia, Markku, 2004, Market-wide impact of the disposition effect: Evidence from IPO trading volume, *Journal of Financial Markets* 7, 207–235.
- Knutson, Brian, and Richard Peterson, 2004, Neurally reconstructing expected utility, *Games and Economic Behavior* 52, 305–315.
- Korniotis, George, and Alok Kumar, 2006, Does investment skill decline due to cognitive aging or improve with experience? Working paper, University of Notre Dame.
- Kumar, Alok, and Charles M.C. Lee, 2006, Retail investor sentiment and return comovements, *Journal of Finance* 61, 2451–2486.
- Lakonishok, Josef, and Seymour Smidt, 1986, Volume for winners and losers: Taxation and other motives for stock trading, *Journal of Finance* 41, 951–974.
- Lee, Charles M.C., Andrei Shleifer, and Richard H. Thaler, 1991, Investor sentiment and the closed-end puzzle, *Journal of Finance* 46, 75–109.
- Lemmon, Michael, and Evgenia Portniaguina, 2006, Consumer confidence and asset prices: Some empirical evidence, *Review of Financial Studies* 19, 1499–1529.
- Lettau, Martin, and Harald Uhlig, 1999, Rules of thumb versus dynamic programming, *American Economic Review* 89, 148–174.
- Linnainmaa, Juhani, 2006, Learning from experience, Working paper, University of Chicago.
- Ljungqvist, Alexander, Vikram Nanda, and Rajdeep Singh, 2006, Hot markets, investor sentiment, and IPO pricing, *Journal of Business* 79, 1667–1702.
- Loughran, Tim, and Jay R. Ritter, 2002, Why don't issuers get upset about leaving money on the table in IPOs? *Review of Financial Studies* 15, 413–443.
- Loughran, Tim, and Jay R. Ritter, 2004, Why has IPO underpricing changed over time? *Financial Management* 33, 5–37.

- Lowry, Michelle B., 2003, Why does IPO volume fluctuate so much? *Journal of Financial Economics* 67, 3–40.
- Lowry, Michelle B., Micah S. Officer, and G. William Schwert, 2006, The variability of IPO initial returns, Working paper No. FR 06-06, Simon School.
- Lowry, Michelle B., and G. William Schwert, 2002, IPO market cycles: Bubbles or sequential learning? *Journal of Finance* 57, 1171–1198.
- Malmendier, Ulrike, and Stefan Nagel, 2008, Depression babies: Do macroeconomic experiences affect risk-taking? Working paper, University of California, Berkeley, and Stanford University.
- Merlo, Antonio, and Andrew Schotter, 1999, A surprise-quiz view of learning in economic experiments, *Games and Economic Behavior* 28, 25–54.
- Nelson, Philip, 1970, Information and consumer behavior, *Journal of Political Economy* 78, 311–329.
- Odean, Terrance, 1998, Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775–1798.
- Peng, Lin, and Wei Xiong, 2006, Investor attention, overconfidence and category learning, *Journal of Financial Economics* 80, 563–602.
- Pouget, Sebastien, 2007, Adaptive traders and the design of financial markets, *Journal of Finance* 62, 2835–2863.
- Qiu, Lily, and Ivo Welch, 2006, Investor sentiment measures, Working paper, Brown University.
- Rajan, Raghuram G., and Henri Servaes, 1997, Analyst following of initial public offerings, *Journal of Finance* 52, 507–529.
- Roth, Alvin E., and Ido Erev, 1995, Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term, *Games and Economic Behavior* 8, 164–212.
- Schotter, Andrew, and Antonio Merlo, 2003, Learning by not doing: An experimental investigation of observational learning, *Games and Economic Behavior* 42, 116–136.
- Seru, Amit, Tyler Shumway, and Noah Stoffman, 2006, Learning by trading, Working paper, University of Michigan.
- Shleifer, Andrei, 2000, *Inefficient Markets: An Introduction to Behavioral Finance* (Clarendon Lectures in Economics, Oxford University Press, New York).
- Simonsohn, Uri, Niklas Karlsson, George Loewenstein, and Dan Ariely, 2008, The tree of experience in the forest of information: Overweighing experienced relative to observed information, *Games and Economic Behavior* 62, 263–286.
- Statman, Meir, Steven Thorley, and Keith Vorkink, 2006, Investor overconfidence and trading volume, *Review of Financial Studies* 19, 1531–1565.
- Thorndike, Edward L., 1898, Animal intelligence: An experimental study of the associated processes in animals, *Psychological Monographs* 2, 1–109 (excerpted in *American Psychologist* 1998, 53, 1125–1127).
- Tversky, Amos, and Daniel Kahneman, 1973, Availability: A heuristic for judging frequency and probability, *Cognitive Psychology* 5, 207–232.