RFS Advance Access published March 17, 2011

Do Investors Learn from Experience? Evidence from Frequent IPO Investors

Yao-Min Chiang
Department of Finance, National Chengchi University

David Hirshleifer
The Paul Merage School of Business, University of California, Irvine

Yiming Qian
Department of Finance, University of Iowa

Ann E. Sherman
Department of Finance, DePaul University

We examine how experience affects the decisions of individual investors and institutions in IPO auctions to bid in subsequent auctions, and their bidding returns. We track bidding histories for all 31,476 individual investors and 1,232 institutional investors across all 84 IPO auctions during the period from 1995 to 2000 in Taiwan. For individual bidders, (1) high returns in previous IPO auctions increase the likelihood of participating in future auctions; (2) bidders' returns decrease as they participate in more auctions; (3) auction selection ability deteriorates with experience; and (4) those with greater experience bid more aggressively. These findings are consistent with naïve reinforcement learning wherein individuals become unduly optimistic after receiving good returns. In sharp contrast, there is little sign that institutional investors exhibit such behavior. (*JEL* G15, G24, G32)

In the fields of economics and finance, there has been growing interest in the effects of experience on decision-making. On the one hand, economic agents can learn to make better decisions as they acquire better information about the environment or the quality of their information signals (e.g., Arrow 1962; Grossman, Kihlstrom, and Mirman 1977; Mahani and Bernhardt 2007; and Linnainmaa 2010). On the other hand, there are psychological biases that can hamper the learning process, which can result in decision quality deteriorating systematically with increased experience.

We thank the editor (Alexander Ljungqvist), an anonymous referee, Matt Billett, Redouane Elkamhi, Eric Lie, Ashish Tiwari, Tong Yao, and Yu Yuan for helpful discussion and comments. Send correspondence to Ann E. Sherman, Department of Finance, DePaul University, 1 E. Jackson Blvd., Suite 6108, Chicago, IL 60604; USA; telephone: (312) 362-5499. E-mail: ann.sherman@depaul.edu.

doi:10.1093/rfs/hhq151

For example, there is evidence that people learn through naïve reinforcement, wherein they expect the gains or losses that they have personally experienced to recur, even in situations where such an expectation is not logically justified. In other words, they expect success in an endeavor to indicate favorable prospects for future success, while failure foreshadows future failure. Although a rational Bayesian may indeed draw such inferences to some extent, under naïve reinforcement learning individuals overweight their personal experience relative to information obtained by communication with or observation of others (Cross 1973; Arthur 1991; Roth and Erev 1995; and Camerer and Ho 1999).[1]

A good understanding of how investors learn is essential for modeling financial and economic decisions. In traditional economic models, economic agents are often assumed to solve complex decision problems flawlessly. A common justification for this approach is that in repeated settings individuals eventually learn to play correct strategies. However, if learning does not improve decision-making, this conclusion is not justified. It is therefore important to examine evidence from a variety of settings to determine in what contexts individuals are able to learn their way out of bias, and in what contexts learning exacerbates bias.

In this article, we explore whether investors learn to improve their strategies through experience (rational learning), or whether experience causes them to invest less effectively (naïve reinforcement learning). To do so, we use a dataset with complete bid information for all the 84 IPO auctions during the period from 1995 to 2000 in Taiwan. We are able to track the bidding histories of 31,476 individual investors and 1,232 institutional investors during that period.

The fact that this dataset includes the first IPO auction in Taiwan and *all* the IPO auctions during the period is important for studying the effects of experience, because it provides the complete bidding history of all bidders prior to each IPO auction. Furthermore, these data include large numbers of both individual and institutional investors. As a result, we are able to explore whether past experience influences investors with different degrees of sophistication in different ways.

We hypothesize that under both types of learning, investors are more likely to bid in future IPO auctions if they received high returns from past auctions. However, rational learning will lead to improved bidding strategies and hence better return performance, while naïve reinforcement learning will lead to deterioration in bidding strategies and therefore worse return performance (see Section 1 for detailed hypothesis development).

[1] One possible explanation for over-extrapolation of past gains or losses into the future is the representativeness heuristic, and the resulting “law of small numbers” (Tversky and Kahneman 1971). Under this effect, a small sample observed from past experience is viewed as representative of the underlying probability distribution, and therefore is overweighted.

We first examine the relationship between investors' past returns from previous IPO auctions and their inclination to participate in future IPO auctions. We find that individual investors tend to bid in the future if they received high past returns and tend to stop bidding if they received poor past returns. In sharp contrast, we find that the decisions of institutional investors to bid are much less, if at all, influenced by their past returns.

To differentiate between the rational and naïve reinforcement learning hypotheses, we examine investors' return performance as they gain more experience. Consistent with the naïve reinforcement learning hypothesis, we find that individual investors' returns steadily decrease as they gain more experience. For example, the mean return decreases by −2.88% from their first deals to second deals, and mean dollar profit decreases by $NT 213,860 (a 39% decrease).[2] In contrast, institutional investors' returns do not decrease, so there is no indication that institutional investors are subject to naïve reinforcement learning.

The effect of experience on returns can reflect changes in two kinds of bidding skills needed for investors to achieve good returns from participating in IPO auctions (Sherman 2005; Chiang, Qian, and Sherman 2010): the ability to judge firm quality (i.e., auction selection), and the ability to shave bids sufficiently (to address the winner's curse and to compensate for information costs). We therefore investigate how each of these skills changes as investors learn from experience.

For individual investors, we find that auction selection ability deteriorates with experience, and that experience causes greater aggressiveness in bid prices. So, the decline in individual investor returns with experience reflects a failure of learning for both kinds of skills. For institutional investors, however, neither of these bidding skills deteriorates with experience.

Finally, we investigate whether in the long run individual investors ever start to learn more effectively from experience. There is some indication that given long enough (at least 24 auctions in our tests), individual investors start to benefit from experience. This is quite a lot of experience, however, and only a very small proportion (0.2%) of individual investors in our sample achieve that level of experience.

This study therefore makes two contributions to the learning literature. First, we provide evidence that individual investors in IPO auctions are subject to *naïve* reinforcement learning. The decline in their performance with experience is not consistent with rational Bayesian learning. Second, we document that a potentially more sophisticated set of players, institutional investors, do not seem to be subject to this learning bias.

Our article is far from the first to address the issue of how decision-makers learn from experience. Several studies examine learning in laboratory experiments (Erev and Roth 1998; Camerer and Ho 1999; and Charness and Levin

2 All NT$ values in the article are in constant year 2000 NT$. The end of 2000 exchange rate was US$1 = NT$32.99.

2005). Using survey data, Malmendier and Nagel (2009) document that expectations about future inflation rates are largely influenced by individual experiences of inflation, consistent with reinforcement learning.

Other papers have studied the effects of learning using actual trading and asset allocation data. Consistent with the rational-learning hypothesis, Feng and Seasholes (2005) and Dhar and Zhu (2006) find that investors' trading experience reduces the behavioral bias of the disposition effect. Nicolosi, Peng, and Zhu (2009) and Seru, Shumway, and Stoffman (2009) present evidence that individual investors learn from trading experience and improve performance. On the other hand, consistent with reinforcement learning, Barber, Odean, and Strahilevitz (2010) document that investors tend to repurchase stocks they previously sold for a gain and shun stocks they previously sold for a loss. Choi, Laibson, Madrian, and Metrick (2009) find that an investor's 401(K) contribution rate increases more if her account has recently experienced a high return or low return variance, and that such behavior is not welfare-improving.[3] These papers all focus on individual investors. A distinctive aspect of our article is that we examine whether individual versus institutional investors differ in how they react to past experience and in their abilities to improve their performance.

Kaustia and Knupfer (2008) provide evidence suggesting that individual investors are more likely to participate in IPOs after good returns from past IPOs. This evidence is valuable in documenting that belief updating about IPO investment does occur in response to experience. However, their tests are not able to distinguish the key alternative hypotheses of our paper—whether investors rationally learn to improve their profits, or whether investors update naïvely, and thereby "learn to fail". Although the authors interpret their results as consistent with naïve reinforcement learning, their evidence is also consistent with the hypothesis that investors learn rationally. If investors learn about their investing abilities (or the accuracy of their information signals) through experience, then those who have high abilities will continue more often than those with low abilities (Mahani and Bernhardt 2007; Seru, Shumway, and Stoffman 2009), consistent with Kaustia and Knupfer's finding.

There are two key differences between our article and that of Kaustia and Knupfer (2008). First, by examining performance over time, we assess whether investors are engaged in rational learning, or become less competent owing to naïve reinforcement learning. The IPO auctions in our sample are discriminatory price auctions (i.e., bidders pay what they bid when they win). This implies that winning bidders will receive different initial returns even from

[3] There is evidence consistent with biased learning for other market participants: Billett and Qian (2008) find that CEOs that frequently engage in acquiring other firms see more negative wealth effects in their later deals while their insider trading prior to those deals becomes more bullish. Hilary and Menzly (2006) find that analysts who have predicted earnings more accurately in the recent past tend to be less accurate and further from the consensus in subsequent earnings predictions, although there are mixed results regarding the effect of long-term experience on analyst forecast accuracy (see Mikhail, Walther, and Williams 1997; and Jacob, Lys, and Neale 1999).

the same auction. These auctions therefore provide a great deal of relevant information for exploring the effects of personal experience on returns. In contrast, Kaustia and Knupfer (2008) use data from offerings in which retail investors play no price-setting role, since they choose only whether or not to order shares. All investors pay the same offer price and receive the same initial returns from the same IPO.

Second, Kaustia and Knupfer (2008) focus on individual investors, whereas we test whether individual and institutional investors differ in their abilities to learn valid lessons from past experience. An added advantage of our dataset is that, in contrast with IPOs that use bookbuilding or fixed price public offer methods, underwriters in IPO auctions have no discretion over either pricing or allocation. Thus, our study is not complicated by the interfering effects of underwriter discretion.

Our findings also have implications for IPO design. Chiang, Qian, and Sherman (2010) document that institutional investors are informed bidders with sophisticated bidding strategies, while individual investors are not. The authors thus raise the question of whether individual investors as a group have the sophistication to price highly risky securities such as IPO stocks. Our article differs in studying how different kinds of investors *learn from experience* in IPO auctions. There is little sign in our evidence that experience with the IPO auction method makes individuals better investors in IPO auctions. This raises the question of whether regulatory protections are needed for individual investors participating in IPO auctions.[4]

1. Hypotheses

The model of Sherman (2005) gives the optimal bidding strategy in an IPO environment in which investors decide whether and how to participate in a multiple unit sealed bid common value auction.[5] In the model, bidders decide whether or not to incur a fixed cost to acquire information about the offering. Upon receiving a signal, bidders decide how much to bid.

In equilibrium it is unprofitable to bid without acquiring information. Bidders follow a mixed information acquisition (entry) strategy, randomizing between obtaining information and then bidding, or simply avoiding the auction entirely. Bidders enter to the point where on average they just recover their

[4] Jagannathan, Jirnyi, and Sherman (2010) offer further evidence regarding the computational challenges involved in bidding in IPO auctions, and point out that unsophisticated bidders following suboptimal bidding strategies increase risk even for sophisticated bidders.

[5] Sherman (2005) models both discriminatory and uniform price auctions. In a discriminatory auction, all winning bidders pay the price that they bid. In a uniform price or “Dutch” auction, all bidders pay the same price. Although in this article we examine a discriminatory auction sample, the hypotheses we develop also apply to uniform price auctions such as those done in the U.S. The scenario we discuss differs from the Sherman model in allowing for a reservation price, so that a bidder with a sufficiently low signal value finds it unprofitable to bid. This allows us to consider how rational learning affects auction selection.

evaluation costs. Having observed their information signals, those that enter follow mixed strategies in choosing the levels of their bids. Bid levels reflect a balance between the costs and benefits of a higher bid (higher chance of winning, but lower expected return to winning), so that a bidder has the same overall expected return from any bid within the equilibrium range. However, if an investor receives a signal value so low that the preferred bid level is below the auction reservation price, the investor does not bid. In equilibrium, bidders shave their bids relative to their true valuations, in part to compensate for the winner's curse.

In this approach, bidders know how much they should rely on their private signals relative to public signals (i.e., they know their own signal precision), so there is no scope for a bidder to learn through personal experience. If we consider this setting repeated each period without learning, we do not expect to see systematic trends in bidder behavior or performance (such as expected returns and bidding strategies) in relation to bidder experience.

We consider a rational hypothesis that allows for bidder learning by modifying the environment described above to allow for uncertainty about how much a bidder should rely on private signals. Under rational Bayesian updating, investors learn over time how much to trust their private signals relative to public information and adjust the weight they give to their private signals when updating their expectations. Each investor begins with a prior about the firm's value based on public information. The degree to which she updates her prior upon receiving her private signals depends on the perceived accuracy of the signal. As an investor observes her bidding outcomes over time, she rationally updates her estimate of her signal accuracy, which affects how much weight she places upon her signal relative to public information.

In this framework, an investor rationally infers that her signal is on average more accurate if she received high returns from previous IPO auctions. Placing undue weight on inaccurate signals is unprofitable, because she is very likely to "win" low-value shares when she bids too high, and is unlikely to win high-value shares when she bids too low.

Sherman (2005) shows that, all else equal, the equilibrium entry probability is higher if the signal is more accurate. Intuitively, a more accurate signal, ceteris paribus, increases the expected benefit of entering. Higher returns to a bidder in past auctions cause a rational inference that her signals are accurate. Thus, investors are more likely to bid in future auctions if they receive higher returns from past auctions.

Consider next an alternative scenario in which investors are subject to naïve reinforcement learning. Such investors update their beliefs about their prospects in an investment domain (e.g., investments in IPO auctions) as an increasing function of their past success in this domain, regardless of whether or not such updating is logically justified. Thus, they are also more likely to bid again in future auctions if they receive higher returns from past auctions.

In summary, we have the following hypothesis.

Hypothesis 1. If investors learn from experience, either by rational Bayesian updating or naïve reinforcement learning, those that receive higher returns from past auctions are more likely to continue to place bids in future auctions.

However, these two types of learning offer very different predictions about the relation of the bidding experience to future return performance. Under rational Bayesian updating, those who bid again (i.e., experienced bidders) tend to be those who rationally learn that their signals are likely to be of high accuracy. In consequence, such bidders will rationally place greater weight on their signals. Experienced bidders thus tend to have more accurate posterior expectations than inexperienced bidders. Greater precision improves decision accuracy and expected returns (Sherman 2005), so expected returns should also increase with experience.

Furthermore, a bidder's expected profits depend on two factors: auction selection ability and bid aggressiveness (whether she shaves her bids adequately). To evaluate selection ability, a measure is needed of the quality of the investment opportunity in a given auction. An IPO auction's average initial return measures the average profits to auction investors, and is therefore an ex post proxy for the quality of the auction as an investment opportunity. Furthermore, Sherman (2005) shows that when investors are informed and bidding optimally, on average the IPO initial return is an increasing function of firm quality, because in equilibrium there is partial adjustment of the auction issue price to information.

We refine this proxy by measuring bidder i's auction selection ability as the value-weighted average return to *other bidders* in an auction in which bidder i bids. Since all winning bidders pay what they bid, other bidders' average return depends on the quality of the IPO (which is what selection ability is about) and not on how aggressively investor i bids in that auction.

Under rational updating, as bidders gain experience, they learn their information precision more accurately, which allows them to form more accurate posterior beliefs and therefore to select auctions more accurately. We therefore predict that greater experience is also associated with high subsequent returns to other bidders who participate in the same subsequent auctions as bidder i.

Hypothesis 2. If investors are learning by rational Bayesian updating, then as the experience of a bidder increases:

a) the bidder's expected return will increase; and
b) the expected returns of other bidders who participate in the same subsequent auctions as the given bidder will also increase.

Alternatively, if investors are subject to naïve reinforcement learning, they become unduly optimistic about firm quality and the IPO expected return following higher returns to past auctions. Such optimism has two consequences. First, investors are more likely to bid in all offerings, including low-value offerings. In other words, they become less selective regarding which auctions to participate in. This poorer selection ability implies *lower* expected returns to other bidders participating in the same subsequent auctions as the bidder. Second, when they bid, they tend to offer higher prices. In other words, their bids will be more aggressive.[6] Both consequences contribute to lower expected returns to the bidder.

Hypothesis 3. If investors are engaged in naïve reinforcement learning, then as the experience of a bidder increases:

a) the bidder's expected return will decrease;
b) the expected returns of other bidders who participate in the same subsequent auctions as the given bidder will also decrease; and
c) the bidder's bids will become more aggressive.

2. Data and Sample

2.1 IPO auctions in Taiwan

IPO auctions in Taiwan are multiple-unit, multiple-bid discriminatory-price sealed-bid auctions. A bidder can submit multiple bids (different combinations of price and quantity) and will pay what she bid if she wins.[7] Both individuals and institutions can participate in these IPO auctions. A bidder is allowed to win no more than 6% of the auction shares, which for an average auction is NT$58.0 million at the average winning price. Institutions bid in all but two of the auctions in our sample.

Bidders submit sealed bids during a pre-announced bid-tender period that lasts for four business days. At the time of submitting bids, investors need to pay a transaction fee of NT$500 for each bid and a bid deposit that is no less than 20% of the bid size (i.e., bid price times bid quantity). During the bid-tender period, underwriters are not allowed to open the sealed bids and are

[6] Intuitively, under rational Bayesian learning, experience tends to be associated with placing more weight on one's private signals, because bidders who rationally place less weight on their private signals are less likely to continue bidding. Placing more weight on one's private signals leads to a wider spread of expected values (higher for good signals, lower for bad signals) relative to the prior expected value before the signal was received. Bids for the lowest value shares are less likely to be placed at all, since the optimal bid is more likely to be below the reservation price, meaning that the average bid that is placed is likely to be higher. Thus, rational Bayesian learning will tend to lead to more aggressive bidding with experience. However, optimal bid-shaving depends on the level of the signal and is generally greater for higher signals, which can also affect the relation of optimal bidding aggressiveness to experience. It is not clear whether the net effect will lead the optimal bidding aggressiveness to increase or decrease with experience under rational learning.

[7] IPO auctions in Taiwan are a hybrid IPO method—half of the IPO shares are to be sold in an auction, and then the other half of the shares are to be sold in a subsequent fixed-price public offer. Chiang, Qian, and Sherman (2010) show that there are no strategic interaction between the two stages and they are essentially two independent sales from investors' point of view. Hence we focus only on auctions in this article.

explicitly forbidden from revealing bid information (Taiwan Securities Association Rules Governing Underwriting and Resale of Securities by Securities Firms, Article 14). By 9:00 AM on the morning after the auction closes, the bid log is delivered to the Taiwan Securities Association, where the bids are opened and allocations are determined. We call this date the auction date.

The clearing price is the maximum bid price that clears the supply. All bids that are strictly above the clearing price are filled in full, while bids at the clearing price are awarded by lottery. Bidding results including the clearing price, subscription ratio, and winning price and quantity for each winner are then announced. Information on losing bids is not made publicly available (but is available in our data).

Shares are seldom able to trade freely during their first official day on the aftermarket, due to limits on daily price changes. In Taiwan, a daily return limit of 7% in each direction is imposed on all publicly traded stocks, including IPO shares, during our sample period. IPO shares frequently hit this limit for the first few days in a row. The first day when the stock price falls within the limit is known as the first non-hit day. We compute an investor's initial return (or IPO underpricing) based on the closing price of the first non-hit day. This initial return is comparable to an IPO's first-day return in the United States, where IPOs do not face daily price limits. In our sample, the "honeymoon period" (the time from the first trading day to the first non-hit day) has a mean (median) of 5.4 (3) trading days and ranges between 1 and 28 trading days. For more institutional details on IPO auctions in Taiwan, see Chiang, Qian, and Sherman (2010).

2.2 Sample

The sample includes all 84 IPO auctions in Taiwan during the 1995 to 2000 period.[8] We obtain detailed bidding information on each auction from the Taiwan Securities Association, including bidder IDs and the bid price and quantity of every bid by each bidder. The format of the bidder ID tells us whether the bidder is an institutional or individual investor. A bidder uses the same ID across auctions, and therefore the dataset allows us to track the bidding history of each bidder. The dataset also includes information on the auction size, the reservation price, the clearing price, and the auction proceeds.

Background information about the IPO firms such as assets, venture capital ownership, and P/E ratio is collected from the firms' prospectuses, which are available from the Taiwan Securities & Futures Information Center database. Stock returns for individual stocks and the market are from the *Taiwan Economic Journal* (TEJ).

Table 1 displays summary statistics of the sample, with variable definitions given in the Appendix. Panel A reports firm characteristics. The average IPO

8 There are seven IPO auctions that came after 2000: three in 2001, two in 2002, one in 2003, and one in 2008. Bid data for these auctions are not available due to the new privacy policy of the Taiwan Securities Association.

Table 1
Summary statistics

Panel A. IPO auction firm characteristics, *N*=84

	Mean	Median	Std. Dev.	25th percentile	75th percentile
% of IPOs on TSE	55.95				
% of IPOs in high tech industry	53.57				
Assets (in millions of NT$)	10,382.46	1,936.35	49,900.46	1,189.78	3,414.02
Auction proceeds (in millions of NT$))	880.58	417.72	2,522.74	149.87	886.28
VC ownership (%)	13.46	6.82	16.24	0.00	23.18
P/E	18.34	16.24	13.26	11.82	20.11
% of shares auctioned	6.51	6.65	2.88	3.75	9.81

Panel B. IPO initial returns, *N*=84

	All	1995	1996	1997	1998	1999	2000
IPO initial return (%)	7.25	6.82	35.46	0.72	4.21	11.98	-11.51
IPO initial return – individuals (%)	7.19	6.82	35.89	0.36	4.01	12.18	-11.54
IPO initial return – institutions (%)	8.70		37.03	3.34	5.31	15.15	-11.31
# of individual bidders per auction	676.77	235	1,595.27	763.05	373.52	452.87	771.44
# of institutional bidders per auction	31.99	2	49.82	24.37	27.45	33.73	41.33
# of IPO auctions	84	1	11	19	29	15	9

Panel C. Bidder activities and returns (bidder-auction observations)

	Individuals	Institutions	Difference
N (# of bidder-auction obs.)	56,849	2,687	
Average NT$ bid size (in thousands)	2,343.32	23,144.97	-20,801.7***
Average NT$ investments (winning size) (in thousands)	2,857.78	28,306.70	-25,448.9***
Mean initial return (%)	5.50%	11.54%	-6.04%***
Median initial return (%)	0	4.61%	-4.61%***
Standard deviation of initial return (%)	22.83%	26.08%	
Minimum initial return (%)	-50.00%	-33.77%	
Maximum initial return (%)	110.59%	110.59%	
% of positive initial returns	49.44%	64.09%	

Panel D. Bidder activities and returns (at the level of unique bidder)

	Individuals	Institutions	Difference
N (# of bidders)	31,476	1,232	
Average # of IPO auctions	1.81	2.18	-0.37***
Average # of winning IPO auctions	0.51	0.79	-0.28***
Mean NT$ bid size (in thousands)	2,325.05	27,732.80	-25,407.75***
Mean NT$ investments (winning size) (in thousands)	2,822.16	34,328.88	-31,506.72***
Mean value-weighted initial return (%)	3.71%	10.87%	-7.17%***
Median value-weighted initial return (%)	-3.37%	4.22%	-7.59%***
Mean dollar profits per auction (thousands)	218.23	1,309.20	-1,090.97*
Median dollar profits per auction (thousands)	-17.50	295.91	-313.41***

The sample includes 84 IPO-auctions in Taiwan during the 1995 to 2000 period. Panel A reports firm characteristics. NT$ refers to New Taiwan Dollars. All NT$ values are deflated to constant year 2000 NT$. The exchange rate at the end of year 2000 is US$1 = NT$32.99. Panel B shows by year the average initial returns of these IPO auctions. We compute the initial return for each auction as the closing price on the first non-hit day over the quantity-weighted average of the winning price. Panels C and D report bidder activities and returns. Panel C treats each bidder-auction combination as an observation. Panel D treats each unique bidder as an observation (the mean values are obtained by first averaging across auctions for each bidder, then averaging across bidders). ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

firm has assets of NT$10.4 billion and raises NT$880.6 million in the auction. Panel B reports the means of IPO initial returns and numbers of bidders for the whole sample period and by year. We compute the initial return for each auction as the closing price on the first non-hit day over the quantity-weighted average winning price minus one (i.e., the value-weighted average initial returns of all winning bidders). We also compute the average initial return in each auction separately for individuals and institutions. The mean IPO initial return over the sample period is 7.3%. On average, 676.8 individuals and 32.0 institutions bid in each IPO auction.

Not surprisingly, the mean return for individual investors is very close to that for the whole sample since most bidders are individuals. Institutional investors on average earn higher returns, with a mean return of 8.7% versus 7.2% for individuals. The time-series variations in returns are similar for individual and institutional investors. There is no obvious time trend for investor participation.

Panels C and D in Table 1 report bidder activities and returns. Unlike in Panels A and B, where each IPO auction is one observation, in Panel C (D) we treat each auction bidder (each unique bidder) as an observation. There are 31,467 (1,232) unique individual (institutional) bidders with 56,849 (2,687) bidder-auction observations. Of these, there are 16,037 (971) winning individual (institutional) bidder-auction combinations for which returns can be calculated. The average individual bid size (i.e., sum of bid price times bid quantity across a bidder's bids) is NT$2.3 million, and investment size (i.e., sum of bid price times winning quantity across a bidder's winning bids) is NT$2.8 million. In comparison, the average institutional bid size is NT$23.1 million and investment size is NT$28.3 million. Individual investors earn a mean (median) initial return of 5.5% (0%), whereas institutional investors earn a mean (median) initial return of 11.5% (4.6%). All the differences are significant at the 1% level.

Despite the positive mean (nonnegative median) returns, investments in these IPOs are very risky. The standard deviation of returns is 22.8% for individual investors and 26.1% for institutional investors. Individual investors' returns range from −50.0% to 110.6%, and institutional investors' returns from −33.8% to 110.6%. Individual investors receive positive returns only 49.4% of the time, and institutional investors 64.1% of the time.

Panel D of Table 1 examines bidder activities and returns at the unique bidder level (i.e., we first average across auctions for each bidder, and then average across bidders). The average individual investor participates (i.e., bids) in 1.8 auctions and wins shares in 0.5 auctions; the average institutional investor participates in 2.2 auctions and wins shares in 0.8 auctions. The differences in bid size and investment size between individuals and institutions are similar to those in Panel C. For bidding results, we calculate each bidder's value-weighted return and average dollar profits. The contrast in returns between individuals and institutions is even more striking than that in Panel C. The median value-weighted return for individuals is negative (−3.4%). In terms of the mean (median) dollar profits per auction, institutions earn NT$1.1 (NT$0.3) million higher profits than individuals, with the difference significant at the 10% (1%) level. The large difference between the mean and median profits suggests that dollar profits are highly skewed.

Most investors in our sample (74.8% of individuals and 66.5% of institutions) bid in only one auction. The proportion of investors bidding in two auctions is 11.7% for individuals and 11.9% for institutions. However, bidders with two or more auctions account for 60.1% of all bidder-auction observations.[9]

3. Past Returns and the Likelihood of Participating in Further Auctions

We test Hypothesis 1 by examining the relationship between a bidder's past returns and her likelihood of bidding again. Using a procedure similar to that of Kaustia and Knupfer (2008), we divide the sample into two subsamples with similar numbers of winning bidder-auction observations, and investigate whether returns received in the first half affect a bidder's likelihood of bidding in the second half. The first half of the sample includes 44 (out of 84) auctions and 8,442 (out of 17,008) winning bidder-auction observations. We compute a bidder's return in the first half as the investment-weighted average of her returns from all the auctions for which she wins shares during that period. We examine the probability of bidding in the second half for each quartile of returns in the first half, for individual and institutional investors, respectively.[10]

We find that individual bidders are more likely to bid in the second half if they receive higher returns in the first half. The probability of bidding again in the second half jumps from 29.8% in the lowest return quartile to 38.5% in the 2nd quartile, and increases to 41.9% in the 3rd quartile, and then levels off at 40.0% in the 4th quartile. For institutions, on the other hand, there is no such positive association between past returns and the probability of bidding in the second half, although institutions on average have a higher tendency to bid again than individuals. The probability of bidding in the second half is 46.3%, 37.8%, 43.9%, and 47.3% for the 1st, 2nd, 3rd, and 4th return quartiles, respectively (results not tabulated).

Different investors may have different propensities to bid (even apart from past returns). To control for this, we examine separately bidders with different numbers of auctions in the first half. Panel A of Table 2 shows the

9 As will be discussed later, due to the time overlaps of some auctions, bidders might have multiple first auctions. In other words, a bidder may have two auctions but both auctions may be defined as her first auction. Bidders with second or later auctions count for 50.4% of all bidder-auction observations.

10 Although we examine individual and institutional investors' probability of bidding separately, the return quartiles are determined for all bidders together.

Table 2
Past return and probability of subsequent bidding

Panel A. Probability of bidding in the second period by number of auctions and return quintile in the first period

		Individuals			Institutions		
# of auctions In the first half	Return Quartile	*N*	Avg. Past Return	Prob. of bidding	*N*	Avg. Past Return	Prob. of bidding
1	Lowest	660	-9.59%	18.64%	31	-6.49%	45.16%
	2	660	4.86%	26.52%	68	5.59%	22.06%
	3	757	18.19%	30.38%	30	17.11%	33.33%
	Highest	930	51.32%	33.66%	55	67.90%	40.00%
2	Lowest	172	-8.78%	28.49%	5	-12.09%	60.00%
	2	123	5.31%	43.09%	19	5.61%	47.37%
	3	132	17.66%	60.61%	8	14.91%	62.50%
	Highest	105	35.77%	60.95%	4	66.99%	75.00%
≥3	Lowest	368	-7.81%	50.54%	5	-1.57%	40.00%
	2	348	5.98%	59.48%	24	5.80%	75.00%
	3	296	16.89%	62.84%	19	17.24%	52.63%
	Highest	132	33.40%	68.18%	15	45.93%	66.67%

Panel B. Two-subsample analysis: the dependent variable equals one if the bidder bids again in the second period and zero otherwise.

	Individuals		Institutions	
	Estimates	*t*	Estimates	*t*
Past return	1.3068	10.89***	0.4256	1.04
Log(1+# of auctions in 1st period)	0.8425	20.01***	0.7996	4.16***
Intercept	-1.1760	-25.27***	-0.7381	-4.12***
Log likelihood	-2,849.64***		-183.59***	
Pseudo R^2	0.08		0.05	
N	4,683		283	

The sample is divided into two halves with similar numbers of winning bidder-auction observations. The first half of the sample includes 44 (out of 84) auctions and 8,442 (out of 17,008) winning bidder-auction observations. A bidder's past return in the first half is the investment-weighted average of her initial returns from all the auctions she wins during that period. Return quintiles are determined for all bidders together. In Panel B, the dependent variable is a dummy equal to one if a bidder bids again in the second period. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

bidding probability by return quartiles for bidders with one, two, and three or more auctions. We find that bidders with more auctions in the first half indeed are more likely to bid again in the second half, holding the return quartile fixed. Nonetheless, for each bidder category, past returns still have a large positive effect on individual investors' decisions to participate in future auctions. For individual investors with any number of auctions in the first half, their bidding probability in the second half steadily increases across return quartiles. In contrast, for institutional investors in each auction-number group, we do not observe such a relationship between past returns and bidding probabilities.

Panel B of Table 2 reports logit regressions that test whether past returns affect the likelihood of bidding again. The dependent variable is a dummy equal to one if the bidder bids again in the second half of the sample. The regressor of interest is a bidder's past return in the first half. We control for the number of auctions in which she participates in the first half. Specifically, we use the natural log of one plus the number of auctions. Consistent with the univariate tests, the coefficient on past return is significantly positive for individuals, but not significant for institutions. For individual investors, when the past return moves from the 25th percentile to the 75th percentile, holding the control variable at its mean, the probability of bidding increases from 32.3% to 39.1%.

In addition to the two-subsample analysis, we perform additional logit regression analyses that include all winning bidder-auction observations, except for those of the last two auctions since there are no future auctions after them in our sample (see Online Appendix, Table A1). In the first alternative test, the dependent variable is a dummy equal to one if the bidder ever bids again in the sample after winning the current auction and zero otherwise. In the second alternative test, the dependent variable equals one if the bidder bids in the next auction and zero otherwise. In both tests, the main variable of interest, *past return,* is a bidder's investment-weighted average return up through the current auction. Results are qualitatively similar to those in Table 2, Panel B: For individual investors, the coefficient on *past return* is significantly positive; for institutional investors, the coefficient is insignificant.

In summary, we find that individual investors' propensity to bid again in an IPO auction increases with their returns from previous auctions. This is consistent with the findings in Kaustia and Knupfer (2008) for individual investors in fixed-price IPO offerings; and with either rational learning or naïve reinforcement learning. We further find, in sharp contrast, that institutional investors' bidding decisions are not affected by their own past returns.

4. Bidder Experience and Returns

To distinguish rational learning from naïve reinforcement learning, we examine whether the bidder's return performance in subsequent IPO auctions increases with experience (Hypothesis 2a) or decreases with experience (Hypothesis 3a).

4.1 Univariate tests

We use *auction order* to measure bidding experience. Each bidder-auction is assigned an *auction order.* An auction is a bidder's first (second, third, etc.) auction if the bidder has 0 (1, 2, etc.) previous IPO auctions. Thus, auction order is a bidder's number of past auctions plus one. A given auction may be one bidder's third auction but another bidder's first auction. An auction is counted as a previous auction if its first non-hit day occurs before the current auction's auction date (so that we can compute the initial return from the previous auction).[11]

To examine bidder returns as auction order increases, Table 3 reports the mean (and median) return of bidders in their first, second, third, and higher-order auctions. Panel A reports results for the entire sample. For individual bidders, returns in general decrease when investors' auction order increases from second to subsequent auctions. We test the significance of the differences in returns for second versus first, third versus second, fourth versus third, fifth versus fourth, and higher-order versus fifth auctions. The differences are significantly negative for third versus second, fourth versus third, and higher-order versus fifth auctions. These results suggest that individual investors' returns tend to deteriorate with experience in bidding.[12] Panel B in Table 3 reports the mean (and median) returns of bidders excluding the outlier auction described in footnote 12. With this exclusion, for individual investors, the average return for second auctions is significantly lower than for first auctions, by 2.88%, and again returns steadily decrease as the auction order increases. These findings indicate that greater bidding experience is associated with *worse* return performance. This is consistent with Hypothesis 3a based upon naïve reinforcement learning, and inconsistent with the predictions of Hypothesis 2a based upon rational learning.[13]

As for institutional investors, returns are much more stable across auction orders. Table 3, Panel B, shows that the differences in the mean returns are all insignificant for third versus second, fourth versus third, fifth versus fourth, and higher-order versus fifth. The only exception is that the mean return for second auctions is significantly lower than for first auctions. For median returns, none of the differences are significant. Taken together with the previous result that institutions' bidding probability is not significantly influenced by their past returns, our findings suggest that institutions' bidding decisions (whether to bid and how high to bid) depend much less on their own experience.

---

11 Under this definition, a bidder may have multiple auctions that have the same auction order. For example, if a bidder participates in a total of three auctions, and for the first two, the periods from auction date to first non-hit day overlap partially. In this case, the bidder's auctions have the following auction order, respectively: first, first, and third.

12 A notable exception is that the average return for second auctions is significantly higher than for first auctions. This particular result, however, is driven entirely by one outlier auction, Chung Hwa Telecom, which is a privatization of a state-owned company and the biggest IPO in Taiwan ever (including other IPO methods). Chung Hwa has an asset value of NT$443.1 billion compared to NT$10.4 billion for the average IPO firm in our sample, and it raises NT$22.7 billion compared to the average NT$0.9 billion. It also attracts the highest number of bidders: 4,286 compared to an average of 708.7 bidders, among whom 4,240 (98.9%) are first-time bidders. In addition, all bidders in this case won shares and all of them received negative returns ranging from −37.6% to −3.4%. Not surprisingly, inclusion of this auction greatly reduces the mean return of first-time bidders.

13 Although steadily declining, returns remain positive even in individual bidders' sixth or higher-order auctions. However, declining returns are inconsistent with rational learning. Moreover, the slightly positive returns for individual bidders' high-order auctions do not necessarily mean that these investors benefit from participating in more auctions. According to IPO auction theory, positive returns are required to compensate for information and participation costs (Sherman 2005). It is unclear whether the returns these experienced bidders receive are adequate compensation for these costs.

Table 3
Bidder returns by auction order

Panel A. Whole samples

	Individuals					Institutions				
	Mean			Median		Mean			Median	
Auction order	*N*	Return	*n*th-(*n*-1)th	Return	*n*th-(*n*-1)th	*N*	Return	*n*th-(*n*-1)th	Return	*n*th-(*n*-1)th
1	10,366	5.22%		-3.55%		613	11.45%		3.48%	
2	1,559	10.89%	5.67%***	7.80%	11.36%***	108	12.26%	0.81%	7.83%	4.35%***
3	870	7.76%	-3.13%***	4.75%	-3.05%***	63	14.83%	2.57%	5.85%	-1.97%
4	547	4.75%	-3.01%***	2.82%	-1.93%***	50	13.30%	-1.53%	6.33%	0.48%
5	424	5.02%	0.27%	2.38%	-0.44%	42	13.01%	-0.29%	6.49%	0.16%
≥6	2,271	2.48%	-2.54%**	1.59%	-0.79%	95	7.56%	-5.46%	4.99%	-1.50%

Panel B. Excluding the largest auction

	Individuals					Institutions				
	Mean			Median		Mean			Median	
Auction order	*N*	Return	*n*th-(*n*-1)th	Return	*n*th-(*n*-1)th	*N*	Return	*n*th-(*n*-1)th	Return	*n*th-(*n*-1)th
1	6,307	14.21%		10.00%		422	19.83%		9.63%	
2	1,524	11.34%	-2.88%***	8.33%	-1.68%	107	12.41%	-7.43%***	7.85%	-1.78%
3	870	7.76%	-3.58%***	4.75%	-3.57%***	63	14.83%	2.43%	5.85%	-2.00%
4	547	4.75%	-3.01%***	2.82%	-1.93%***	50	13.30%	-1.53%	6.33%	0.48%
5	424	5.02%	0.27%	2.38%	-0.44%	42	13.01%	-0.29%	6.49%	0.16%
≥6	2,271	2.48%	-2.54%**	1.59%	-0.79%	95	7.56%	-5.46%	4.99%	-1.50%

This table shows the mean and median values of bidder initial returns by auction order. The column labeled as (*n*th-(*n*-1)th) shows the difference in mean or median returns in *n*th and (*n*-1)th auction order. We use *t*-test for differences in means, and Wilcoxon-Mann-Whitney test for differences in medians. ***, **, and * denote that the difference is significant at the 1%, 5%, and 10% levels, respectively.

To ensure that the declining returns across auction orders are indeed due to differences in bidding experience rather than differences in the characteristics of bidders who engage in different numbers of auctions, we compute a bidder's *own* change in returns across auctions. Results are consistent with those in Table 3 (see Online Appendix, Table A2).

In summary, we find that individual returns steadily decline as they gain more experience, consistent with Hypothesis 3a implied by naïve reinforcement learning. In sharp contrast, there is little evidence that institutions are subject to such naïve learning.

An alternative explanation for the evidence that returns decrease for individuals as they gain bidding experience is that the average auction returns happen to decrease over time during our sample period. However, several data oppose this. First, the average returns from auctions do not steadily decrease. As can be seen in Panel B of Table 1, the average bidder performs best in 1996, which is followed by a bad year in 1997. A similar mini-cycle is repeated in 1999 and 2000. Second, this argument does not explain why institutional returns *across auction orders* do not exhibit the same pattern. The time-series pattern of the *average institutional return* is very similar to that of the individuals. Finally, we performed robustness checks by excluding the four auctions with the highest and lowest average returns, or excluding auctions in year 2000 (in which most auctions have negative average returns; overall it is the worst year during the sample period). The results are similar to those in Table 3.

4.2 Multivariate tests

In order to control for other factors that may affect returns, we run regressions of bidders' returns on the experience variable, natural log of auction order, and a set of controls based on those in Chiang, Qian, and Sherman (2010). We control for the following firm-specific characteristics: firm size measured as the natural logarithm of assets, VC ownership in the firm, P/E ratio measured as the auction's reservation price over earnings per share, percentage of shares auctioned, a dummy equal to one if the firm is in a high-tech industry, and a dummy equal to one if the firm is traded on the Taiwan Stock Exchange (TSE) as opposed to on the OTC. To control for market conditions, we include market volatility in the three months prior to the auction, and recent auction return computed as the weighted average of the returns of previous IPO auctions with weights based on (720-N), where N is the number of days between a previous auction's first non-hit day and the recent auction's auction day, as well as year dummies.[14,15]

We also include unexpected entry (i.e., unexpected number of bidders) of institutions and individuals as well as the average bid premium of institutions and

[14] The return of an IPO auction refers to the value-weighted average return of all winning bidders in the auction.

[15] We lose the first two auctions in our sample for including the variable *recent auction return*. We lose the second auction because its auction date occurred before the first IPO shares began to trade.

Table 4
Regressions: The effects of experience on returns

	Overall sample				Frequent Bidders			
	Individuals		Institutions		Individuals		Institutions	
Independent variables	Estimates	*t*	Estimates	*t*	Estimates	*t*	Estimates	*t*
Log(auction order)	-0.0155	-2.34**	0.0026	0.22	-0.0187	-2.55**	-0.0006	-0.06
Log (assets)	0.0574	3.07***	0.0861	3.90***	0.0723	3.35***	0.0871	3.08***
VC ownership	-0.0780	-0.64	-0.0442	-0.37	-0.0645	-0.51	-0.0561	-0.49
P/E	-0.0044	-3.09***	-0.0056	-3.44***	-0.0059	-3.98***	-0.0059	-3.06***
High-tech dummy	0.2184	4.59***	0.2750	5.64***	0.2283	4.98***	0.2385	5.19***
TSE dummy	-0.1893	-2.97***	-0.2340	-2.98***	-0.1954	-2.74***	-0.2346	-2.38**
% shares auctioned	2.4536	2.20**	3.2086	2.09**	2.2368	2.08**	2.9909	2.00*
Market volatility	27.7903	2.22**	47.7255	2.62**	28.4687	2.11**	50.0980	2.85***
Recent auction return	0.1234	0.39	0.2265	0.52	0.1029	0.33	0.2427	0.59
Unexpected entry of individuals	-0.1280	-1.99*	-0.1554	-2.52**	-0.1691	-2.57**	-0.2016	-3.12***
Bid premia of individuals	-0.5940	-2.10**	-1.0459	-2.83***	-0.7563	-2.14**	-1.1941	-3.11***
Unexpected entry of institutions	0.0992	2.57**	0.1354	4.16***	0.1189	3.11***	0.1288	3.95***
Bid premia of institutions	0.6214	2.14**	1.1154	2.85***	0.8480	2.33**	1.4400	3.80***
Intercept	-1.6173	-2.95***	-1.6219	-2.39**	-1.1880	-2.35**	-1.8739	-3.05***
Year dummies	yes		yes		yes		yes	
R^2	0.62		0.69		0.60		0.68	
N	15,820		971		7,880		545	

The dependent variable is a bidder's initial return in an auction. Frequent bidders are those whose highest auction order is more than one. t-statistics are adjusted for auction clustering and heteroscedasticity. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

individuals.[16] These variables can reflect either private information or investor sentiment about the IPO. They could also reflect other public information not captured by the other controls. To ensure that our results are not driven by the specific form of the unexpected entry measures, in unreported tests we used raw entry instead of unexpected entry. Alternatively, we excluded unexpected entry and bid premium of institutions and individuals from the tests. The findings are similar using either of these alternative sets of control variables.

Table 4 reports the results for regressions estimated separately for individual and institutional investors, and for all bidders versus frequent bidders, where frequent bidders are those whose highest auction order is greater than one. In each regression, we compute t-statistics with standard errors adjusted for auction clustering and heteroscedasticity. The coefficients on the control variables are consistent with those of Chiang, Qian, and Sherman (2010)—both

[16] Unexpected entry of institutions is measured as the residual from the following regression (following Chiang, Qian, and Sherman 2010, Table 3): log(number of institutional bidders) $= b_1 + b_2$*log(assets) + b_3*VC ownership + b_4*P/E ratio + b_5*High tech dummy + b_6*TSE dummy + b_7*% shares auctioned + b_8*Market volatility + b_9*recent auction return + ε. Unexpected entry of individuals is similarly measured.

individual and institutional returns increase with the unexpected entry and average bid premium of institutions, and decrease with unexpected entry and average bid premium of individuals, which they interpret as evidence that institutional investors bid based on information and shave bids optimally, whereas individual investors as a group do not bid optimally.

For individual bidders, the coefficient on our main test variable, bidder experience, is -0.0155 for all bidders and -0.0187 for frequent bidders only, both significant at the 5% level. As experience increases by one standard deviation (0.87 for all individual bidders and 0.95 for frequent individual bidders), the expected return decreases by 1.3% for the sample of all individual bidders and by 1.8% for the sample of frequent individual bidders.

In contrast, experience has no effect on returns for institutional investors. The coefficient on experience is 0.0026 ($t = 0.22$) for the sample of all institutional investors and -0.0006 ($t = -0.06$) for the sample of frequent institutional bidders. The difference in the coefficient on experience between individual and institutional investors (for either the "all" or "frequent bidders" samples) is significant at the 5% level.

In an alternative regression specification, we define the dependent variable to be a positive return dummy, which equals one if the bidder earns a positive return and zero for a negative return, and estimate a logit regression of this dummy on experience. The results lead to similar conclusions about the effect of experience.

As additional robustness checks, we use dummies for different auction orders (second, third, fourth, and higher) instead of log(auction order) (see Online Appendix, Table A3, Panel A). We find that for individual investors, the coefficients on these dummies are all negative, suggesting that returns of bidders' higher-order auctions are lower than those of their first auctions. For example, the coefficient on the dummy for the second auction is -0.0299, suggesting that the average return decline is 2.99% from first to second auctions, which is very similar to the univariate result (Table 3, Panel B). With five test variables, the degrees of freedom and therefore the power to identify effects is reduced. Nevertheless, for the full sample of individual bidders, the coefficients are significant for dummies of auction order $= 2$ and ≥ 6; for frequent individual bidders, the coefficients are significant on all auction order dummies. In contrast, for both samples of institutional investors, none of the coefficients on auction order dummies is significant.

We also performed tests that include auction fixed effects (see Online Appendix, Table A3, Panel B). This removes variation in investors' return that derives from their decisions about which auctions to participate in, leaving only variations in return that derive from the choice of how aggressively to bid within an auction. As before, we find that the coefficients on auction order variables are in general significantly negative for samples of individual investors, but they are insignificant for samples of institutional investors. The magnitudes of these coefficients for samples of individuals become smaller, which can be explained by the fact that auction fixed effects remove the effects of bidders' auction selection ability.[17]

In summary, our results show that returns to individual investors decrease as they participate in more auctions, whereas the returns to institutional investors do not. For individual investors, the evidence is consistent with naïve reinforcement learning (Hypothesis 3a), and inconsistent with rational learning (Hypothesis 2a).[18] In sharp contrast, there is little evidence of naïve reinforcement learning on the part of institutional investors.

5. Effects of Experience on Auction Selection and Bid Aggressiveness

In this section, we explore the effect of experience on two types of bidding skills that contribute to returns: auction selection and bid-shaving (inverse of bid aggressiveness). We therefore examine Hypothesis 2b versus Hypothesis 3b/3c.

5.1 Auction selection ability

We regress the measure of auction selection ability—other bidders' return (i.e., the value-weighted average return to other bidders in an auction in which bidder i bids)—on bidder i's experience measured by the natural log of auction order.[19] If auction selection ability improves with experience, we expect to see a positive regression coefficient on experience. We include the same control variables as those in Table 4.

A bidder who is more certain about her information about firm quality should bid more aggressively and should be more likely to win shares. Hence these auctions should better reflect her selection ability than auctions in which she participates but does not win. We therefore include only auctions in which the current bidder (bidder i) wins. We also estimate the same regressions including all auctions in which the current bidder (bidder i) bids (but does not necessarily win). Results (available in the Online Appendix, Table A4, Panel A) support the same conclusions.

Table 5 reports the regression results for individual investors, institutional investors, frequent individual investors, and frequent institutional investors separately. For individual investors, the coefficient on experience is significantly negative for both samples, suggesting that their selection ability declines with

[17] This suggests, as documented more formally in Section 5, that return declines of individual investors are partly due to deterioration in auction selection abilities.

[18] There are other theories of bias in learning processes, such as the hypothesis that investors overweight information that suggests that they are skillful (as modeled in Daniel, Hirshleifer, and Subrahmanyam 1998; and Gervais and Odean 2001). We do not rule out the possibility that other forms of learning bias could help explain our findings. Even if so, investors would still be learning to be less successful instead of learning to succeed as implied by rational learning.

[19] Results are robust with respect to the use of an alternative measure for bidder i's auction selection ability, the ratio of the closing price on the first non-hit day over the reservation price in the auction in which she participates. This measures the "underpricing" of the shares at the reservation price.

Table 5
The effects of experience on auction selection ability

	Overall Samples				Frequent Bidders			
	Individuals		Institutions		Individuals		Institutions	
Independent variables	Estimates	*t*	Estimates	*t*	Estimates	*t*	Estimates	*t*
Log(auction order)	-0.0119	-1.82 *	0.0021	0.19	-0.0152	-2.11 **	-0.0025	-0.27
Log (assets)	0.0565	3.07 ***	0.0848	3.86 ***	0.0694	3.28 ***	0.0847	3.01 ***
VC ownership	-0.0748	-0.63	-0.0459	-0.37	-0.0469	-0.38	-0.0424	-0.37
P/E	-0.0043	-3.07 ***	-0.0056	-3.50 ***	-0.0058	-3.96 ***	-0.0058	-3.10 ***
High-tech dummy	0.2205	4.73 ***	0.2776	5.65 ***	0.2307	5.17 ***	0.2421	5.36 ***
TSE dummy	-0.1826	-2.90 ***	-0.2262	-2.85 ***	-0.1841	-2.62 **	-0.2213	-2.25 **
% shares auctioned	2.4219	2.17 **	3.0683	2.00 **	2.1202	1.96 *	2.8543	1.90 *
Market volatility	-0.1328	-2.17 **	-0.1530	-2.53 **	-0.1762	-2.83 ***	-0.1986	-3.24 ***
Recent auction return	-0.5824	-2.14 **	-1.0776	-2.86 ***	-0.7480	-2.23 **	-1.2365	-3.11 ***
Unexpected entry of individuals	0.1006	2.71 ***	0.1334	4.13 ***	0.1200	3.25 ***	0.1285	4.00 ***
Bid premia of individuals	0.6015	2.16 **	1.1504	2.90 ***	0.8293	2.40 **	1.4825	3.78 ***
Unexpected entry of institutions	26.0815	2.10 **	46.6306	2.61 **	26.5065	2.01 **	49.0621	2.84 ***
Bid premia of institutions	0.1308	0.42	0.1904	0.44	0.1175	0.38	0.2284	0.57
Intercept	-1.5636	-2.93 ***	-2.4767	-3.53 ***	-1.7722	-3.34 ***	-2.7160	-4.37 ***
Year dummies	yes		yes		yes		yes	
R^2	0.64		0.71		0.62		0.70	
N	15,820		971		7,880		545	

The dependent variable is the quantity-weighted average return of other bidders in an auction the current bidder wins. Frequent bidders are those whose highest auction order is greater than one. *t*-statistics are adjusted for auction clustering and heteroscedasticity. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

experience. For both samples of institutional investors, the coefficient on experience is insignificant, suggesting that their selection abilities do not decline with experience. The coefficient on experience is more negative for the sample of all (frequent) individual bidders than for all (frequent) institutional bidders, with the difference significant at the 5% (10%) level.

As a robustness check, we also estimate auction-level regressions (see Online Appendix, Table A4, Panel B), in which the dependent variable is an auction's weighted average return. The main variables of interest are the average experiences of individual and institutional bidders in the auction, measured as the natural logarithm of one plus the average number of previous auctions for individual (institutional) investors. Control variables as in Table 4 are included. Consistent with the results in Table 5, auction selection abilities are negatively related to the average experience of individual bidders, but are not significantly related to that of institutional bidders.

5.2 Bid aggressiveness

We measure bid aggressiveness (i.e., an investor's tendency to bid a high price) as the percentile of a bidder's bid price out of all bids (including losing bids)

Table 6
Regressions: The effects of experience on bid aggressiveness

	Overall Sample				Frequent Bidders			
	Individuals		Institutions		Individuals		Institutions	
Independent variables	Estimates	*t*	Estimates	*t*	Estimates	*t*	Estimates	*t*
Log(auction order)	5.9512	8.41***	0.2558	0.27	3.0045	5.10***	-2.6188	-2.41**
Log (assets)	0.4633	2.20**	-1.2905	-1.35	1.1729	2.49**	1.9091	1.58
VC ownership	-0.9743	-0.90	4.6828	0.61	1.0118	0.31	-4.9323	-0.66
P/E	-0.0122	-1.08	0.0567	1.25	-0.004	-0.14	-0.0396	-0.89
High-tech dummy	-0.3022	-0.90	3.3359	1.42	2.3774	2.25**	3.7297	1.50
TSE dummy	-0.2277	-0.41	-0.5752	-0.22	1.018	1.00	-1.2576	-0.48
% shares auctioned	11.3133	1.06	-44.8958	-0.87	-13.4392	-0.65	24.6489	0.54
Market volatility	1.0576	2.09**	6.0325	2.39**	1.4617	1.55	6.7588	3.22***
Recent auction return	1.5418	0.66	-99.2375	-5.45***	-15.6421	-2.59**	-115.4919	-6.51***
Unexpected entry of individuals	-0.4929	-1.71*	-1.6779	-0.77	0.5295	0.83	-0.5784	-0.33
Bid premia of individuals	-2.4408	-0.97	103.2498	5.74***	15.3264	2.65**	113.8192	6.66***
Unexpected entry of institutions	-65.448	-0.86	572.8137	1.15	-422.2514	-2.31**	312.6872	0.81
Bid premia of institutions	3.8001	1.95*	11.7931	0.83	13.4022	3.19***	22.3685	1.99*
Intercept	42.1041	8.64***	56.1617	2.31**	35.5223	3.95***	15.0381	0.64
Year dummies	yes		yes		yes		yes	
R^2	0.02		0.10		0.02		0.11	
N	56,165		2,682		27,886		1,731	

The dependent variable is the percentile of a bidder's bid price in an auction. Frequent bidders are those whose highest auction order is more than one. *t*-statistics are adjusted for auction clustering and heteroscedasticity. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

in an auction.[20] Results are similar when bid aggressiveness is measured as auction-mean (or median) adjusted bid premia.[21] We relate bid aggressiveness to the bidder's experience, and include the same control variables as in Table 5.

Table 6 reports the regression results. For individual bidders, the coefficient on experience is 5.95 for all bidders and 3.00 for frequent bidders only, both significant at the 1% level. As the explanatory variable increases by one standard deviation (0.87 for the sample of all individual bidders and 0.95 for the sample of frequent individual bidders), an individual bidder's bid increases by 5.2 percentiles for the sample of all individual bidders and by 2.9 for the sample of frequent individual bidders. In the average auction, one percentile increase in bid price (relative to clearing price) has a mean of 0.6%.[22]

In contrast, institutional investors do not bid more aggressively as they participate in more auctions. The coefficient on experience is insignificant for the

[20] For bidders with multiple bids, we use their quantity-weighted average bid price.

[21] As these are discriminatory or pay-what-you-bid auctions, there is no free-rider problem. In uniform price or "Dutch" auctions, on the other hand, the free-rider problem might lead some investors to bid unrealistically high amounts simply to be "first in line" (see Sherman 2005).

[22] The smaller coefficient for the frequent bidder sample suggests that they on average tend to bid more aggressively than one-time bidders regardless of their bidding histories.

sample of all institutional investors, and significantly negative for the sample of frequent institutional bidders. The difference in the coefficient on experience between the samples (all or frequent bidders) of individual versus institutional investors is significant at the 1% level. Furthermore, in regressions using auction order dummies instead of log(auction order), similar conclusions apply.

In summary, we find evidence that individual investors become less selective in entering auctions as they gain more experience. Moreover, they tend to bid more aggressively in their later auctions. These results are consistent with the predictions of naïve reinforcement learning (Hypotheses 3b and 3c).[23] Institutional investors' bidding abilities, in contrast, do not deteriorate as they participate in more auctions.

5.3 Dollar profits

An alternative explanation for individual investors' declining returns with experience is that they bid more aggressively to obtain a greater allocation of shares and therefore earn higher dollar profits despite lower percentage returns. This hypothesis does not explain the evidence that auction selection ability declines with experience. Nevertheless, we now examine the relationship between dollar profits and bidder experience.

Table 7 reports the mean and median values of bid size (in Panel A), investment size (i.e., winning size) (in Panel B), and dollar profits (in Panel C) by auction order, excluding the outlier auction Chung Hwa. For individual investors, we saw some evidence that bid size increases with auction order, but there is no evidence that their investment size (or winning size) increases.[24]

In Table 7, Panel C suggests that individual investors' dollar profits decrease as investors' experience increases. For example, the mean profit of their first auction is NT$543,400 while that of their second auction is NT$329,540 (a 39% decrease), the difference significant at the 1% level. The median dollar profits also decrease with auction order, with the differences significant at the 5% level for third versus second (a difference of NT$23,650), fourth versus third (a difference of NT$16,250), and higher-order versus fifth (a difference of NT$4,950). These differences are economically meaningful compared to the median profits of NT$38,950 for the sample of individuals.[25] In contrast, for institutional investors, we do not observe any obvious patterns in bid size, investment size, or profits as auction order increases.

23 The bid aggression result is potentially consistent with rational learning as well, although this is not one of our formal hypotheses (see footnote 6).

24 This may suggest that experienced bidders crowd larger bids into later auctions and make winning more difficult. However, neither the total subscription ratio nor the number of investors participating in an auction generally increases over time.

25 This median value (NT$38,950) is different from that in Table 1, Panel D (NT$-17,500), because: (1) Here (in Table 7) we exclude the outlier auction of Chung Hwa. Including Chung Hwa leads to a sample median of zero for all auction-bidder observations. (2) Instead of using auction-bidder observations, Panel D in Table 1 uses unique-bidder observations (i.e., first averaging across auctions for each bidder, and then averaging across bidders).

Table 7
Bid size, investment size, and dollar profits by auction order

Panel A. Bid size (in thousands of NT$)

	Individuals					Institutions				
	Mean			Median		Mean			Median	
Auction order	*N*	Bid size	*n*th-(*n*-1)th	Bid size	*n*th-(*n*-1)th	*N*	Bid size	*n*th-(*n*-1)th	Bid size	*n*th-(*n*-1)th
1	34,146	2,245.94		603.77		1,324	19,833.36		12,046.15	
2	6,129	2,336.16	90.22	759.38	155.61***	366	20,320.75	487.39	11,214.13	-832.02
3	2,987	2,555.65	219.49	843.82	84.44***	187	18,309.11	-2011.64	10,631.25	-582.88
4	1,770	2,551.47	-4.18	850.43	6.61	166	18,341.58	32.47	10,848.18	216.93
5	1,317	2,591.62	40.15	803.54	-46.89	120	18,129.11	-212.47	12,224.11	1,375.93
≥6	6,406	2,809.51	217.89	924.60	121.06***	332	17,322.72	-806.39	13,127.74	903.63

Panel B. Investment size (in thousands of NT$)

	Individuals					Institutions				
	Mean			Median		Mean			Median	
Auction order	*N*	Investments	*n*th-(*n*-1)th	Investments	*n*th-(*n*-1)th	*N*	Investments	*n*th-(*n*-1)th	Investments	*n*th-(*n*-1)th
1	6,307	3,435.72		878.83		422	18,432.43		9,876.46	
2	1,524	2,758.46	-677.26***	762.19	-116.64*	107	17,133.54	-1,298.89	10,631.25	754.79
3	870	2,751.72	-6.74	795.05	32.86	63	12,098.96	-5,034.58	7,971.41	-2,659.84
4	547	2,643.06	-108.66	833.06	38.01	50	18,764.20	6,665.24	11,893.14	3,921.73**
5	424	3,216.36	573.30	818.17	-14.89	42	14,501.37	-4,262.83	11,556.65	-336.49
≥6	2,271	2,689.22	-527.14	817.97	-0.20	95	15,430.62	929.25	11,694.63	137.98

(*continued*)

Table 7
Continued

Panel C. Dollar profits (in thousands of NT$)

	Individuals					Institutions				
	Mean			Median		Mean			Median	
Auction order	*N*	Dollar profits	*n*th-(*n*-1)th	Dollar profits	*n*th-(*n*-1)th	*N*	Dollar profits	*n*th-(*n*-1)th	Dollar profits	*n*th-(*n*-1)th
1	6,307	543.40		60.75		422	4,610.81		950.80	
2	1,524	329.54	-213.86***	56.94	-3.81	107	3,115.64	-1,495.17	733.69	-217.11
3	870	260.30	-69.24	33.29	-23.65***	63	2,453.06	-662.58	273.38	-460.31*
4	547	236.88	-23.42	17.04	-16.25**	50	3,441.28	988.22	778.48	505.11
5	424	317.32	80.44	13.05	-3.99	42	1,778.34	-1,662.94	725.06	-53.43
≥6	2,271	147.46	-169.86**	8.10	-4.95**	95	1,306.03	-472.31	412.55	-312.51

This table shows the mean and median values of bidders' bid size, investment size, and dollar profits by auction order. The sample excludes the largest auction, Chung Hwa. The column labeled as (*n*th-(*n*-1)th) shows the difference in mean or median returns in *n*th and (*n*-1)th auction order. We use *t*-tests for differences in means, and Wilcoxon-Mann-Whitney tests for differences in medians. ***, **, and * denote that the difference is significant at the 1%, 5%, and 10% levels, respectively.

We then estimate a multivariate regression of dollar profits on bidder experience, controlling for auction characteristics and market conditions. Since profits are highly skewed, we use median regression, which is less sensitive to non-normality than OLS regressions. We find that individual investors' profits decrease with bidder experience, whereas institutional investors' profits do not depend on experience. Our estimates show that a one-standard-deviation increase in experience leads to a drop of NT$4,630 in median profits, which is substantial compared to the sample median of NT$38,950.[26]

The evidence that individual investors' dollar profits decrease with experience is consistent with naïve reinforcement learning leading to suboptimal bidding strategies in subsequent auctions. It is not consistent with the alternative hypothesis that individual investors follow optimal bidding strategies, which would imply non-decreasing profits.

6. Do Bidders Ever Learn to Be Rational?

A natural question is whether the skill deterioration phenomenon we document ever reverses or at least levels off, given enough experience. We would expect leveling off to occur, since experiments on reinforcement learning find that the learning curve is initially steep and then flattens as experience increases (Roth and Erev 1995). In this section, we further examine whether and/or when bidders in IPO auctions stop naïve learning.

For this purpose, we add a quadratic term of experience to our main tests (those in Tables 4, 5, and 6). That is, we run regressions of returns, auction selection ability, and bid aggressiveness on *auction order* and *auction order squared*. If individual bidders' naïve learning levels off or even reverses itself, we expect the coefficient on the squared term to be of the opposite sign to that on the linear term. If the strength of naïve learning does not decline, the coefficient on the squared term should be insignificant.

Regression results suggest that individual bidders may eventually learn.[27] For the sample of all individual bidders, returns are negatively associated with *auction order* (the coefficient is −0.0048) and positively related with *auction order squared* (the coefficient is 0.0001). This suggests that returns first decline with experience, but that this effect levels off, and may eventually turn to increase with experience. The point at which the marginal effect of *auction order* becomes zero is when the auction order is 24 (i.e., when $-0.0048 + 2*0.0001*\textit{auction order} = 0$). It is of course hard to know how well the quadratic specification continues to fit for very high-order auctions. Taken at face value, this calculation implies that it takes around 24 auctions for individual bidders to reverse their naïve reinforcement learning and start to benefit

[26] Results of median regressions are tabulated in the Online Appendix, Table A5. Unreported OLS regression results lead to the same conclusions. In OLS regressions, as experience increases by one standard deviation, the mean dollar profits decrease by NT$117,157 for individuals. In comparison, the sample mean of dollar profits for individuals is NT$398,130.

[27] Regression results for the main variables of interest are tabulated in the Online Appendix.

from experience. This is a considerable amount of experience; in our sample, only 0.2% of individual bidders' highest auction orders are at least 24. The results are very similar for the sample of frequent individual bidders. In contrast, institutional investors' returns are not significantly related to either *auction order* or *auction order squared.*

In addition, our estimates suggest that individual investors' auction selection ability decreases with experience until auction order reaches 14 (for the sample of all individual bidders) or 19 (for the sample of frequent individual bidders). Their bid aggressiveness increases with experience until auction order reaches 20 for both samples of individual investors. In contrast, institutional investors' auction selection ability and bid aggressiveness do not depend on experience.

Consistent with theory, individual investors' naïve reinforcement learning levels off. These bidders may eventually learn to be more rational, given enough experience; however, for most bidders who have participated in limited numbers of IPO auctions, past successful experience leads them to unduly optimistic expectations and causes them to follow suboptimal bidding strategies. Institutional investors' bidding strategies and performance, on the other hand, do not seem to depend on past experience.

7. Conclusion

We examine how bidding experience affects a bidder's decision to participate in an IPO auction, as well as the resulting return performance. Individual investors are more likely to bid in the future if they receive high returns from their previous IPO auctions. However, their returns steadily decline as they participate in more auctions. For example, the mean return decreases by −2.88% from their first to second deals, and mean dollar profit decreases by $NT 213,860 (a 39% decrease). Furthermore, as individual investors gain more experience, their auction selection ability deteriorates, and they become more aggressive in the levels of their bids.

This evidence indicates that individual investors are subject to naïve reinforcement learning. When individual investors receive high returns from previous auctions, they become more optimistic about receiving high returns from future auctions, making them less selective in participating in future auctions and more aggressive in their bidding.

In sharp contrast, there is little sign that institutional bidders are subject to this bias. Their decisions to participate in an IPO auction are unrelated to their past returns. Furthermore, their returns do not decline with experience as they bid in more auctions, and their auction selection and bid-shaving abilities do not deteriorate with experience.[28]

[28] We do not find evidence of rational learning for institutional investors either. A possible reason is that institutions already bid optimally, and hence do not improve with experience. Results in Chiang, Qian, and Sherman (2010) are consistent with this notion. Alternatively, institutions may also be subject to naïve reinforcement learning, but not as strongly as individual investors. If so, even though institutions avoid losing skill through experience, they also do not improve as much as might otherwise be expected.

Teachers often tell their students that the important thing that education provides is not a given set of facts, but an ability to “learn how to learn.” Overall our findings indicate that individual investors do not know how to learn from their past experience (at least not from even fairly long experience), and their naïve overweighting of personal experience causes their performance to decline. However, a more optimistic message from this article is that it may be *possible* for investors to avoid this deterioration in skill, as reflected in the fact that institutional investors do not seem to be subject to naïve reinforcement learning. A direction for future research is to understand what it is about institutional managers that makes them better at avoiding pernicious pseudo-learning, and whether there are ways to help individual investors, despite their more limited resources, to train themselves to improve their learning skills.

Appendix

Table 8
Variable definitions

Variables	Definition
Auction order	An auction is a bidder’s first (second, third, etc.) auction if the bidder has 0 (1, 2, etc) previous IPO auctions. An auction is counted as a previous auction if its first non-hit day occurs before the current auction’s auction date.
Initial return	The closing price of the first non-hit day over the winning price minus one.
VC ownership	The percentage of shares held by venture capitalists prior to the IPO.
P/E	The ratio of the reservation price of the auction to the annual earnings per share prior to the IPO.
High-tech dummy	The dummy equals one if the firm is categorized as in electronic sector by the exchange and zero otherwise.
TSE dummy	The dummy equals one if the firm is listed on Taiwan Stock Exchange and zero otherwise.
% of shares auctioned	The number of shares to be auctioned divided by the total number of shares outstanding.
Market volatility	The standard deviation of daily market returns during the three months prior to the auction day.
Recent auction return	The weighted average initial return of IPO auctions for which returns have been observed, with weights based on $(720 - N)$ (zero weight if $720 - N < 0$), where N is the number of days between a previous auction’s first non-hit day and the current auction’s auction day. An auction is counted as a previous auction if its first non-hit date occurs before the current auction’s auction date.
Unexpected entry of institutions (individuals)	The unexpected number of institutional (individual) bidders as constructed in Chiang, Qian, and Sherman (2010).
Bid premium of institutions (individuals)	The quantity-weighted average bidding price of all institutional (individual) bids (including losing bids) relative to the reservation price.

Supplementary Data

Supplementary data are available online at http://www.sfsrfs.org/addenda/.

References

Arrow, K. 1962. The Economic Implications of Learning by Doing. *Review of Economic Studies* 29:155–73.

Arthur, B. W. 1991. Designing Economic Agents That Act Like Human Agents: A Behavioral Approach to Bounded Rationality. *American Economic Review* 81:353–59.

Barber, B. M., T. Odean, and M. Strahilevitz. 2010. Once Burned, Twice Shy: Naïve Learning, Counterfactuals, and the Repurchase of Stocks Previously Sold. Unpublished paper, University of California, Berkeley.

Billett, M., and Y. Qian. 2008. Are Overconfident Managers Born or Made? Evidence of Self-attribution Bias from Frequent Acquirers. *Management Science* 54:1037–51.

Camerer, C. F., and T.-H. Ho. 1999. Experience-weighted Attraction Learning in Normal Form Games. *Econometrica* 67:827–74.

Charness, G., and D. Levin. 2005. When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect. *American Economic Review* 95:1300–9.

Chiang, Y., Y. Qian, and A. Sherman. 2010. Endogenous Entry and Partial Adjustment in IPO Auctions: Are Institutional Investors Better Informed? *Review of Financial Studies* 23:1200–30.

Choi, J., D. Laibson, B. Madrian, and A. Metrick. 2009. Reinforcement Learning and Savings Behavior. *Journal of Finance* 64:2515–34.

Cross, J. G. 1973. A Stochastic Learning Model of Economic Behavior. *Quarterly Journal of Economics* 87:239–66.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam. 1998. Investor Psychology and Security Market Under- and Over-reactions. *Journal of Finance* 53:1839–85.

Dhar, R., and N. Zhu. 2006. Up Close and Personal: An Individual-level Analysis of the Disposition Effect. *Management Science* 52:726–40.

Erev, I., and A. E. Roth. 1998. Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria. *American Economic Review* 88:848–81.

Feng, L., and M. Seasholes. 2005. Do Investor Sophistication and Trading Experience Eliminate Behavioral Biases in Financial Markets? *Review of Finance* 9:305–51.

Gervais, S., and T. Odean. 2001. Learning to Be Overconfident. *Review of Financial Studies* 14:1–27.

Grossman, S. J., R. E. Kihlstrom, and L. J. Mirman. 1977. A Bayesian Approach to the Production of Information and Learning by Doing. *Review of Economic Studies* 44:533–47.

Hilary, G., and L. Menzly. 2006. Does Past Success Lead Analysts to Become Overconfident? *Management Science* 52:489–500.

Jacob, J., T. Lys, and M. Neale. 1999. Expertise in Forecasting Performance of Security Analysts. *Journal of Accounting and Economics* 28:51–82.

Jagannathan, R., A. Jirnyi, and A. E. Sherman. 2010. Why Don't Issuers Choose IPO Auctions? The Complexity of Indirect Mechanisms. Unpublished paper, Northwestern University.

Kaustia, M., and S. Knupfer. 2008. Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions. *Journal of Finance* 63:2679–702.

Linnainmaa, J. 2010. Why Do (Some) Households Trade So Much? Unpublished paper, University of Chicago.

Mahani, R., and D. Bernhardt. 2007. Financial Speculators' Underperformance: Learning, Self-selection, and Endogenous Liquidity. *Journal of Finance* 62:1313–40.

Malmendier, U., and S. Nagel. 2009. Learning from Inflation Experiences. Unpublished paper, University of California at Berkeley.

Mikhail, M., B. Walther, and R. Williams. 1997. Do Security Analysts Improve Their Performance with Experience? *Journal of Accounting Research* 35:120–31.

Nicolosi, G., L. Peng, and N. Zhu. 2009. Do Individual Investors Learn from Their Trading Experience? *Journal of Financial Markets* 12:317–36.

Roth, A. E., and I. Erev. 1995. Learning in Extensive-form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. *Games and Economic Behavior* 8:164–212.

Seru, A., T. Shumway, and N. Stoffman. 2009. Learning by Trading. *Review of Financial Studies* 23:705–39.

Sherman, A. 2005. Global Trends in IPO Methods: Book Building versus Auctions with Endogenous Entry. *Journal of Financial Economics* 78:615–49.

Tversky, A., and D. Kahnerman, 1971. Belief in the Law of Small Numbers. *Psychological Bulletin* 76:105–10.