The Effects of Experience on Investor Behavior: Evidence from India's IPO Lotteries*

Santosh Anagol† Vimal Balasubramaniam‡ Tarun Ramadorai§

March 2015

Abstract

We exploit the randomized allocation of stocks in 57 Indian IPO lotteries to 1.7 million investors between 2007 and 2012 to provide new estimates of the causal effect of investment experiences on future investment behavior. We find that investors experiencing exogenous gains in IPO stocks (the treatment) are more likely to apply for future IPOs, increase trading in their portfolios, exhibit a stronger disposition effect, and tilt their portfolios towards the sector of the treatment IPO. Treatment effects are stronger for smaller accounts, and inherit the sign of the IPO first-day return.

*We gratefully acknowledge the Alfred P. Sloan Foundation for financial support, and the use of the University of Oxford Advanced Research Computing (ARC) facility for processing. We thank Sushant Vale for dedicated research assistance, and Samar Banwat, John Campbell, Rajesh Jatekar, Karthik Muralidharan, Nayana Ovalekar, Bhupendra Patel, Ben Ranish, P.S. Reddy, Shubho Roy, Susan Thomas, and seminar participants at the Saïd Business School and the IGIDR Emerging Markets Finance conference for useful comments and conversations.

†Anagol: Wharton School of Business, Business Economics and Public Policy Department, University of Pennsylvania, and Oxford-Man Institute of Quantitative Finance. Email: anagol@wharton.upenn.edu

‡Balasubramaniam: Saïd Business School, Oxford-Man Institute of Quantitative Finance, University of Oxford, Park End Street, Oxford OX1 1HP, UK. Email: vimal.balasubramaniam@sbs.ox.ac.uk

§Ramadorai: Saïd Business School, Oxford-Man Institute of Quantitative Finance, University of Oxford, Park End Street, Oxford OX1 1HP, UK and CEPR. Email: tarun.ramadorai@sbs.ox.ac.uk

1 Introduction

Workhorse economic models typically assume that agents have stable preferences and well-founded beliefs. In these models, preferences are "deep" parameters that are not influenced by states of the world, and beliefs are defined using all past data and updated according to Bayes rule. More recent work in economics, however, takes the view that the preferences and beliefs of individuals are more malleable. One interesting approach in this vein has been to model agents' preferences and beliefs as being particularly influenced by their own personal experiences (see, for example Kőszegi and Rabin, 2007; Rabin and Weizsäcker, 2009; Camerer and Ho, 1999; Roth and Erev, 1995; Ellison and Fudenberg, 1993).

Empirical validation for the effects of personal experience on economic decision-making has been growing, especially in the area of investments, where data on agents' choices involving risk is readily available. One strand of this emerging empirical literature relates personal economic experiences to long run risk-taking in financial markets, finding that experiences with low stock returns, inflation, and unemployment all have major effects on stock market participation even decades after the experience (Malmendier and Nagel, 2011, 2009; Knüpfer et al., 2014). A second strand studies how recent portfolio experiences shape short-term decisions, finding that savings and stock trading decisions appear to respond to investors' personal experiences.[1]

A fundamental challenge confronting empirical work in this area is the fact that most investment experiences are determined endogenously by the investor. For example, if we observe investors who have recently experienced gains and exhibit subsequent changes in investment behavior, we might be tempted to conclude that these return experiences have changed these investors' risk preferences. However, it is entirely possible in this scenario

[1] A large literature relates investors' experienced asset returns to future investment behavior. One strand focuses on how prior gains and losses affect risk-taking, see Thaler and Johnson (1990) for the first analysis and Gamble and Johnson (2014). Another strand looks at specific types of investor experiences in certain asset classes, see Andersen et al. (2014) for a recent review. For recent work using Indian data in this context, see Campbell et al. (2014).

that the initially experienced gains were themselves a result of an increase in risk-taking by the investor, caused by an unobservable change in the investor's risk preferences. Another plausible possibility is that the initially experienced gains may reflect changes in the investor's skill, an attribute which is notoriously difficult to measure.

Empirical work in this area has been careful to control for various investor and time characteristics in an attempt to isolate the experience-behavior relationship. However, it is ultimately impossible to test whether confounding factors have been suitably controlled for. The ideal research design in this case would be to find a setting in which investment experiences are randomly assigned to investors, and to then track how this random assignment of experience affects future behavior.[2]

This paper introduces a new research design for estimating the causal relationship between investor experiences and future behavior. We exploit the fact that (owing to excess demand) shares in initial public offerings (IPOs) are often allocated to retail investors using randomized lotteries. By comparing allocated versus non-allocated applicants, we can identify the causal effect of how the experience of IPO initial returns (which are often high, and vary substantially across IPOs) changes future investment behavior.

We apply this research design to India, where we have data from 57 different IPOs in which 1.7 million investor accounts experienced randomized allocation in lotteries between 2007 and 2012. For all 605,435 treatment and 1,093,969 control accounts, we are able to track the details of investment in their equity portfolios on a monthly basis both prior to and following treatment. Given the large number of IPO experiments that we observe, we also have substantial power to test how different types of return experiences affect investors. Moreover, we are able to estimate heterogenous treatment effects, i.e., estimate how investors'

[2]An assumption underlying many of the specifications estimated in the literature is that variations in expected returns due to risk-taking or skill are likely to be swamped by variations in unexpected returns caused by luck. However it is worth noting that this is simply an assumption, which can be tested if econometricians have access to truly random variation in investment gains and losses and are able to track outcomes in response to these random gains and losses.

responses to experience vary with investor characteristics such as the size of the pre-existing portfolio of the investor. To our knowledge, this is the first paper to estimate the causal effect of return experiences using the randomized allotment of real securities.

While our specific data and analysis focus on India, we also note that this research design could be applied to many countries that use lottery systems to allocate IPO shares, including Bangladesh, Brazil, China, Germany, Hong Kong, Singapore, Sweden, and Taiwan. In addition, several brokerages, such as TD Ameritrade and E-Trade in the United States, allocate shares to individual investors using random assignment;[3] our methodology could also be applied to data from such individual brokerages. We argue that this approach has the potential to produce a large set of credibly identified results on how short-term experiences affect the behavior of investors.

We begin our analysis by testing whether investors that are randomly allotted shares are more likely to apply for future IPOs. Our results confirm previous non-experimental results that personal experience in the IPO market appears to lead to reinforcement learning (see Kaustia and Knüpfer, 2008; Chiang et al., 2011). We find that conditional on applying for an IPO, investors who "win the IPO lottery," i.e., those that are randomly allocated IPO shares with positive returns, are significantly more likely to apply for future IPOs, and investors randomly allocated IPOs with negative returns are significantly less likely to apply for future IPOs. These results are unlikely to be driven by wealth effects, as they appear even for accounts where the allocation of a given IPO is a small proportion of the total account value. Given the random assignment, it is difficult to explain these results without some appeal to reinforcement learning.

We next test whether investors' randomized IPO return experiences cause substantially different trading decisions in their *non-IPO* portfolios. We view these as our most interesting

[3] See, for example, https://www.tdameritrade.com/investment-products/stocks/IPOs.page.

analyses because testing for experience effects beyond IPO subscriptions takes the greatest advantage of the experiment that we study.

To be more precise, when using non-experimental variation in experiences, we might naturally become more concerned about unobservable investor or time characteristics as we try to explain behavior that is further removed from the original experience. For example, if we find that IPO investors who had positive experiences in this setting are more likely to subscribe to future IPOs, it seems plausible that learning from personal experiences is the main driver of this result, rather than unobservable investor or time heterogeneity. However, if we find using non-experimental data that successful IPO investors are more likely to increase their future trading volume across all stocks, we would quite naturally be more concerned about whether our inferences are contaminated by unobserved investor or time heterogeneity. Even if it is true that investor experience in the IPO market greatly influences a broad variety of investment behaviors, identification remains a challenge as it is ultimately very difficult to control for all of the factors that might jointly determine IPO experiences and trading behavior. The random assignment of experiences in our design allows us to precisely identify experience effects on a wide range of investor decisions.

We uncover a new set of facts regarding the causal effect of investor experiences on investor portfolios as a whole. We find that the exogenous shock of receiving a gain in an IPO security strongly increases treated investors' propensity to trade stocks, exacerbates the disposition effect, which is the tendency to sell stocks in the portfolio which have experienced gains rather than losses, and causes small but precisely estimated increases in the fraction of the investor's portfolio that is invested in the industry sector of the treatment IPO stock. We also find that there is a small but significant increase in the number of stocks held by investors experiencing IPO lottery wins.[4]

[4]All these results are estimated removing the direct allocation of the IPO stock that treatment accounts have because they were "winners" of the lottery.

Treatment effects are not homogenous across investors or experiences. The effects of experience appear to be stronger for smaller accounts, suggesting that more financially sophisticated, larger investors are less susceptible to these effects. Treatment effects are generally positive when the experience (measured by the first-day IPO return) is positive, and generally estimated to be negative when first-day returns are negative. This positive correlation between the experience shocks and the estimated treatment effects is statistically significant across all outcome variables estimated across all 57 IPOs that we study.

We view our results as important for the development of both behavioral and rational theories of investor behavior. Many previous behavioral models assume that investors narrowly frame stocks separately when evaluating performance, and in this sense ignore the potential for cross-security effects within investor portfolios (see, for example, Barberis et al., 2006). For example, current models of "realization utility," the idea that investors receive utility jolts at the time of selling an investment, generally assume that utility is defined at the asset level and generally ignore the possibility that there may be cross-asset realization utility effects (see Barberis and Xiong, 2012; Frydman et al., 2014). However, it seems plausible that realizing a gain in one stock might make an investor more willing to realize a loss in another because utility jolts are bracketed together. Our findings suggest that experiences arising from one stock in a portfolio has a *causal* effect on decisions regarding other securities, or put differently, we find that there can be contagion effects even *within* an investor's portfolio.

On the other hand, it is difficult to square our results with fully rational theories of economic decision-making, as it is difficult to explain them using mechanisms such as wealth effects or rational portfolio rebalancing. Much like the related literature in this field that uses non-experimental variation for identification, we find strong evidence consistent with reinforcement learning behavior by investors in financial markets.

Finally, our results are also related to the recent literature (see, for example, Parker et al., 2011; Agarwal et al., 2007; Bertrand and Morse, Bertrand and Morse) which uses micro-data to study the consumption response to unanticipated income shocks. Most of these studies harness the power of experimental or quasi-experimental variation to reject the predictions of the rational expectations life-cycle/permanent income hypothesis. Our results are similar in approach, and complement this literature by showing that there are good reasons to believe that shifts in beliefs and preferences caused by exogenous variation in gains and losses have effects on investment (and not just consumption) behavior.

The next section describes the natural experiment that we study, describing the details of the Indian IPO lottery process. Section (3) describes the data that we employ, Section (4) describes how we estimate treatment effects on a range of investment behaviors using these lotteries, Section (5) describes the results, and finally, Section (6) concludes.

2 The Experiment: India's IPO Lotteries

2.1 Details of Regulation and the IPO Process

As with many other details of regulation in the country, the Indian regulatory process for IPOs is quite complex. Several papers (e.g., Anagol and Kim, 2012; Campbell et al., 2015), have used this complexity of the Indian regulatory process to cleanly identify a range of economic phenomena.

Our experiment uses the Indian retail investor IPO lottery as an identification mechanism. This lottery arises in situations in which an IPO is oversubscribed, and the use of a proportional allocation rule to allocate shares would violate the minimum lot size set by the firm. In such cases, the lottery is run to give investors their proportional allocation in expectation. In this lottery, some investors will receive the minimum lot size and others will receive zero shares.

The fundamental reason for the lottery is that in India, regulations require that a firm must set aside a maximum of 50% (and, more importantly, a minimum of 35%) of its shares to be available for allocation to retail investors at the time of IPO.[5] For the purposes of the regulation, "retail investors" are defined by a numerical cutoff, as those seeking to purchase shares below a particular regulatory ceiling for a given issue. At the time of writing, this regulatory ceiling is Rs. 200,000 (roughly US $3400), but this has varied over time.[6]

The allocation process of shares begins with the lead investment bank, which sets an indicative range of prices. The upper bound of this range (the "ceiling price") cannot be more than 20% higher than the lower bound (or "floor price"). Importantly, a minimum number of shares (the "minimum lot size") that can be purchased at IPO is also determined at this time. All IPO allocations are constrained to be integer multiples of this minimum lot size.

Retail investors can submit two types of bids for IPO shares. The simplest type of bid is a "cutoff" bid, where the retail investor commits to purchasing a stated multiple of the minimum lot size at the final issue price that the firm chooses within the price band. To submit a cutoff bid, the retail investor must deposit an amount into an escrow account, which is equal to the ceiling price of the price band multiplied by the number of shares bid for. If the investor is allotted shares in the case in which the issue price is less than the ceiling price, the difference between the deposited and required amounts is returned to the investor.

[5]The Securities Exchange Board of India details the process of allotment in section 7.6.1.1 and 11.3.5 of its Disclosure and Investor Protection (DIP) Guidelines until 2009, and Chapter II of the Issue of Capital and Disclosure Requirements (ICDR) regulations since 2009. They can be accessed at `http://www.sebi.gov.in/guide/sebiidcrreg.pdf` and `http://www.sebi.gov.in/guide/DipGuidelines2009.pdf`. This 35% minimum amount was increased from 25% on 4 April 2005. Over our sample period, however, this regulation does not change.

[6]The Indian regulator, SEBI, introduced the definition of a retail investor on August 14, 2003 and capped the amount that retail investors could invest at 50,000 rupees per brokerage account per IPO. This limit was increase to 100,000 rupees on March 29, 2005, and again increased to 200,000 rupees on November 12, 2010. Note that this regulatory definition technically permits institutions to be classified as retail when investing small amounts, but over our sample period, independent account classifications from the depositories reveal that this accounts for a miniscule proportion of retail investment in IPOs. We remove these aberrations from our analysis, and discuss this in the internet appendix to the paper.

It is also possible for retail investors to submit a full demand schedule, i.e., the number of lots that they would like to purchase at each possible price within the indicative range. The investor must also deposit the maximum monetary amount consistent with their demand schedule at the time of submitting their bid. In this case, if allotted shares, the investor's order will be filled at the stated demand associated with the issue price.

Once all bids have been submitted, the firm and investors jointly determine the level of retail (and total) investor oversubscription. The two inputs to this are total retail demand, and the firm's chosen supply to retail investors. Firm supply is restricted by the overall number of shares that the firm decides to issue (this is fixed prior to the commencement of the application process for the IPO). It is also restricted by the mandated lower and upper bounds of 35% and 50% for the retail investor fraction in IPO allotment, in which range the firm chooses a point. Once this choice is made, retail oversubscription is the ratio of total retail demand to firm supply, i.e., total shares made available to retail investors.

There are then three possible cases:

A. If retail oversubscription at the ceiling price is less than or equal to one, then all retail investors are allotted shares according to their demand schedules.

B. If retail oversubscription at the ceiling price is greater than one, shares are allocated to investors in proportion to their stated demands. There is no lottery involved in this case.

C. If retail oversubscription is *far* greater than one (the issue is substantially oversubscribed), then the situation may well arise that a number of investors would, under a proportional allocation scheme, receive an allocation which is lower than the minimum lot size. This is not permitted, and therefore such investors are entered into a lottery. In this lottery, the probability of receiving the minimum lot size is proportional to the number of shares in the original bid.

It is this third case, in which the lottery takes place, that constitutes our experiment.

It is worth noting that far from being an unusual case, this lottery scenario affects roughly 1.7 million Indian investors over the 2007 to 2012 period that we study. Note that the minimum allocation (minimum lot size times issue price) determines the maximum stake in our experiment along with the listing return, i.e., the difference between the price at listing and the issue price. The minimum allocation is the base on which gains and losses for the treatment group can be accrued, relative to the control group.

We now provide a more formal description of the process, and illustrate it with an example from an actual Indian IPO.

2.2 The Probability of Treatment

Let S be the total supply of shares that the firm decides to allocate to retail investors. Let $c = 1, ..., C$ index "share categories," which are integer multiples of the minimum lot size x for which investors can bid. The set of possible amounts of shares for which investors can bid is therefore: x,$2x$,...,Cx.[7] Let a_c be the total number of applications received for share category c. The total demand D for an IPO with C share categories is then:

$$D = \sum_{c=1}^{C} cxa_c. \tag{1}$$

Retail oversubscription v is then defined as:

$$v = \frac{D}{S}. \tag{2}$$

As described in case (2.1) above, if $v \leqslant 1$ at the ceiling price, then all investors get the shares for which they applied, and if $v > 1$, one of cases (2.1) or (2.1) will be in force.

[7]All Indian IPOs have a minimum lot size, which is also the mandatory lot size increment.

In the latter two cases, the first step is to compute the allocations for each share category under a proportional allocation rule, and compare these allocations to the minimum lot size x.

Let $J \leq C$ be the share category such that share categories $c \in [J, ..., C]$ receive proportional allocations which are greater than x, and share categories $c' \in [1, ..., J)$ receive proportional allocations which are less than x. If $J = 1$ then we are in case (2.1), otherwise we are in case (2.1).

In either case, investors in share categories $c \geqslant J$ receive a proportional allotment $\frac{cx}{v}$, and a total number of shares equalling $\sum_{c=J}^{C} \frac{cxa_c}{v}$. However, investors in share categories $c' \in [1, ..., J)$ cannot receive the minimum of x shares (since J is the cutoff share category, i.e., $\frac{(J-1)x}{v} < x$). Let Z be the remainder of shares to be allotted, i.e.,

$$Z = S - \sum_{c=J}^{C} \frac{cxa_c}{v}. \tag{3}$$

These are the shares allocated by lottery in case (2.1). Note that in this lottery, the possible outcomes are winning the minimum lot size x with probability p_c, or winning nothing with probability $1 - p_c$.

By regulation, the probability of winning in share categories $c' \in [1, ..., J)$ must be exactly proportional to the number of shares applied for, meaning that in expectation, investors will receive their proportional allocation. That is, for share categories $c' \in [1, ..., J)$:

$$\frac{p_{c'}}{p_{c'-1}} = \frac{c'x}{(c'-1)x} = \frac{c'}{c'-1}. \tag{4}$$

The combination of equation (4) and the fact that the total remaining shares are described by equation (3) gives us:

$$\sum_{c'=1}^{J-1} (p_{c'})xa_{c'} + \sum_{c'=1}^{J-1} (1 - p_{c'}) = Z. \tag{5}$$

Solving (5), we get that $p_{c'} = \frac{c'}{v}$ of winning exactly x shares in share categories $c' \in$

$[1, ..., J)$. We show the solution in the appendix to the paper.

In general, the probability of winning increases proportionally with the number of shares bid for c, and decreases with the overall level of over-subscription v. Note that this implies that the probability of winning will vary across share categories within IPOs, as well as across IPOs. However, conditional on two investors applying for the same share category in the same IPO, the investor chosen to actually receive the shares will be random. This is the major source of variation we exploit in estimating experience effects in portfolio decisions.

2.3 An Example: Barak Valley Cements IPO Allocation Process

We now provide an example to illustrate this process. Barak Valley Cements' IPO opened for subscription on October 29, 2007, and remained open for subscription through November 1, 2007. The stock was simultaneously listed on the National Stock Exchange (NSE) and the Bombay Stock Exchange (BSE) on November 23, 2007. The listing price of the stock was Rs. 42 per share, and the stock closed on the first day of listing at Rs. 56.05 per share, for a 33.45% listing day gain.

The retail oversubscription rate v for this issue was 37.62. Given this high v, all investors that applied for this IPO were entered into a lottery, i.e., $J = C$.

Table 1 shows the official retail investor IPO allocation data for Barak Valley Cements.[8] Each row of column (0) of the table shows the share category c, associated with a number of shares bid for given in column (1), which, given the minimum lot size $x = 150$ for this offer is just cx. In this case, $C = 15$, meaning that the maximum retail bid is for 2,250 shares. This is because $C = 16$ would give a number of 2,400 shares, and a maximum subscription amount of Rs. 100,800 at the listing price of Rs. 42, which violates the maximum retail

[8]These data are obtained from `http://www.chittorgarh.com/ipo/ipo_boa.asp?a=134`

investor application constraint of Rs. 100,000 rupees per IPO.[9] Column (2) of the table shows the total number of retail investor applications received for each share category, and column (3) is simply the product of columns (1) and (2).

Column (4) shows the investor allocation under a proportional allocation rule, i.e., $\frac{cx}{v}$. As $v = 37.62$, this proportional allocation is less than the firm's minimum lot size of 150 shares per investor for all share categories, i.e., $J = C$. By regulation, the firm is now required to conduct a lottery to decide share allocations.

Column (5) shows the probability of winning the lottery for each share category c, which is $p = \frac{c}{v}$. For example, 2.7% of investors that applied for the minimum lot size of 150 shares will receive this allocation (this is the treatment group in this share category), and the remaining 97.3% of investors applying in this share category (the control group) will receive no shares. In contrast, 40.6% of investors in share category $c = 15$ receive the minimum lot size $x = 150$ shares. For this particular IPO, *all* retail investors are entered into the lottery, and will ultimately receive either zero or 150 shares of the IPO.

Column (6) shows the total number of shares ultimately allotted to investors in each share category, which is the product of x, column (5) and column (2). Columns (7) and (8) show the total sizes of the treatment and control groups (number of retail investors) in each share category for the Barak Valley Cements IPO lottery. Across all share categories, 12,953 investors are treated, and 55,669 are in the control group.

It is perhaps easiest to think of our data as comprising a large number of experiments, where each experiment is a share category within an IPO. *Within* each experiment the probability of treatment is the same for all applicants, and we exploit this source of randomness, combining all of these experiments together to estimate the causal effect of experiencing the IPO listing return on future investment behavior.

[9]In practice, each brokerage account was counted as an individual retail investor, meaning that a single investor could in practice exceed this threshold by subscribing using multiple different brokerage accounts. We are able to capture this behavior in our data, however, as our data are aggregated by the anonymized tax identification number of the investor.

3 Data

To understand the causal effects of experience on investment behavior in this setting, we require two major sources of data. First, we need data on the full set of investors who applied for each IPO, i.e., both successful and unsuccessful applicants. These data are used to define our treatment and control groups. Second, we require investor-level data on portfolio allocations and trades to measure how investing behavior changes in response to the treatment, i.e., the experience in the IPO lottery.

3.1 Data on IPO Applications

When an individual investor applies to receive shares in an Indian IPO their application is routed through a registrar, who, in consultation with one of the stock exchanges, performs the randomization to determine which investors are allocated. We obtain data on the full set of applicants to 57 Indian IPOs over the period from 2007 to 2012 from one of India's largest share registrars. This registrar handled the largest number of IPOs by any one firm in India since 2006, covering roughly a quarter of all IPOs between 2002 and 2012, and roughly a third of the IPO market by number over our sample period.

For each IPO in our sample, we observe whether or not the applicant was allocated shares, the share category c in which they applied, the geographic location of the applicant by pin-code,[10] the type of bid placed by the applicant (cutoff bid or full demand schedule), the share depository in which the applicant has an account (more on this below), whether the applicant was an employee of the firm, and other application characteristics such as whether the application was supported by a blocked amount at a bank.[11]

[10]PIN codes in India are postal codes managed and administered by the Indian Postal Service department of the Government of India. They are similar to postcodes in the UK, although cover a larger region in India.

[11]An application supported by blocked amount (ASBA) investor is one who has agreed to block the application money in a bank account which will be refunded should she not be allocated the shares in an IPO.

3.2 Data on IPO Applicants' Equity Portfolios

Our second major data source allows us to characterize the equity investing behavior of these IPO applicants. We obtain these data from a broader sample of information on investor equity portfolios from Central Depository Services Limited (CDSL). Alongside the other major depository, National Securities Depositories Limited (NSDL), CDSL facilitates the regulatory requirement that settlement of all listed shares traded in the stock market must occur in electronic form. CDSL has a significant market share – in terms of total assets tracked, roughly 20%, and in terms of the number of accounts, roughly 40%, with the remainder in NSDL. While we do also have access to the NSDL data (these data are used extensively and carefully described in Campbell et al., 2014), we are only able to link the CDSL data with the IPO allocation information, as we describe below.

The sensitive nature of these data mean that there are certain limitations on the demographic information provided to us. While we are able to identify monthly stock holdings and transactions records at the account level in all equity securities in CDSL, we have sparse demographic information on the account holders. The information we do have includes the pincode in which the investor is located, and the type of investor. We use investor type to classify accounts as beneficial owners, domestic financial institutions, domestic non-financial institutions, foreign institutions, foreign nationals, government, and retail accounts. This paper studies only the category of retail accounts.

As described in Campbell et al. (2014), the share of direct household equity ownership in India in total equity investment is very large (roughly 80%-95%), relative to the share of indirect equity holdings using mutual funds, unit trusts, and unit-linked insurance plans. This means that we observe roughly the entire equity portfolio of the household in our analysis, allowing us to interpret the treatment effects of experience that we estimate as effects on household portfolio choice. This distinguishes our study of investment behavior from those attempting to detect effects of experienced returns on trading behavior, such as (Seru et al., 2010; Strahilevitz et al., 2011).

3.3 Constructing the Final Sample

In order to match the application data to the CDSL data on household equity portfolio choice, we obtain a mapping table between the anonymous identification numbers of household accounts from both data sources. We verify the accuracy of the match by checking common geographic information fields provided by both data providers such as state and pincode.

Every applicant for an IPO must register an account number with either of the two depositories. The option to receive allocated shares in an IPO in physical form does not exist. Therefore, we observe all allocations made to investors in IPOs after the selection process managed by share registry firms in the CDSL data. We also observe accounts that applied for an IPO, but due to randomized allocation did not get allocated any share in an IPO, thus observing the universe of counterfactuals in the IPO randomized lottery. We are also able to observe allotments (but not applications) to particular household accounts from CDSL, which we use in some of our analysis below.

All CDSL trading accounts are associated with a tax related permanent account number (PAN), and regulation requires that an investor with a given PAN number can only apply once for any given IPO.[12] Consistent with this, we observe that there are no two trading accounts in any single IPO that are associated with the same (anonymized) PAN number. Thus no investor account may simultaneously belong to both the control and treatment group, or be allocated twice in the same IPO. However, it is possible that a household with multiple members with different PAN numbers could submit multiple applications for a given IPO; we discuss how this possibility changes the interpretation of our results as we present them.

[12] In July 2007 it became mandatory that all applicants provide their PAN information in IPO applications. SEBI circular No.MRD/DoP/Cir-05/2007 came into force on April 27, 2007. Accessed at http://goo.gl/OB61M2 on 19 Sep 2014.

3.4 Summary Statistics

Between March 2007 and March 2012, the common sample period for our total dataset, we observe 85 IPOs (of a total of roughly 246), 57 of which have at least one share category with randomized allocation, which is roughly 70% of the sample. This compares to 73.3% of all 246 IPOs over the period which were over-subscribed. Figure 1 shows the coverage of IPOs in our sample relative to that in the total universe of IPOs. Our sample coverage closely tracks aggregate IPO waves, with a severe decline in 2009, and high numbers of IPOs in 2008 and 2010.

Table 2 presents summary statistics on the 57 IPOs in our sample. The table shows that these IPOs account for 23% of all IPOs over this period by number, and US$ 2.65 BN or roughly 9% of total IPO value over the period.

Between 32% and 35% percent of shares in these IPOs are allocated to retail investors.[13] The average IPO in our sample was 11.5 times oversubscribed, leading to an average of 10,622 treatment accounts and 19,192 control accounts per IPO. The majority of IPOs are in the manufacturing sector, with fourteen in other services and five in each of technology and retail respectively.

Figure 2 plots the mean and distribution of first-day returns for our 57 IPOs across the five years of our sample. The figure shows that our sample contains significant dispersion in experiences, with IPOs generating both high negative ($< -50\%$) and high positive returns ($> 150\%$) and a range in-between. The second panel shows the first day variability of the IPO stocks in our sample, measured by the first day high price minus the first day low price divided by the issue price. Our IPO stocks also show large dispersion in first day return volatility, with intra-day dispersion of 50% not uncommon.

[13] This is slightly below the mandatory 35 percent allocation to retail investors because we do not include employees in this calculation as employees are not randomly assigned shares.

Table 3 Panel A shows summary statistics on the investor accounts in our sample. For each IPO applicant in our sample, we calculate these statistics in the month prior to the IPO for which the account applied.[14]

The table shows that our sample primarily comes from IPOs in 2007 and 2010, with 2011 and 2008 also contributing to the sample, reflecting the IPO waves occurring during the period documented in Figure 1. Of the total of 1.7 million IPO applicants, approximately 370,000 were new (rookie) accounts at the time of applying for the IPO. This is consistent with previous work suggesting that participating in IPOs is a common way for retail investors to begin participating in the stock market – which also suggests that our analysis allows us to uncover insights into the much-studied question of equity market participation (see, for example, Mankiw and Zeldes, 1991; Haliassos and Bertaut, 1995; Blume and Zeldes, 1994; Vissing-Jorgensen, 2002).

The pooled mean equity portfolio size in the sample is US$ 8,400, with a median value of roughly US$ 2,200, suggesting the presence of a few very large accounts in the set of applicants. We also break accounts down by their age since account opening (in months). Our sample contains a reasonably large distribution of account ages, with roughly a quarter overall being less than five months old, but close to half greater than two years old. This pattern varies interestingly across years, with rookie and young investors more prevalent in the "boom" years of 2007, 2008, and 2011, and seasoned investors more prevalent in IPO participation in the relatively low-aggregate equity market performance years of 2009 and 2010.

In terms of geographic distribution, the accounts are focused on areas with greater economic activity overall, including the major Indian states of Gujarat, Maharashtra, Rajasthan, as well as the capital state Delhi.

[14] An account that applied to two separate IPOs would therefore appear in these data twice.

Table 3, Panel B shows summary statistics for various characteristics of household portfolios that we use as outcome variables for estimating treatment effects later in the paper. All of these measures are calculated based on the full sample of treatment and control accounts for a total of 13 months, 6 before and after the IPO, and the month of the IPO.

The first row of this panel of the table shows that on average, across all 13 months, the accounts in our data applied for, or were allotted shares in, 0.31 IPOs.

The second row of the table shows summary statistics of our monthly measure of trading activity, which we simply measure as the number of transactions (purchases plus sales) in the account. This averages 4.3 transactions per account, varying from a low of 2.9 to a high of 6.7 across years, with considerable cross-sectional variability within each year. The median is considerably lower, and zero in some years, because of the substantial number of rookie accounts in the dataset.

The third row of the table shows our measure of diversification, which is simply measured as the number of securities held in the account. In our sample this measure averages 11.3 securities held per account, with a median of 5 securities, suggesting a right-skewed distribution.

The fourth row of the table shows our measure of the disposition effect, which is percentage paper gains realized minus percentage paper losses realized. This difference averages 6.67% in our full sample. This magnitude is comparable to, but slightly larger than that documented in studies of the US market.

Finally, we compute the investor's portfolio tilt towards the sector of the IPO firm, and find that on average, the retail investors in our sample held 8.7% of their portfolio in the IPO stock's sector across all event-months in our sample.

The next section describes our methodology to combine all of the experiments that we observe from the 57 IPOs.

4 Methodology

4.1 Robustness of the Randomized Allocation

Our first step is to perform statistical tests of the robustness of the randomization process in our sample of 57 IPOs.

For each IPO share category that underwent randomized allocation in our sample, we estimate the following cross-sectional regression:

$$y_{i,j,c} = \alpha_{j,c} + \beta_{j,c} I_{\{success_{i,j,c}=1\}} + u_{i,j,c}. \tag{6}$$

Here, $y_{i,j,c}$ for investor i in IPO j, share category c is successively, a set of characteristics of the IPO applicant such as whether the applicant was a cut-off bidder, the state in India in which the applicant is located, the type of payment mode selected by the applicant, and whether the applicant chose CDSL or NSDL as their depository.

$I_{\{success_{i,j,c}=1\}}$ is a treatment dummy, which takes the value of one if the applicant is allotted shares in the lottery and zero otherwise. $\beta_{j,c}$ is therefore the estimated difference between the means of the characteristic $y_{i,j,c}$ between the treatment and control groups within each IPO share category. Under the null hypothesis that the randomization is truly robust, we expect $\beta_{j,c}$ to be statistically insignificant.

We conduct these tests across a large number of randomized IPO share categories for a range of applicant characteristics. We also expect, therefore, (given a sufficiently large set of tests) that a standard normal distribution will accurately describe the distribution of the t-statistics of the $\beta_{j,c}$ coefficients under the null hypothesis of no significant differences in the characteristics of treatment and control accounts. If the null is rejected, that is, if certain applicant characteristics are systematically associated with wins in the IPO allocation lottery, then we would expect to find more t-statistics in the tails of the distribution.

We conduct a substantial number of randomization checks, and simply present figure 3 to represent the results. The figure shows the distribution of the $\beta_{j,c}$ t-statistics alongside a standard normal distribution, and shows that the standard normal appears well able to characterize the empirically estimated distribution. This provides reassurance that the randomization is truly robust across all of the lottery experiments that we observe.

4.2 Estimating Treatment Effects

We can view each randomized share category in each IPO as a separate experiment with a different probability of being allotted shares. The idea of our empirical specification is to pool all of these experiments in order to maximize statistical power, while ensuring that we exploit the randomized variation of treatment status within each IPO share category. Our strategy is similar to that employed in Black et al. (2003), who estimate the impact of a worker training program that was randomly assigned within 286 different groups of applicants.

Intuitively, this approach proceeds by simply stacking the different applicants from all of the experiments together into a single dataset, and then including a fixed effect for each experiment. These experiment-level fixed effects ensure that our identification of the treatment effect stems solely from the random variation in treatment within each experiment.[15]

In particular, we estimate the causal effect of the experience of winning an IPO lottery on an outcome variable by estimating the cross-sectional regression in each (event) month t:

$$y_{i,j,c,t} = \alpha + \rho_t I_{\{success_{i,j,c}=1\}} + \gamma_{j,c} + \beta X_{i,j,t} + \epsilon_{i,j,c,t}. \tag{7}$$

Here, $y_{i,j,c,t}$ is successively, an outcome variable of interest (for instance, the number of times the individual i applies for subsequent IPOs) for applicant i in IPO j, share category c, at event month t (we measure time in relation to the date of the lottery). $I_{\{success_{i,j,c}=1\}}$ is an indicator variable that takes the value of 1 if the applicant was successful in the lottery

[15]See Chapter 3 of Angrist and Pischke (2008) for a discussion of how regression with fixed effects for each experimental group identifies the parameter of interest using only the experimental variation.

for IPO j in category c (investor is in the treatment group), and 0 otherwise (investor is in the control group). $\gamma_{j,c}$ are fixed effects associated with each experiment, i.e., each IPO share category in our sample. Angrist et al. (2013) refers to these experiment-level fixed effects as "risk group" fixed effects. $X_{i,j,t}$ are account-level control variables. Conditional on the inclusion of the risk-group fixed effects, variation in treatment is random, meaning that the inclusion of controls should have no effect on our point estimates of ρ_t. Nevertheless, we include these controls to soak up additional variation in the dependent variable to increase the statistical precision of our estimates.

Specification (7) identifies ρ_t as the causal impact of the experience of winning the IPO lottery on the outcome variable $y_{i,j,c,t}$.

To account for the possibility that there may be heterogenous treatment effects across different groups of investors, we also estimate specifications of the form:

$$y_{i,j,c,t} = \alpha + \rho_t^G I_{\{success_{i,j,c}=1\}} I_{\{G(X)\}} + \gamma_{j,c} + \beta X_{i,j,t} + \epsilon_{i,j,c,t}. \tag{8}$$

In the above equation, $I_{\{G(X)\}}$ is an indicator variable for membership of a particular group, where group membership depends on account-level characteristics, for example, portfolio size. We also separately estimate treatment effects across the spectrum of IPO experiences, for example we estimate treatment effects for IPOs with positive first-day returns separately from those for IPOs with negative first-day returns.

A few notes on estimation. First, we cluster all standard errors by calendar-month, to pick up potential correlations of the error terms $\epsilon_{i,j,c,t}$ across all IPOs occurring in the same month, as well as correlations across share categories within IPOs. Second, we estimate all treatment effects for $t \in [-6, ..0, .. + 6]$ where $t = 0$ is the month in which the lottery takes place, with leads and lags of up to 6 months. The +1 to +6 window identifies the causal impact of the experience on future outcomes. Tracking the -1 to -6 outcome variable serves as a convenient "placebo" test in addition to the randomization checks which we conduct above. If the lottery is truly randomized, we should find that receiving treatment at time zero does not, on average, predict outcomes in time periods before treatment was actually assigned. This placebo test is particularly useful because many outcomes are highly serially correlated over time, so we would be likely to pick up any selection bias into treatment (if it exists) by inspecting the behavior of treatment and control groups in the pre-treatment periods. For example, if particular applicants figure out a way to "game" the lottery then we might find that their treatment at time zero actually predicts their behavior in the -1 to -6 window.

We now turn to discussing the results from estimating equation (7) for a range of outcome variables.

5 Results

5.1 Treatment Effects on Future IPO Subscription

We begin by testing how the treatment, i.e., receiving a randomized allocation of IPO stock, affects an investor's propensity to apply for other IPOs in the subsequent six months. This outcome has been studied in previous work, (see, for example, Kaustia and Knüpfer, 2008; Chiang et al., 2011), but always in non-experimental contexts in which randomized variation of the type that we exploit is not available. As a result, this outcome variable is a useful cross-check on whether our empirical approach confirms the results in prior work.

Table 4 presents these results, in which the outcome variable in equation (7) is a dummy variable which captures whether or not the account applied for an IPO in a given month within the event window.

The construction of this outcome variable warrants further explanation. For IPOs where our data provider was the registrar, we can directly measure whether or not an account applied. For IPOs where our data provider was not the registrar, we can observe whether the account was allotted shares since we see allotments for the entire universe of IPOs from the CDSL data. We set the outcome variable to one in either case – if we see an application for IPOs for which our data provider was the registrar, or if we see an allotment for IPOs not covered by our registrar – and zero otherwise.[16]

We follow the format of table 4 in all subsequent tables of results from estimating equation (7) for a range of outcome variables. Each panel of the table shows results for a set of applicants for the window $t \in [-6, ..., 0, ..., +6]$ where $t = 0$ is the month of the lottery. The first row of numbers in each panel shows the coefficients ρ_t, which are the estimated treatment effects from the cross-sectional regressions estimated for each event-time t in the window shown in the column header. The second row of numbers in each panel shows standard errors, and the third row of numbers in each panel, in square parentheses, shows the mean of the outcome variable for the control group, which we use to interpret the magnitudes of the treatment effects.

Panel A of the table shows results for the entire sample of applicants. We find that there is a significant relationship between treatment status in the outcome five and two months prior to treatment, however the signs are in opposite directions and there is no clear pattern amongst the other coefficients. Further, these correlations are not found when we split the sample in different ways suggesting that these are just chance occurrences. Based on these results, as well as similar tests using the other outcome variables we study, we conclude that treated applicants are not systematically different from non-treated applicants after including the risk group fixed effects. Note that by chance some of the pre-period treatment effects will show up as significant, and that given our very large sample sizes we will be able to statistically detect relationships that are not particularly economically meaningful.

Panel A shows that in the month of treatment, accounts that received a randomized allocation are 0.17 percentage points (p.p.) more likely to apply to an IPO. In the month

[16]For the set of IPOs for which we can observe allotments but not applications, our measure is noisy, because although an account had to apply to receive shares, there are also accounts which applied but did not receive shares. We focus on this combined measure because it includes all of the information available to us, but we note that our results likely under-estimate the full impact of IPO experiences on future IPO application behavior.

after treatment, treated accounts are 0.85 p.p. more likely to have applied for an IPO, and this effect is significant at the five percent level. This corresponds to a roughly 2% increase in the probability of applying for an IPO relative to the base rate probability of applying in the control group (43.68%). The effect size in month two is substantial, raising the probability of applying relative to the base rate by 3%. The effect sizes in months three through five are smaller in levels (between 0.15 and 0.27 p.p. when significant), but are similar in magnitude to the effect sizes in the first few post-treatment months relative to the base rate of applying for IPOs (they all represent roughly a 2% increase in the base rate of applying). Cumulatively, simply assuming that these probabilities are independent, we see an increase in the probability of applying to a future IPO of roughly 12% relative to the base rate in the control group (in month zero) over the six months following the IPO.[17] Panel A overall suggests a significant causal effect of exogenous IPO experience on future IPO applications, and are a useful validation of our estimation approach given their qualitative similarity to previous work using non-randomized allocation of IPOs.

The remaining panels of the table consider the possibility that there may be heterogenous treatment effects, and allow us to dig deeper into the economic sources of the experience effects that we estimate.

One possible component of the experience effects that we estimate stems from wealth effects. In particular, the channel here would be that treatment investors make money on the IPO they randomly receive, and feel less wealth-constrained, thus applying in greater numbers to future IPOs. One simple way to assess the importance of wealth effects to our results is to separately estimate effect sizes based on the size of the portfolio in the month prior to treatment. If we find meaningful effect sizes for large accounts for which the IPO allocation would represent a smaller fraction of wealth, then it is unlikely that our results are primarily due to wealth effects. Note that in general the wealth gains associated with

[17]As mentioned earlier, these are likely under-estimates of the true effect as we only observe allotments and not applications for IPOs that were not handled by our data provider.

the IPO allotment will be small because lottery winners only receive the minimum allotment of shares, and the net gain to the winners is only the amount of appreciation in the IPO stock from the time that the shares are allocated (we know that both treatment and control investors have the wealth associated with the value of the IPO at issuance because they are required to deposit more than this as part of the application.)[18]

Panels B and C of table 4 split the sample by the portfolio size of the applicant. For each IPO share category (i.e., experiment) we split the sample based on the median portfolio size (our measure of wealth) of the applicant, and then estimate equation 7 separately for each of these samples (above (below) median in Panel B (C)). This change means that the set of IPOs, and therefore IPO features such as first day returns, timing, and so on are the same in both Panels B and C. In the month of treatment we find a statistically significant treatment effect only for the below median wealth accounts. However, we find significant effects for both above and below median wealth accounts in month one. The effects are statistically significant in the below median portfolio size group until month six, with the effects in the below median wealth accounts approximately double the size of the effect in the above median wealth accounts. While the effect of treatment continues to be positive for the above median wealth group, they do become small and statistically insignificant in the later periods (months +4 and +6).

In Panels D and E we split the sample into IPOs that had positive and negative first-day returns, respectively. Note that there are many more IPOs (and therefore accounts) in the group of first-day positive return IPOs.

[18] For example suppose an IPO had an issue price of 10 rupees and the minimum lot size was 10 shares. Both lottery winners and losers had to have deposited 100 rupees (or more) to enter the lottery. The only difference is that the winners get this back in shares and the losers get this back in cash. Thus, the only wealth effect here would be if the IPO stock appreciates in value; if it appreciated by 10 percent in the first month then the wealth effect would only be 10 rupees. Note also that technically, control investors could immediately buy the stock in the secondary market. This would limit the wealth gains to the difference between the initial listing price and the issue price of the IPO.

We find that treatment effects on future IPO applications are positive when first-day IPO returns on the treatment IPO are positive, consistent with the idea that positive experiences make investors more likely to seek out similar experiences in the future. Given the fact that most of the sample experienced positive returns, these results are similar to the full sample results in Panel A. We also find that experiencing a negative return in an IPO has negative impacts on future participation, though these effects are quite imprecisely estimated. Interestingly, these negative impacts are quite persistent – even five months after the negative IPO experience treatment accounts are 5.7% less likely to apply to an IPO, relative to a base rate probability of 12.1% in the control group. Figure (4) shows that this positive relationship between the sign and size of the experience measured by first-day returns is positive across all of the outcome variables that we discuss in greater detail below.

5.2 Treatment Effects on Trading Activity

We now move to testing whether the experience of the IPO lottery allocation has an impact on the investor's portfolio outside the narrow sphere of the IPO market. We view these as our most interesting tests, because they allow us to explore to what extent experiences in particular stocks spillover to other parts of an investor's portfolio.

While in this draft of the paper, we do not have a formal model of portfolio choice, we view these results as meaningful for questions of how experiences affect economic agents' beliefs and preferences, and ultimately, their decision-making. In the specific domain which we consider, namely, retail investor portfolio choice, our results help to shed light on whether investors are better modeled as making separate stock-by-stock decisions (i.e. they "narrowly bracket" their utility changes from the IPO allocation in the sense of Rabin and Weizsäcker (2009) from those experienced on other components of their portfolio), or whether there are within-portfolio utility spillovers. When we find the latter, we go further when analyzing heterogenous treatment effects in an attempt to understand how these within-portfolio spillovers manifest themselves.

We begin by testing whether the treatment makes investors more likely to trade stocks *other than the IPO stock*. A large literature has found a strong correlation between trading volume and returns across stock, bond, and housing markets. Leading theories for this phenomenon include loss aversion (Genesove et al., 2001), investor over-confidence (Statman et al., 2006), and down payment constraints (in housing markets - see Stein, 1995). While some progress has been made in empirically testing these theories, identification is an important challenge because rising markets are potentially different from flat or falling markets in many ways (for example, margin constraints may be looser in rising markets). Our experiment allows us to focus more precisely on the more behavioral channel of how experiencing exogenous gains affect investors' propensities to trade. While we do not argue that this evidence alone can identify which theory of the relationship between trading volume and asset returns is correct, we believe that it is potentially very useful to know whether there is a causal relationship between this kind of short-term experience and trading volume.

Investigating the effects of treatment on trading behavior is also interesting in light of feedback models of asset prices. Most feedback models only consider price feedback; i.e., price increases attract certain types of investors to purchase assets, leading in turn to price impact and additional increases in prices which complete the feedback loop, (see, for example, Shiller, 2015; Barberis et al., 1998; De Long et al., 1990). These models are often based on the assumption that investors have extrapolative expectations. Testing for the presence of such expectations using price and investment flow data is difficult because in most models, prices, and investment decisions are jointly determined in equilibrium.[19] Having the ability to utilize exogenous variation in gains and losses in the portfolio confers a significant advantage in this setting.

[19]Note that this mechanism is not mutually exclusive to the others mentioned above; for example it is possible that positive experiences make investors overconfident, which then leads to greater trading volume as in Statman et al. (2006)

In table 5, the dependent variable is the logarithm of the total number of purchase and sale transactions plus one (to account for zero transaction accounts). In Panel A of the table, we report results both for the dependent variable measured including the treatment IPO stock (the "with IPO security" results) as well as a version of the dependent variable measured using all stocks other than the IPO stock for which the investor applied. In the remaining panels we concentrate on the latter measure, not including the treatment IPO stock.

When we include the IPO stock we see that the amount of trading activity increases substantially in month zero – treated investor make roughly 47% more trades than the control group. This makes sense from a simple portfolio re-balancing perspective – many investors sell the stock immediately. These effects slowly decline as treated investors sell their allocation in the months following treatment.

The more interesting measure does not include the IPO stock. For this measure, we find that the number of transactions increases by approximately 1.7% in the month after receiving the IPO, and remains high and statistically significant through six months after the treatment IPO. Relative to the amount of trading in month zero, this cumulates to an approximate six percent increase in trading over the six months after allotment. This result has a number of interesting implications for models of trading and liquidity, since it says that exogenous variation in gains and losses (for example, those engendered by cash-flow relevant news releases) are associated with changes in investors' propensity to trade.

The remaining panels of this table show that this effect is approximately twice as large for smaller accounts (those below the median portfolio size), but still significant and economically meaningful for the accounts with above-median portfolio size. The results also show that trading activity increases for treatments involving IPOs with positive first-day returns, and (less statistically significantly) decreases for those with negative first day returns.

5.3 Treatment Effects on the Disposition Effect

While a large empirical literature documents the disposition effect across a wide variety of contexts, there is little empirical work testing how the disposition effect responds to exogenous variation in investor experiences. Empirical evidence on how the disposition effect responds to investor experiences is useful for a variety of reasons, not least because it allows us to separate different potential causes for this effect, including loss-averse preferences (see, for example, Barberis and Xiong, 2009), or an irrational belief in mean-reversion (Weber and Camerer, 1998). For example, if the disposition effect is driven by investors' irrational belief in mean reversion, we should see no difference in the disposition effect across our treatment and control investors, because in terms of information sets, these groups should be exactly the same; both chose to apply for the IPO in question, but one was simply lucky to have been allotted. It seems implausible that the experience of receiving one IPO would cause an investor to start believing more (or less) in mean reversion.

We define the disposition effect as the percent of paper gains in the portfolio realized during the month minus the percent of paper losses in the portfolio realized during the month. For example, suppose an account had 4 stocks on paper with gains, and 5 stocks on paper with losses at the beginning of the month. Further suppose that the account sold 1 stock of both gains and losses respectively. Then, our disposition effect measure would be 5%, i.e., 25% of gains realized minus 20% of losses realized.

Table 6 presents our results on the treatment effects on the disposition effect. As before, only the top panel of the table shows results using measures computed including the IPO that was randomly allocated, with the remaining results generated from all of the other stocks in the investor's portfolio.

The table shows that in the month following the IPO, there is a 0.77 p.p. increase in treated investors' disposition effect relative to a base rate of 10% in the control group. In other words, there is roughly an 8% increase in the disposition effect across a treated investor's remaining portfolio due to random allocation of the IPO security, and treated investors behave as if they were more loss averse following the positive realization. One interpretation of this finding is that gains have the effect of shifting investors' utility "reference-points" (see Tversky and Kahneman, 1991) up across the board for all stocks. This finding echoes that of Campbell et al. (2014), who find that overall account outperformance relative to the market is associated with increases in the disposition effect, using a different (non-experimental) approach.

The table also shows that there are heterogenous treatment effects on the disposition effect. Treatment causes increases in the disposition effect for both above and below-median wealth groups in the first month after the allocation, but persists beyond the first month only for the below-median wealth group. We also find that the effects are stronger for those IPOs that had positive first day returns, while experiencing negative returns does not appear to have a consistent impact on the disposition effect over the six months after allocation.

5.4 Treatment Effects on Familiarity in Portfolio Choice

A large literature documents that investors demonstrate a preference for familiarity, i.e., they tend to invest in firms that are located physically close to them, or those that have some relationship with the investor's occupation (see, for example, Coval and Moskowitz (2001)). One potential explanation for this familiarity effect is that investors believe they have private information about stocks that they are familiar with (although whether they actually out-perform in those stocks is unclear − see Massa and Simonov (2006); Seasholes and Zhu (2010). A simpler way for investors to become familiar with a sector is to simply own a stock in that sector. Consistent with this, Huang (2012) finds using data from a large discount broker in the U.S. over the period 1991 - 1996, that individuals are more likely to buy a stock in an industry in which they previously experienced a gain. Our design allows us to test this idea using exogenous variation in sectoral experience that is unlikely to be conflated with other investor or time-varying characteristics.

We therefore test whether treated investors are more likely to invest in the sector of the randomly allocated IPO lottery stock. The outcome variable here is the percentage of the portfolio invested in stocks in the same industry sector as that of the IPO lottery stock.[20]

Panel A of Table 7 first checks the mechanical increase in the outcome variable when it is constructed using the actual IPO stock that was allocated to the investor. We then move to analyzing the treatment effects on portfolio choice for all stocks other than the IPO stock, and find that there is a small but statistically significant increase in the fraction of the portfolio invested in the sector of the IPO stock. This effect is most prominent in months three through six following the IPO, and corresponds to a 5 to 7 basis point increase in the fraction of the portfolio in the sector. As a percentage of the base rate, which is the control group average allocation to the corresponding sector of approximately 8%, this corresponds to a 1% increase in the fraction of the portfolio allocated to this sector for the treatment group relative to the control group. These effects do appear to be quite persistent despite being small in magnitude.

Turning to the heterogenous treatment effects, we find once again that these results are positive (negative) for IPOs that experienced positive (negative) first-day returns. We also find that the results appear to be larger in magnitude, though not as precisely estimated, for investors with below-median portfolio sizes.

Taken together, these results lend credence to models that assume that investors extrapolate their experiences to their beliefs about other related securities, such as Barberis et al. (2015).

[20]Sectoral allocation is defined by the Indian National Industrial Classification Code (NIC code) as of 2004 for all sectors of the Indian economy. Using the NIC classification, we use the third-level aggregation to define 42 sectors in the economy. The details of this classification is available at http://mospi.nic.in/Mospi_New/upload/nic_alphabetic_5digit2004.html

5.5 Treatment Effects on Diversification

Table 8 reports results on the effect of the randomized IPO allocation on the diversification of the investor's portfolio as a whole. We use a simple working definition of diversification, which is the number of securities held in the portfolio (in particular the logarithm of the number of securities in the portfolio plus one). We find little evidence that our treatment and control groups are unbalanced on this measure of diversification in the months prior to receiving treatment.

The table shows that treated accounts hold approximately 0.64 p.p. more stocks in the month after the IPO allocation, increasing to 0.7 p.p. more stocks two months after the allocation, decreasing to approximately 0.5 p.p. more stocks six months after the allocation. These results while signalling a tiny increase in diversification, are nonetheless precisely estimated.

The results are significant in both the above and below median wealth samples, although the point estimates are larger in the above median versus below median sample. These results are larger when we focus on the sample of IPOs that experienced positive first- day returns, and generally negative (although statistically insignificant) for IPOs that experienced negative first day returns.

These results are interesting in light of work on reference-dependent risk attitudes (see Kőszegi and Rabin, 2007). If experienced gains affect attitudes towards risk, causing expectations of future risk to reduce, then buying an additional stock viewed in isolation, i.e., as an additional gamble, is expected utility increasing. Put differently, if you are randomly allocated a loss in the IPO lottery, this might increase your expectation of future risk in stock investing, somewhat perversely causing future gambles to be less aversive in the language of Kőszegi and Rabin (2007).

6 Conclusion

Our paper exploits the randomized allocation of stocks in 57 Indian IPO lotteries to 1.7 million investors between 2007 and 2012, and provides new estimates of the causal effect of investment experiences on future investment behavior. To our knowledge, this is the first paper to estimate the causal effect of return experiences using the randomized allotment of real securities.

We find that investors experiencing exogenous gains in IPO stocks (the treatment) are more likely to apply for future IPOs, increase trading in their portfolios, exhibit a stronger disposition effect, and tilt their portfolios towards the sector of the treatment IPO. We also find that these treatment effects are stronger for smaller accounts, and increase in magnitude with the experience itself, i.e., the IPO first-day return in each case. We view our results as having implications for a wide range of empirical and theoretical work on the effects of experience on economic decision making.

We plan to extend this first version of our paper in a number of ways. Two of these ways are obvious extensions of our work in this draft. First, we intend to investigate whether experience solely has effects on beliefs, or preferences, or whether it affects both aspects of economic decision-making, by more efficiently using the vast range of variation present in our experimental data. Second, we intend to investigate a wider range of outcomes, including the effects of experience on net investment in the equity market, and whether experiences generate a preference for skewness.

References

Agarwal, S., L. Chunlin, and N. S. Souleles (2007). The reaction of consumer spending and debt to tax rebates-evidence from consumer credit data. *Journal of political economy 115*(6), 986–1019.

Anagol, S. and H. H. Kim (2012). The impact of shrouded fees: Evidence from a natural experiment in the Indian mutual funds market. *American Economic Review 102*(1), 576–93.

Andersen, S., T. Hanspal, and K. M. Nielsen (2014). Once bitten, twice shy: Do personal experiences or wealth changes affect risk taking? *Working Paper*.

Angrist, J. D., P. A. Pathak, and C. R. Walters (2013). Explaining charter school effectiveness. *American Economic Journal: Applied Economics 5*(4), 1–27.

Angrist, J. D. and J.-S. Pischke (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Barberis, N., R. Greenwood, L. Jin, and A. Shleifer (2015). X-CAPM: An extrapolative capital asset pricing model. *Journal of Financial Economics 115*(1), 1–24.

Barberis, N., M. Huang, and R. H. Thaler (2006). Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *The American economic review*, 1069–1090.

Barberis, N., A. Shleifer, and R. Vishny (1998). A model of investor sentiment. *Journal of financial economics 49*(3), 307–343.

Barberis, N. and W. Xiong (2009). What drives the disposition effect? An analysis of a long-standing preference-based explanation. *Journal of Finance 64*(2), 751–784.

Barberis, N. and W. Xiong (2012). Realization utility. *Journal of Financial Economics 104*(2), 251–271.

Bertrand, M. and A. Morse. Trickle-down consumption. *National Bureau of Economic Research, Working Paper No. 18883*.

Black, D. A., J. A. Smith, M. C. Berger, and B. J. Noel (2003). Is the threat of reemployment services more effective than the services themselves? evidence from random assignment in the UI system. *American Economic Review*, 1313–1327.

Blume, M. and S. P. Zeldes (1994). Household stockownership patterns and aggregate asset pricing theories. *Working Paper, University of Pennsylvania*.

Camerer, C. and T.-H. Ho (1999). Experience-weighted attraction learning in normal form games. *Econometrica 67*(4), 827–874.

Campbell, J. Y., T. Ramadorai, and B. Ranish (2014, March). Getting better or feeling better? how equity investors respond to investment experience. Working Paper 20000, National Bureau of Economic Research.

Campbell, J. Y., T. Ramadorai, and B. Ranish (2015). The impact of regulation on mortgage risk: Evidence from india. *American Economic Journal: Economic Policy, forthcoming*.

Chiang, Y.-M., D. Hirshleifer, Y. Qian, and A. E. Sherman (2011). Do investors learn from experience? evidence from frequent ipo investors. *Review of Financial Studies*, 151.

Coval, J. D. and T. J. Moskowitz (2001). The geography of investment: Informed trading and asset prices. *The Journal of Political Economy 109*(4), 811–841.

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann (1990). Noise trader risk in financial markets. *Journal of political Economy*, 703–738.

Ellison, G. and D. Fudenberg (1993). Rules of thumb for social learning. *Journal of Political Economy*, 612–643.

Frydman, C., N. Barberis, C. Camerer, P. Bossaerts, and A. Rangel (2014). Using neural data to test a theory of investor behavior: An application to realization utility. *The Journal of Finance 69*(2), 907–946.

Gamble, K. J. and B. Johnson (2014). How prior outcomes affect individual investors' subsequent risk taking. *Journal of Personal Finance 13*(1).

Genesove, D., C. Mayer, et al. (2001). Loss aversion and seller behavior: Evidence from the housing market. *The Quarterly Journal of Economics 116*(4), 1233–1260.

Haliassos, M. and C. C. Bertaut (1995). Why do so few hold stocks? *Economic Journal*, 1110–1129.

Huang, X. (2012). Industry investment experience and stock selection. *Available at SSRN 1786271*.

Kaustia, M. and S. Knüpfer (2008). Do investors overweight personal experience? evidence from ipo subscriptions. *The Journal of Finance 63*(6), 2679–2702.

Knüpfer, S., E. H. Rantapuska, and M. Sarvimäki (2014). Labor market experiences and portfolio choice: Evidence from the finnish great depression. Working Paper 2275930.

Kőszegi, B. and M. Rabin (2007). Reference-dependent risk attitudes. *The American Economic Review*, 1047–1073.

Malmendier, U. and S. Nagel (2009). Learning from inflation experiences. *Unpublished manuscript, UC Berkley*.

Malmendier, U. and S. Nagel (2011). Depression babies: Do macroeconomic experiences affect risk taking?*. *Quarterly Journal of Economics 126*(1).

Mankiw, N. G. and S. P. Zeldes (1991). The consumption of stockholders and nonstockholders. *Journal of Financial Economics 29*(1), 97–112.

Massa, M. and A. Simonov (2006). Hedging, familiarity and portfolio choice. *Review of Financial Studies 19*(2), 633–685.

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2011). Consumer spending and the economic stimulus payments of 2008.

Rabin, M. and G. Weizsäcker (2009). Narrow bracketing and dominated choices. *American Economic Review 99*(4), 1508–1543.

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*(1), 164–212.

Seasholes, M. S. and N. Zhu (2010). Individual investors and local bias. *The Journal of Finance 65*(5), 1987–2010.

Seru, A., T. Shumway, and N. Stoffman (2010). Learning by trading. *Review of Financial studies 23*(2), 705–739.

Shiller, R. J. (2015). *Irrational exuberance*. Princeton University Press.

Statman, M., S. Thorley, and K. Vorkink (2006). Investor overconfidence and trading volume. *Review of Financial Studies 19*(4), 1531–1565.

Stein, J. C. (1995). Prices and trading volume in the housing market: A model with down-payment effects. *The Quarterly Journal of Economics 110*(2), 379–406.

Strahilevitz, M. A., T. Odean, and B. M. Barber (2011). Once burned, twice shy: How naïve learning, counterfactuals, and regret affect the repurchase of stocks previously sold. *Journal of Marketing Research 48*(SPL), S102–S120.

Thaler, R. H. and E. J. Johnson (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management science 36*(6), 643–660.

Tversky, A. and D. Kahneman (1991). Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics*, 1039–1061.

Vissing-Jorgensen, A. (2002). Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures. *National Bureau of Economic Research*.

Weber, M. and C. F. Camerer (1998). The disposition effect in securities trading: An experimental analysis. *Journal of Economic Behavior & Organization 33*(2), 167–184.

Table 1

Example IPO allocation process:

Barak Valley Cement IPO Allocation

Each row in this table refers to the set of investors that applied for the number of shares in Column (1) of the Table. Column (2) is the number of retail investor applications received at that share level. Column (3) is the Total shares applied for at that share level (Column (1)*Column (2)). Column (4) is the number of shares an investor at that share level would receive if the allotment was proportional. Column (5) is the probability that an investor in that row's share level would receive an allotment. Column (6) is the total shares allotted to that share level. Column (7) is the number of investors in that share level that would receive an allotment. Column (8) is the number of investors in that share level that would not receive an allotment.

Share Category	Shares Bid For	# Applications	Total Shares	Proportional Allocation	Win Probability	Shares Allocated	# Treatment group	# Control group
(c)	($c \times x$)	a_c	$a_c \times c \times x$	$\frac{cx}{v}$	$\frac{c}{v}$		$\frac{c}{v} \times a_c$	$(1 - \frac{c}{v}) \times a_c$
(0)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1	150	14,052	2,107,800	4	0.027	57,000	380	13,672
2	300	9,893	2,967,900	8	0.054	80,250	535	9,358
3	450	5,096	2,293,200	12	0.081	61,950	414	4,682
4	600	4,850	2,910,000	16	0.108	78,750	525	4,325
5	750	2,254	1,690,500	20	0.135	45,750	305	1,949
6	900	1,871	1,663,900	24	0.162	45,450	304	1,567
7	1050	4,806	5,046,300	28	0.189	136,500	910	3,896
8	1200	2,900	3,480,000	32	0.216	94,050	628	2,272
9	1350	481	649,350	36	0.244	17,550	117	364
10	1500	1,302	1,953,000	41	0.271	52,800	352	950
11	1650	266	436,900	45	0.298	11,850	79	187
12	1800	317	570,600	49	0.325	15,450	103	214
13	1950	174	339,300	53	0.352	9,150	61	113
14	2100	356	747,600	57	0.379	20,250	135	221
15	2250	20,004	45,009,000	61	0.406	1,217,700	8119	11,885

Table 2
IPO Characteristics

This table presents mean and standard deviation of IPO characteristics for IPOs in our sample, by the year in which the IPO took place and in total. Each variable is constructed as follows: Percentage of IPO in numbers is percentage of all IPOs in India that is part of our sample. Percentage of IPOs by value is the percent of total issue value of all IPOs in India in our sample. "% issued (Retail investors)" presents the total issue value set aside to retail investors. "Over-subscription" ratio is measured as the total demand for shares over the total supply of shares for retail investors *without* employees of the firm. Total number of allotted and non-allotted retail investors are computed only for the share categories where randomized allotment took place rounded to the nearest integer. The sectoral composition of the IPOs is based on the Indian National Industrial Classification Code of various sectors of the economy.

	2007	2008	2009	2010	2011	All
	(1)	(2)	(3)	(4)	(5)	(6)
IPOs in sample						
Number of IPOs in sample	10	12	3	24	8	57
Percentage of all IPOs in India	12.04	31.58	17.65	35.82	20.51	23.36
Value of IPOs in sample ($ bn)	0.28	0.42	0.03	1.58	0.34	2.65
Percentage of total value of IPOs in India	2.99	8.77	1.18	14.17	24.62	8.99
% issued (Retail investors excl. employees)						
Mean	32.18	33.25	34.65	31.99	34.87	32.83
Std. Dev	2.51	2.18	0.41	2.77	0.36	2.51
Over-subscription ratio						
Mean	21.95	12.63	1.72	9.28	6.73	11.45
Std. Dev	21.93	17.29	0.77	10.12	5.94	14.59
No. of retail investors ***Allotted***						
Mean	7,157	14,627	998	13,869	2,812	10,622
Std. Dev.	9,154	27,743	1003	22,023	5,062	19,771
No. of retail investors ***Not allotted***						
Mean	22,100	44,688	208	12,586	4,253	19,192
Std. Dev.	26,602	97,730	86	30,697	8,668	50,878
No. of IPOs from different sectors						
Technology	1	2	0	2	0	5
Manufacturing	4	9	3	14	3	33
Other Services	3	1	0	7	3	14
Retail	2	0	0	1	2	5

Table 3

INVESTOR CHARACTERISTICS

Panel (A) presents the sample size, geographic coverage, portfolio size (in US dollars) and age (in months) of investors the month prior to the IPO in our dataset. Panel (B) presents pooled unconditional mean, standard deviation and median of investor characteristics of our sample for 13 months (6 months before and after IPO, and the month of the IPO). Each characteristic is represented within each year of the IPO and in total. Portfolio value is adjusted by end of the month INR-USD exchange rate, rounded to the nearest integer. Age is defined as the number of months since the investor's account was opened. Total number of IPOs applied or allotted is computed as the sum total of investors that have any allocation in IPOs they participated in and applied for in the subset of IPOs for which we observe application information. Gross number of transactions measured as the total number of purchase and sale transactions within the month. Number of securities held measures the end of the month holdings of total number of securities. Disposition is measured as the difference between percentage of gain realized and percentage of losses realized by the investor in each month. Portfolio weight on the IPO sector is computed as the share of the investors' portfolio in the IPO sector.

Panel (A): SAMPLE SIZE, GEOGRAPHIC COVERAGE, INVESTOR SIZE AND AGE

	2007	2008	2009	2010	2011	All
	(1)	(2)	(3)	(4)	(5)	(6)
No. of investors						
No. of rookie investors	250,669	3,026	349	100,270	16,483	370,797
No. of pre-existing investors	726,175	24,477	4,605	536,160	37,190	1,328,607
Total no. of investors	976,844	27,503	4,984	636,430	53,673	1,699,404
States of India						
Gujarat (%)	36.78	45.28	42.90	33.62	28.13	35.44
Maharashtra (%)	22.46	11.49	16.61	20.94	13.91	21.39
Rajasthan (%)	14.44	15.30	17.39	15.89	13.02	15.01
Delhi (%)	4.29	1.92	3.06	4.27	3.52	4.21
Portfolio size (US$) of pre-existing investors						
Mean	8,750	5,633	6,494	8,295	5,175	8,401
Std. Dev.	21,168	15,365	16,259	19,632	12,614	20,265
Median	2,257	1,477	1,963	2,169	1,607	2,181
Age categories % within each year						
Rookie investors	12.98	3.87	1.59	3.66	3.96	9.03
1 − 5 months	22.76	23.03	7.63	7.31	22.73	16.93
6 − 12 months	15.41	22.32	5.45	7.41	9.24	12.30
13 − 24 months	20.61	19.65	27.67	8.35	12.46	15.76
> 24 months	28.24	31.14	57.65	73.27	51.61	45.97

Panel (B): POOLED SUMMARY STATISTICS OF INVESTOR CHARACTERISTICS

	2007	2008	2009	2010	2011	All
	(1)	(2)	(3)	(4)	(5)	(6)
No. of IPOs applied/allotted						
Mean	0.29	0.20	0.52	0.36	0.13	0.31
Std. Dev.	0.73	0.57	0.97	0.81	0.49	0.75
Median	0	0	0	0	0	0
Purchase and sale activity						
Mean	4.27	3.91	6.73	4.36	2.89	4.26
Std. Dev.	17.26	16.22	22.00	25.68	10.26	20.67
Median	1	1	2	1	0	1
Number of securities held						
Mean	9.78	12.32	16.12	13.61	10.35	11.29
Std. Dev.	18.60	23.95	33.91	26.34	21.40	22.12
Median	4	6	7	7	4	5
Disposition (%)						
Mean	5.31	6.48	12.24	8.76	6.19	6.67
Std. Dev.	22.68	23.08	28.39	24.27	21.59	23.34
Median	0	0	0	0	0	0
Portfolio weight on IPO sector (%)						
Mean	7.18	3.91	5.67	10.76	14.42	8.69
Std. Dev.	18.72	14.08	15.53	22.53	27.73	20.59
Median	0	0	0	0	0	0

Table 4
FUTURE IPO PARTICIPATION

Future IPO participation for an IPO investor is measured as an indicator variable taking the value 1 if the investor is either an IPO applicant in our sample of 57 IPOs or has received an IPO allocation in other IPOs that are not in our applications data, but observed in the trades database. This is measured *without* the IPO security (denoted by "w/o IPO security") in which the investor participates in the lottery. This table presents the regression results for the full sample of investors (Panel A), heterogeneous treatment effects by portfolio size (Panel B and C), and first-day return experience (Panel D and E) for a total of 13 months, i.e., 6 months before and after the IPO month and the month of the IPO. The first row presents the treatment effect in percentage-points, with robust standard errors clustered by calendar month parenthesis and the mean of the dependent variable (in percent) for the control group in square-brackets in the following rows. The control variables include age of the investor, whether the bid type was "cutoff" or "full" and the type of application "asba" or cheque-based application.

	Event-time												
	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6
Panel A: Full sample (N= 1699404)													
w/o IPO	0.035	0.137**	-0.021	0.091	-0.114**	0.013	0.167*	0.854***	0.646**	0.269*	0.154**	0.261**	0.110
security	(0.068)	(0.061)	(0.098)	(0.110)	(0.054)	(0.113)	(0.100)	(0.144)	(0.295)	(0.139)	(0.072)	(0.104)	(0.099)
	[19.890]	[31.110]	[18.690]	[20.010]	[28.150]	[31.670]	[50.710]	[43.680]	[21.490]	[12.830]	[8.890]	[12.910]	[6.760]
Panel B: Above median portfolio size (N = 849289)													
w/o IPO	0.087	0.253**	0.067	0.132	-0.119	-0.101	-0.046	0.665***	0.404*	0.184*	0.063	0.172*	0.046
security	(0.111)	(0.100)	(0.133)	(0.130)	(0.105)	(0.139)	(0.145)	(0.109)	(0.212)	(0.096)	(0.040)	(0.101)	(0.067)
	[26.960]	[39.970]	[24.100]	[24.380]	[34.220]	[36.740]	[52.440]	[45.700]	[22.990]	[13.520]	[9.300]	[13.790]	[7.080]
Panel C: Below median portfolio size (N= 850115)													
w/o IPO	-0.003	0.025	-0.080	0.073	-0.113**	0.121	0.377***	1.031***	0.901**	0.362*	0.247**	0.361**	0.176
security	(0.066)	(0.060)	(0.078)	(0.131)	(0.050)	(0.090)	(0.128)	(0.195)	(0.397)	(0.207)	(0.123)	(0.151)	(0.142)
	[12.820]	[22.250]	[13.280]	[15.650]	[22.080]	[26.600]	[48.990]	[41.670]	[19.990]	[12.140]	[8.490]	[12.020]	[6.430]
Panel D: Positive first-day return (N= 1579470)													
w/o IPO	0.064	0.151**	0.027	0.118	-0.082	-0.009	0.174	0.944***	0.724**	0.300**	0.191**	0.320***	0.133
security	(0.081)	(0.073)	(0.115)	(0.108)	(0.060)	(0.120)	(0.106)	(0.151)	(0.299)	(0.146)	(0.089)	(0.117)	(0.106)
	[19.590]	[31.060]	[17.940]	[19.800]	[28.180]	[30.590]	[49.820]	[44.250]	[20.800]	[13.090]	[8.900]	[12.950]	[6.950]
Panel E: Negative first-day return (N= 119934)													
w/o IPO	-0.430	-0.089	-0.802***	-0.350*	-0.624***	0.378*	0.051	-0.599**	-0.617	-0.238	-0.450	-0.694	-0.259
security	(0.302)	(0.313)	(0.284)	(0.204)	(0.179)	(0.220)	(0.303)	(0.262)	(0.386)	(0.225)	(0.449)	(0.463)	(0.200)
	[25.680]	[32.050]	[33.150]	[24.070]	[27.580]	[52.530]	[67.960]	[32.660]	[34.990]	[7.750]	[8.820]	[12.140]	[3.040]

Significance: *** 0.01 ** 0.05, * 0.10,(clustered robust std. errors), [Mean Dep. Variable - Control group]

Table 5
Activity

Activity is measured as the log of total purchase and sale transactions (+1) undertaken by an investor within each month. This is measured both *with* and *without* the IPO security for the full sample results (Panel A). Results are presented for *without* the IPO security in which the investor participates in the lottery for heterogenous treatment effects by portfolio size (Panel B and C) and first-day return experience (Panel D and E). For each set of results reported, the first row presents the treatment effect in percentage-points, with robust standard errors clustered by calendar month in parenthesis and the mean of the dependent variable (in **number of transactions**) for the control group in square-brackets in the following rows. The control variables include age of the investor, whether the bid type was "cutoff" or "full" and the type of application "asba" or cheque-based application.

	Event-time												
	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6
Panel A: Full sample (N= 1699404)													
with IPO	0.140	0.070	0.181	0.077	0.032	0.172*	46.701***	11.130***	2.522***	1.760***	1.266***	0.918***	0.966***
security	(0.148)	(0.192)	(0.161)	(0.161)	(0.145)	(0.101)	(2.045)	(1.275)	(0.380)	(0.382)	(0.312)	(0.228)	(0.239)
	[1.930]	[2.040]	[2.150]	[2.340]	[2.630]	[3.050]	[3.190]	[3.140]	[2.740]	[2.180]	[2.090]	[2.320]	[2.110]
w/o IPO	0.140	0.070	0.181	0.077	0.032	0.171*	0.637**	1.721***	1.553***	1.035***	0.729***	0.593***	0.673***
security	(0.148)	(0.192)	(0.161)	(0.161)	(0.145)	(0.101)	(0.279)	(0.221)	(0.281)	(0.264)	(0.255)	(0.201)	(0.215)
	[1.930]	[2.040]	[2.150]	[2.340]	[2.630]	[3.050]	[3.180]	[3.130]	[2.740]	[2.180]	[2.090]	[2.320]	[2.110]
Panel B: Above median portfolio size (N = 849289)													
w/o IPO	0.471***	0.310	0.334	0.044	0.010	0.078	0.407	1.126***	0.950***	0.712**	0.373**	0.239	0.439*
security	(0.168)	(0.202)	(0.218)	(0.177)	(0.223)	(0.313)	(0.252)	(0.350)	(0.325)	(0.297)	(0.166)	(0.212)	(0.267)
	[2.740]	[2.970]	[3.230]	[3.660]	[4.160]	[4.850]	[4.630]	[4.150]	[3.480]	[2.760]	[2.630]	[2.850]	[2.560]
Panel C: Below median portfolio size (N= 850115)													
w/o IPO	-0.161	-0.119	0.079	0.158	0.112	0.269*	0.885**	2.413***	2.164***	1.412***	1.137**	1.008***	0.939***
security	(0.184)	(0.153)	(0.116)	(0.145)	(0.130)	(0.158)	(0.413)	(0.294)	(0.407)	(0.414)	(0.494)	(0.360)	(0.278)
	[1.360]	[1.400]	[1.440]	[1.490]	[1.660]	[1.930]	[2.190]	[2.370]	[2.150]	[1.720]	[1.660]	[1.890]	[1.730]
Panel D: Positive first-day return (N= 1579470)													
w/o IPO	0.161	0.136	0.241	0.170	0.129	0.207**	0.684**	1.892***	1.762***	1.154***	0.791***	0.709***	0.780***
security	(0.186)	(0.228)	(0.187)	(0.198)	(0.166)	(0.099)	(0.281)	(0.200)	(0.260)	(0.259)	(0.246)	(0.169)	(0.213)
	[1.910]	[2.020]	[2.140]	[2.320]	[2.610]	[3.030]	[3.150]	[3.120]	[2.740]	[2.190]	[2.090]	[2.320]	[2.110]
Panel E: Negative first-day return (N= 119934)													
w/o IPO	-0.190	-0.999	-0.772	-1.413***	-1.514***	-0.406	-0.117	-1.007	-1.778***	-0.872	-0.270	-1.255	-1.044
security	(0.633)	(0.846)	(0.905)	(0.506)	(0.430)	(0.691)	(0.909)	(0.736)	(0.410)	(0.608)	(0.761)	(0.960)	(0.842)
	[2.400]	[2.460]	[2.390]	[2.770]	[3.000]	[3.590]	[3.950]	[3.360]	[2.610]	[2.020]	[2.040]	[2.340]	[1.950]

Significance: *** 0.01 ** 0.05, * 0.10,(clustered robust std. errors), [Mean Dep. Variable - Control group]

Table 6
DISPOSITION

Disposition is measured as the difference between the percentage of gains realized (PLR) and the percentage of losses realized (PGR) at the end of the month by an investor. This is measured both *with* and *without* the IPO security for the full sample results (Panel A). Results are presented for *without* the IPO security in which the investor participates in the lottery for heterogenous treatment effects by portfolio size (Panel B and C) and first-day return experience (Panel D and E). For each set of results reported, the first row presents the treatment effect in percentage-points, with robust standard errors clustered by calendar month in parenthesis and the mean of the dependent variable (in percent) for the control group in square-brackets in the following rows. The control variables include age of the investor, whether the bid type was "cutoff" or "full" and the type of application "asba" or cheque-based application.

	Event-time												
	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6
Panel A: Full sample (N= 1699404)													
with IPO	-0.027	-0.024	0.050	-0.049	-0.049	-0.043	2.135	4.095***	0.150	0.199**	0.245***	0.128**	0.171***
security	(0.045)	(0.063)	(0.040)	(0.060)	(0.039)	(0.041)	(1.350)	(0.890)	(0.128)	(0.092)	(0.072)	(0.056)	(0.051)
	[5.360]	[4.220]	[5.380]	[6.500]	[6.260]	[8.520]	[6.830]	[10.080]	[3.350]	[4.890]	[5.990]	[5.770]	[4.510]
w/o IPO	-0.027	-0.024	0.050	-0.049	-0.049	-0.043	0.363*	0.769***	0.079**	0.042	0.130***	0.061	0.082
security	(0.045)	(0.063)	(0.040)	(0.060)	(0.039)	(0.041)	(0.211)	(0.134)	(0.036)	(0.045)	(0.032)	(0.040)	(0.050)
	[5.360]	[4.220]	[5.380]	[6.500]	[6.260]	[8.520]	[6.830]	[10.060]	[3.330]	[4.880]	[5.980]	[5.770]	[4.510]
Panel B: Above median portfolio size (N = 849289)													
w/o IPO	-0.015	-0.048	0.043	-0.114	-0.046	-0.020	0.111	0.518***	0.056	-0.033	0.060	0.005	0.056
security	(0.056)	(0.099)	(0.068)	(0.071)	(0.052)	(0.067)	(0.167)	(0.098)	(0.039)	(0.047)	(0.046)	(0.050)	(0.068)
	[7.780]	[5.860]	[7.740]	[9.000]	[8.030]	[9.300]	[9.680]	[10.720]	[4.200]	[5.380]	[6.760]	[6.560]	[4.970]
Panel C: Below median portfolio size (N= 850115)													
w/o IPO	-0.041	0.013	0.053	0.015	-0.046	-0.061	0.602**	1.024***	0.105*	0.111*	0.194***	0.113	0.098*
security	(0.046)	(0.048)	(0.055)	(0.060)	(0.050)	(0.101)	(0.278)	(0.217)	(0.054)	(0.059)	(0.067)	(0.069)	(0.054)
	[2.950]	[2.580]	[3.020]	[4.000]	[4.500]	[7.730]	[3.980]	[9.410]	[2.460]	[4.380]	[5.210]	[4.970]	[4.050]
Panel D: Positive first-day return (N= 1579470)													
w/o IPO	-0.025	-0.023	0.059	-0.059	-0.025	-0.038	0.377*	0.819***	0.070*	0.055	0.127***	0.061	0.091*
security	(0.057)	(0.064)	(0.040)	(0.062)	(0.041)	(0.040)	(0.226)	(0.139)	(0.041)	(0.044)	(0.031)	(0.040)	(0.055)
	[5.220]	[4.190]	[5.260]	[6.410]	[6.150]	[8.510]	[6.460]	[10.040]	[3.240]	[4.800]	[5.970]	[5.700]	[4.480]
Panel E: Negative first-day return (N= 119934)													
w/o IPO	-0.065	-0.036	-0.092	0.105	-0.433**	-0.114	0.141	-0.025	0.225**	-0.157	0.172	0.061	-0.065
security	(0.227)	(0.101)	(0.244)	(0.107)	(0.193)	(0.249)	(0.126)	(0.210)	(0.085)	(0.170)	(0.195)	(0.171)	(0.172)
	[8.020]	[4.820]	[7.780]	[8.190]	[8.420]	[8.660]	[13.990]	[10.510]	[5.250]	[6.410]	[6.150]	[7.100]	[5.100]

Significance: *** 0.01 ** 0.05, * 0.10,(clustered robust std. errors), [Mean Dep. Variable - Control group]

Table 7
Familiarity

Familiarity is measured as the portfolio weight (in percent) of the industry to which the IPO belongs to. If the investor applies for a information technology stock in our sample, the portfolio weight on the information technology industry is measured at the end of the month. This is measured both *with* and *without* the IPO security for the full sample results (Panel A). Results are presented for *without* the IPO security in which the investor participates in the lottery for heterogenous treatment effects by portfolio size (Panel B and C) and first-day return experience (Panel D and E). For each set of results reported, the first row presents the treatment effect in percentage-points, with robust standard errors clustered by calendar month in parenthesis and the mean of the dependent variable (in percent) for the control group in square-brackets in the following rows. The control variables include age of the investor, whether the bid type was "cutoff" or "full" and the type of application "asba" or cheque-based application.

	Event-time												
	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6
Panel A: Full sample (N= 1699404)													
with IPO	0.032	0.020	0.031	0.034	0.007	0.048**	21.600***	9.599***	7.332***	6.542***	6.301***	5.625***	5.247***
security	(0.028)	(0.033)	(0.043)	(0.033)	(0.024)	(0.024)	(3.924)	(1.470)	(0.876)	(0.873)	(0.884)	(0.800)	(0.802)
	[4.980]	[5.350]	[5.070]	[4.800]	[8.480]	[6.060]	[7.550]	[6.980]	[8.020]	[7.860]	[8.000]	[8.100]	[7.700]
w/o IPO	0.031	0.020	0.031	0.034	0.007	0.047**	0.036	0.009	0.042	0.071**	0.079**	0.069**	0.048**
security	(0.028)	(0.033)	(0.043)	(0.033)	(0.024)	(0.024)	(0.029)	(0.040)	(0.026)	(0.027)	(0.036)	(0.027)	(0.017)
	[4.980]	[5.350]	[5.070]	[4.800]	[8.480]	[6.060]	[7.500]	[6.860]	[7.900]	[7.750]	[7.890]	[8.000]	[7.600]
Panel B: Above median portfolio size (N = 849289)													
w/o IPO	0.059***	0.001	-0.002	0.011	0.046**	0.032	0.010	0.009	0.013	0.047**	0.069**	0.047**	0.045**
security	(0.021)	(0.037)	(0.065)	(0.024)	(0.023)	(0.020)	(0.019)	(0.024)	(0.025)	(0.023)	(0.026)	(0.023)	(0.022)
	[6.960]	[7.700]	[7.290]	[7.130]	[9.050]	[8.410]	[8.700]	[8.350]	[8.810]	[8.720]	[8.760]	[8.810]	[8.520]
Panel C: Below median portfolio size (N= 850115)													
w/o IPO	-0.011	0.021	0.051	0.044	-0.037	0.050	0.042	-0.007	0.065	0.089**	0.081	0.085	0.043
security	(0.047)	(0.047)	(0.039)	(0.049)	(0.038)	(0.037)	(0.059)	(0.065)	(0.052)	(0.045)	(0.058)	(0.052)	(0.032)
	[3.010]	[3.000]	[2.840]	[2.470]	[7.900]	[3.710]	[6.300]	[5.380]	[7.000]	[6.800]	[7.020]	[7.180]	[6.690]
Panel D: Positive first-day return (N= 1579470)													
w/o IPO	0.039	0.022	0.033	0.038	0.008	0.050**	0.034	0.010	0.047*	0.078***	0.085**	0.077***	0.057***
security	(0.029)	(0.033)	(0.045)	(0.034)	(0.025)	(0.025)	(0.030)	(0.042)	(0.027)	(0.028)	(0.036)	(0.027)	(0.016)
	[5.120]	[5.500]	[5.200]	[4.920]	[8.780]	[6.230]	[7.580]	[7.050]	[8.130]	[7.980]	[8.130]	[8.230]	[7.820]
Panel E: Negative first-day return (N= 119934)													
w/o IPO	-0.088*	-0.029	-0.018	-0.037	-0.017	0.006	0.060	-0.015	-0.033	-0.037	-0.022	-0.056	-0.093
security	(0.052)	(0.055)	(0.064)	(0.070)	(0.074)	(0.090)	(0.068)	(0.045)	(0.063)	(0.081)	(0.080)	(0.056)	(0.064)
	[2.250]	[2.400]	[2.410]	[2.470]	[2.610]	[2.710]	[5.880]	[3.210]	[3.490]	[3.360]	[3.270]	[3.430]	[3.390]

Significance: *** 0.01 ** 0.05, * 0.10,(clustered robust std. errors), [Mean Dep. Variable - Control group]

Table 8

DIVERSIFICATION

Diversification is measured as the log of total number of securities (+1) held at the end of the month by an investor. This is measured both *with* and *without* the IPO security for the full sample results (Panel A). Results are presented for *without* the IPO security in which the investor participates in the lottery for heterogenous treatment effects by portfolio size (Panel B and C) and first-day return experience (Panel D and E). For each set of results reported, the first row presents the treatment effect in percentage-points, with robust standard errors clustered by calendar month in parenthesis and the mean of the dependent variable (presented as **number of securities held**) for the control group in square-brackets in the following rows. The control variables include age of the investor, whether the bid type was "cutoff" or "full" and the type of application "asba" or cheque-based application.

	Event-time												
	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6
Panel A: Full sample (N= 1699404)													
with IPO	0.233	0.168	0.209	0.181	0.059	0.028	24.177***	17.942***	10.541***	9.587***	9.047***	8.617***	8.414***
security	(0.228)	(0.218)	(0.220)	(0.210)	(0.223)	(0.193)	(2.108)	(1.975)	(1.229)	(1.280)	(1.317)	(1.393)	(1.382)
	[3.450]	[3.650]	[3.910]	[4.180]	[4.610]	[5.220]	[6.250]	[6.920]	[7.240]	[7.130]	[7.080]	[7.080]	[7.090]
w/o IPO	0.233	0.168	0.209	0.181	0.059	0.028	0.272	0.643***	0.726***	0.561***	0.511***	0.547***	0.572***
security	(0.228)	(0.218)	(0.220)	(0.210)	(0.223)	(0.193)	(0.228)	(0.127)	(0.108)	(0.137)	(0.156)	(0.142)	(0.153)
	[3.450]	[3.650]	[3.910]	[4.180]	[4.610]	[5.220]	[6.250]	[6.910]	[7.230]	[7.120]	[7.070]	[7.070]	[7.080]
Panel B: Above median portfolio size (N = 849289)													
w/o IPO	0.556***	0.448**	0.368	0.259	0.141	0.033	0.347*	0.392*	0.468**	0.398**	0.389**	0.454**	0.517***
security	(0.197)	(0.184)	(0.215)	(0.171)	(0.168)	(0.165)	(0.184)	(0.209)	(0.196)	(0.198)	(0.193)	(0.220)	(0.171)
	[6.960]	[7.580]	[8.370]	[9.220]	[10.330]	[11.810]	[14.200]	[14.490]	[14.550]	[14.270]	[14.030]	[13.830]	[13.740]
Panel C: Below median portfolio size (N= 850115)													
w/o IPO	-0.033	-0.048	0.107	0.170	0.060	0.076	0.229	0.963***	1.012***	0.739***	0.645**	0.657***	0.624***
security	(0.162)	(0.186)	(0.182)	(0.131)	(0.130)	(0.138)	(0.237)	(0.125)	(0.168)	(0.153)	(0.235)	(0.145)	(0.206)
	[1.710]	[1.760]	[1.830]	[1.900]	[2.060]	[2.310]	[2.750]	[3.300]	[3.600]	[3.550]	[3.570]	[3.610]	[3.650]
Panel D: Positive first-day return (N= 1579470)													
w/o IPO	0.217	0.172	0.215	0.192	0.081	0.034	0.291	0.705***	0.803***	0.649***	0.583***	0.614***	0.638***
security	(0.260)	(0.260)	(0.253)	(0.240)	(0.259)	(0.221)	(0.259)	(0.166)	(0.135)	(0.175)	(0.195)	(0.185)	(0.191)
	[3.410]	[3.600]	[3.860]	[4.130]	[4.550]	[5.140]	[6.150]	[6.820]	[7.160]	[7.060]	[7.010]	[7.010]	[7.030]
Panel E: Negative first-day return (N= 119934)													
w/o IPO	0.489	0.103	0.118	0.007	-0.296	-0.073	-0.049	-0.356	-0.511	-0.859	-0.645	-0.531	-0.487
security	(0.887)	(1.105)	(0.928)	(0.765)	(0.732)	(0.731)	(1.039)	(0.941)	(0.927)	(0.904)	(0.925)	(0.984)	(0.822)
	[4.460]	[4.690]	[5.020]	[5.420]	[5.940]	[6.980]	[8.570]	[8.880]	[8.750]	[8.400]	[8.340]	[8.340]	[8.160]

Significance: *** 0.01 ** 0.05, * 0.10,(clustered robust std. errors), [Mean Dep. Variable - Control group]

Figure 1
IPO FREQUENCY

This figure presents the monthly time series of our sample of IPOs compared with all the IPOs that took place between 2007 and 2011.

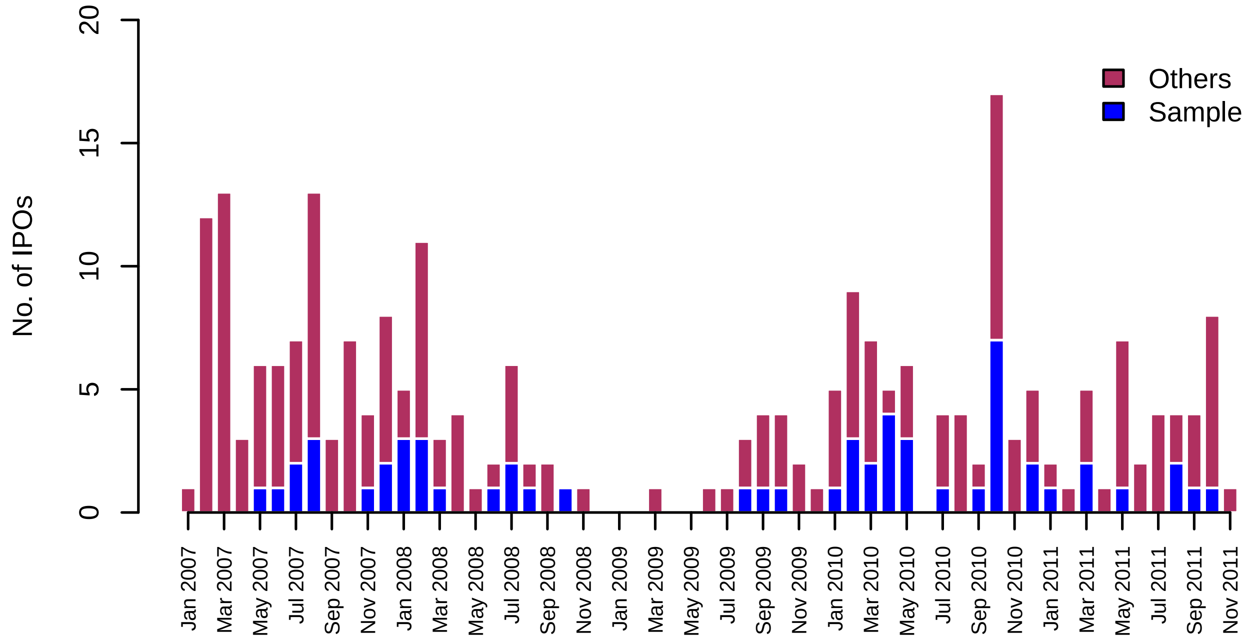

Figure 2
IPO Investor Experience

This figure plots the first-day returns and price-variability experienced by the sample of IPO investors. First-day returns are computed as the first day returns on the issue price of the IPO. First-day price-variability is estimated with the first-day high, low and the issue price of the IPO stocks. The number of IPOs in each year is denoted below each year in "[.]"

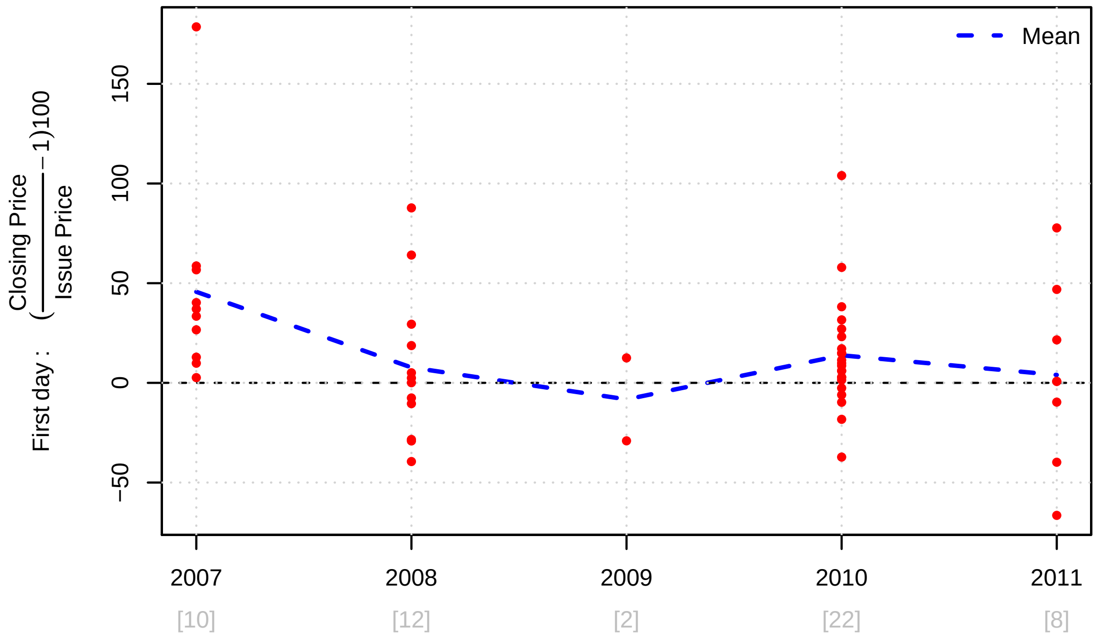

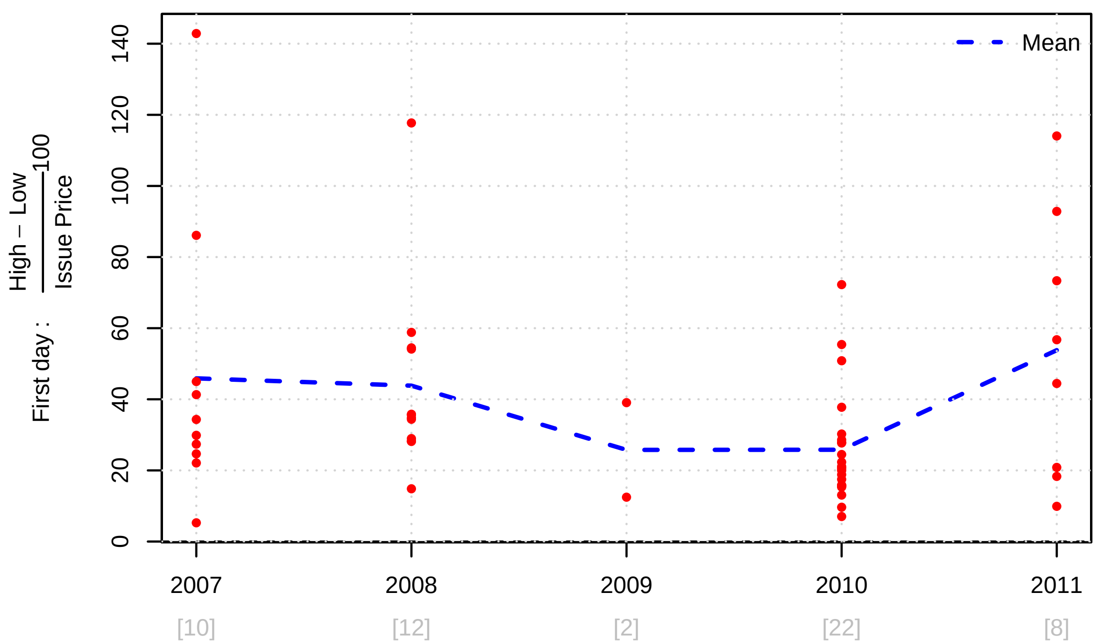

Figure 3

Test of Balance on Application Characteristics of IPO Investors

We run the following regression specification for an array of application characteristics within each of the IPO share categories across 57 IPOs, our sample:

$$y = \alpha + \beta I(success = 1) + u$$

Here, y is the outcome variable such as whether the applicant used CDSL or NSDL as their depository, the nature of the bid - whether it was an application with an explicit demand schedule ("Full" bid) or one just for a cutoff price ("Cutoff" bid), the modality of payment - whether the application was supported by a bank statement (ASBA) or by cheque or other means of financial backing and an indicator variable for the major IPO states of India - Gujarat, Maharashtra and Rajasthan. Within each IPO share category, we expect β, the coefficient on the indicator variable whether the applicant was successful in the lottery, to be statistically insignificant. This figure plots the distribution of the t-statistic obtained from these regressions for all pre-treatment characteristics against the normal distribution (dashed line). We expect the distribution of the t-statistic to be normal, with 2.5% of false-significance on each tail. The vertical dotted lines represent -1.96 and 1.96, at 5% significance. Statistical test of difference between the two distributions using the Kolmogorov-Smirnov Test suggests that the they are not significantly different from each other, with the $ks-$statistic of 0.0265 and a $p-$value of 0.124.

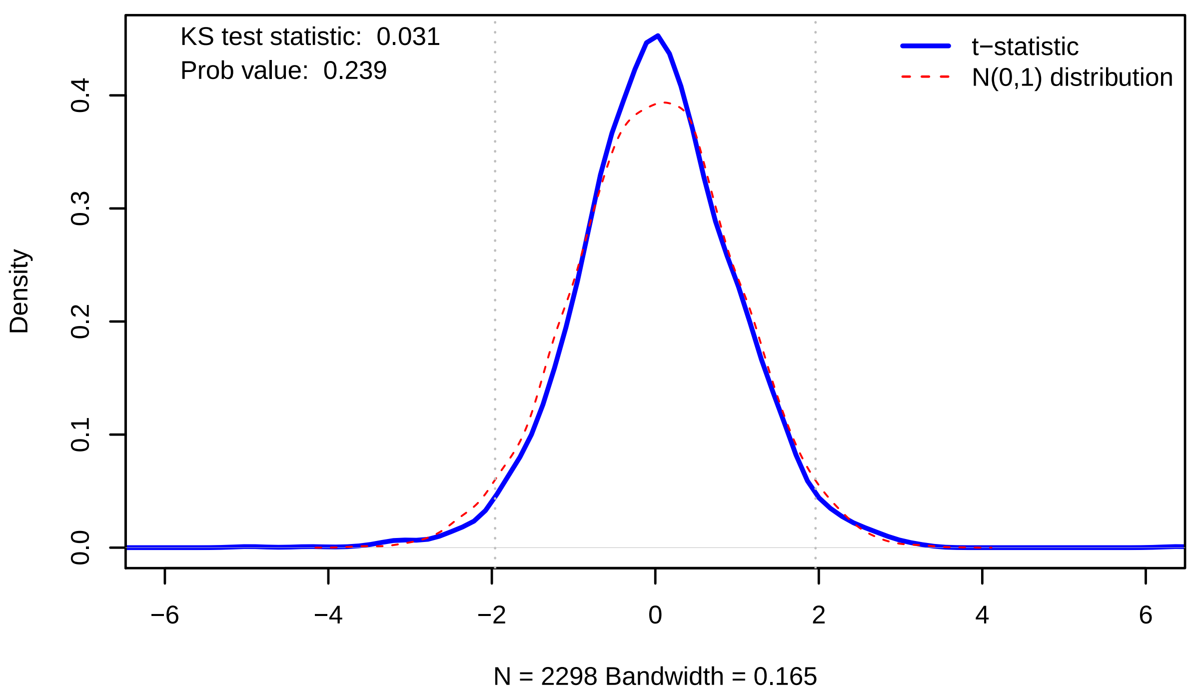

Figure 4
IPO level variation in treatment effects

We present the IPO level variation in treatment effects on investors' likelihood of participating in future ipos (Panel A), the gross (purchase and sale) transactions (Panel B), disposition (Panel C), portfolio weight on the IPO sector (Panel D) and the number of stocks held (Panel E). All reported data points are computed without the IPO stock, one month after the IPO allocation, thus representing the effect on the portfolio of the investor.

Panel A: Future IPO participation

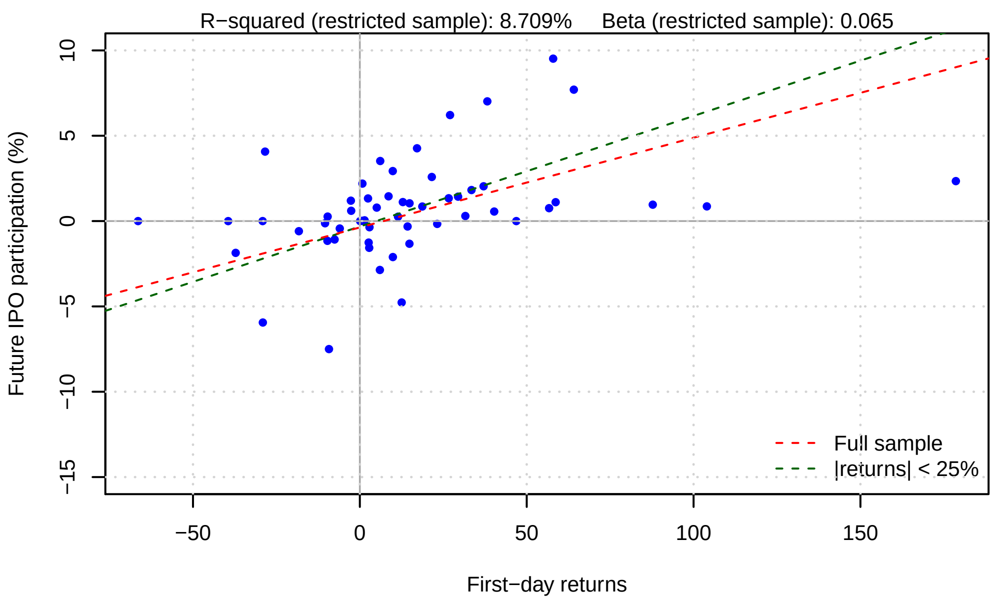

Panel B: Activity

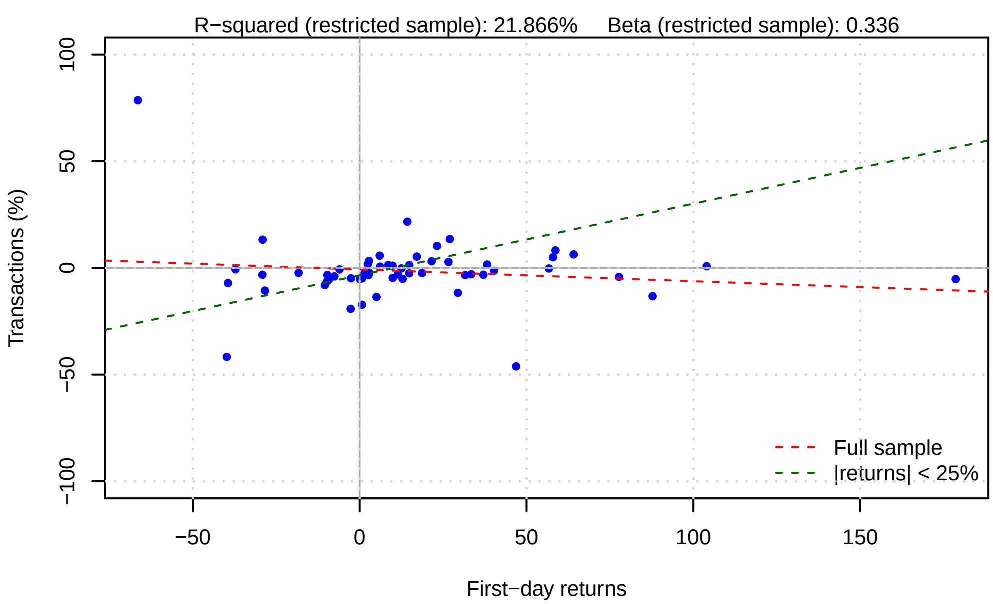

Panel C: Disposition

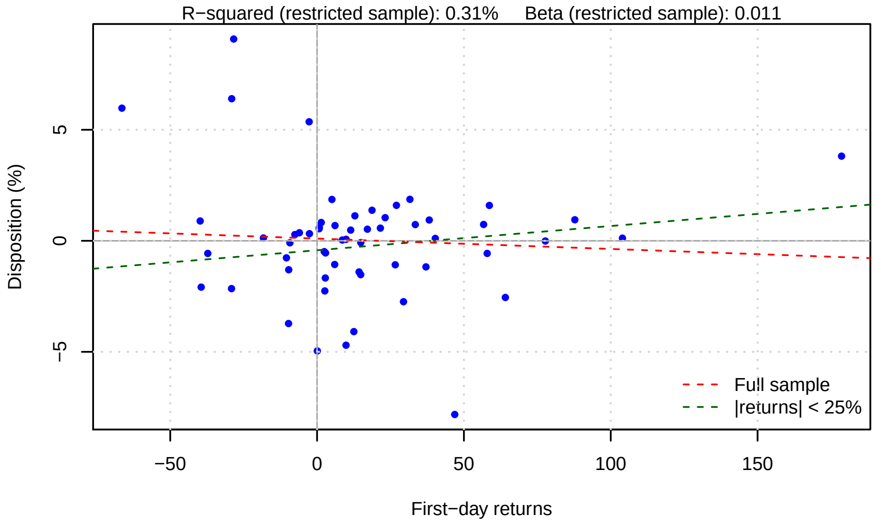

Panel D: Familiarity

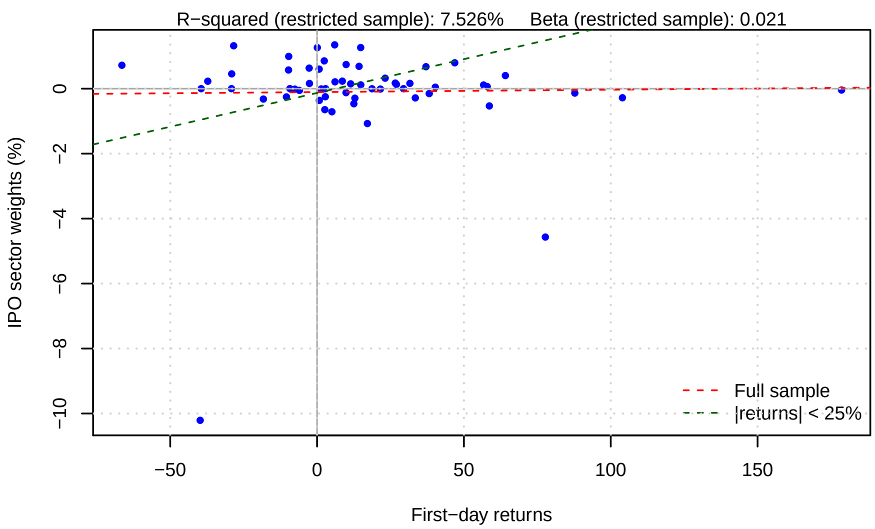

Panel E: Diversification

7 Appendix

7.1 Calculating the Probability of Winning

We begin from equation (5) in the paper. In that equation, we substitute for Z from equation (3) and use equation (4) to re-express $p_{c'}$ for share categories $c' \in [1, ..., J)$ in terms of p_1. We then arrive at:

$$\sum_{c'=1}^{J-1} c' p_1 x a_{c'} = S - \sum_{c=J}^{C} a_c \frac{cx}{v}. \tag{9}$$

Substituting for S from equation (2), we get that:

$$p_1 = \frac{\frac{1}{v} \sum_{j=1}^{C} a_c cx - \sum_{c=J}^{C} a_c \frac{cx}{v}}{\sum_{c'=1}^{J-1} c' x a_{c'}}, \tag{10}$$

which gives $p_1 = \frac{1}{v}$, and $p_{c'} = \frac{c'}{v}$ for randomized share categories $c' \in [1, ..., J)$.

We quickly demonstrate that this probability is well-defined, i.e., $0 < \frac{c'}{v} < 1$. Recall that the regulation requires randomization when the proportional allocation cannot allocate at least the minimum lot size of shares. Consider $c' = (J-1)$, which is the final share category in which proportional allocation is not possible, and random allocation must take place. That is:

$$\frac{(J-1)x}{v} < x \Longrightarrow (J-1) < v$$

This will also be true for all values of $0 < c' < (J-1)$. Further, since $v > 0$ and $c' > 0$, $\frac{c'}{v} > 0$. Thus, $0 < \frac{c'}{v} < 1$.