

Endowment Effects in the Field: Evidence from India's IPO Lotteries*

Santosh Anagol[†] Vimal Balasubramaniam[‡] Tarun Ramadorai[§]

December 11, 2015

Abstract

Winners of randomly assigned initial public offering (IPO) lottery shares are significantly more likely to hold the shares of the company than lottery losers 1, 6, and even 24 months after the lottery allocation. These effects persist in samples of wealthy and highly active players, and we provide additional evidence that this type of “endowment” effect is not driven by wealth effects or inertia alone. The effect decreases as experience in the IPO market increases, but persists even for the most experienced players. These results suggest that players’ preferences and/or beliefs about an asset are not independent of ownership, providing field evidence derived from the behavior of 1.5 million Indian stock investors which is in line with the large laboratory literature that documents such effects.

*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 Eduardo Azevedo, Raj Iyer, Ulrike Malmendier, and Shing-Yi Wang for comments.

[†]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

A large laboratory literature documents that the act of owning an object appears to have an important causal effect on the individual's subsequent valuations of the object. These results have been influential because they call into question the basic assumption, common to most economic theory, that preferences and beliefs are independent of ownership. Apart from a few notable exceptions (List et al., 2003; List, 2011), however, less progress has been made in understanding to what extent these “endowment” effects impact behavior in real market settings. This issue is of particular importance given that laboratory findings of endowment effects appear to be sensitive to laboratory experimental procedures (Zeiler and Plott, 2004).

We study a unique setting where millions of real market participants are randomly assigned an object, and where we can observe the future trading of that object. As a result of regulatory requirements, in many cases Indian initial public offering (IPO) shares are randomly assigned to applicants. Under the assumption of minimal wealth effects (which we discuss in detail later), the winners and losers in these IPO lotteries should have the same preferences and information sets before and after the shares are allotted. Thus, while lottery losers do not have the opportunity to buy the shares at the issue price, once the stock begins to trade, both winner and loser groups have equal opportunities to trade the stock. Given the equivalence of their information sets and background characteristics induced by the random assignment, we should expect that the holdings of this randomly allocated stock should converge rapidly over time. If endowment effects are important, however, we should expect a divergence in the behavior of the randomly chosen winners and losers.

We document that the winners of these IPO lotteries are substantially more likely to hold the randomly allocated IPO shares for many months and even years after the allocation. In our main results we find that, on average, 56.4 percent of IPO winners hold the IPO stock at the end of the month, while only 1.4 percent of losers hold the stock. Six months after the lottery assignment the gap decreases slightly, to 46.6 percent of winners holding the stock and 1.6 percent of the losers holding the stock; but even 24 months after the random assignment

we find that winners are 35 percent more likely to hold the IPO stock than are losers. We term this effect an “endowment effect,” with analogies to the laboratory literature, because the fact that winners of these lotteries are randomly endowed with the IPO stock appears to cause them to be less likely to sell it than the losers are likely to buy it.

The laboratory literature on endowment effects has broadly focused on two categories of causal mechanisms.¹ Mechanisms in the first category assume that agents either have preferences where endowments directly affect valuations (e.g. reference dependent preferences with loss aversion), or belief formation processes where endowments directly affect beliefs. In contrast, causal mechanisms in the second category assume preferences and beliefs that are independent of endowments, and identify other features of the empirical setting (such as transactions costs associated with trading) which cause those endowed to prefer greater holdings of the asset. While the laboratory literature has attempted to completely rule out certain mechanisms while testing for the presence of others, we note that in our field setting (and very likely in most field settings) various mechanisms from both categories could be at play. In our analysis, we are primarily interested in investigating various causes to better understand whether experience appears to eliminate the endowment effect that we detect; in particular, there are some types of causes that we would expect to grow less important as experience is accumulated (e.g., behavioral biases that generate pecuniary losses for investors), and others where experience would matter less (e.g., wealth effects or transactions costs).

Within the first category, the most prominent mechanism explored in the literature is the idea that agents have reference dependent utility with loss-aversion; in this framework agents receive utility from potential transactions relative to a reference point, which in the case of asset markets is typically assumed to be the purchase price. In our setting a lottery winner might “feel” a loss when selling the stock below the issue or open price, and therefore have a tendency to hold it, whereas a loser has no such reference point because they do not own

¹See Ericson and Fuster (2013) for a review of laboratory experiments that estimate the endowment effect and attempt to distinguish causal mechanisms.

the stock. This behavior would be consistent with the well established result that investors are more likely to hold their losses and sell their winners - the “disposition effect” (Shefrin and Statman, 1985; Odean, 1998). One piece of evidence that we detect, consistent with this interpretation, is that lottery winners are substantially more likely to hold the stock if the stock experiences a loss. However, reference dependent preferences with loss aversion of this type, and the related disposition effect literature do not provide an obvious explanation for why lottery *losers* consistently choose *not* to purchase the stock on the open market – this is primarily because this literature has not been able to observe a “control” group of investors with similar background characteristics that did not own the stock. We discuss other theories that might explain why lottery losers rarely ever choose to purchase the IPO stock.

The second category of causal mechanisms assumes that individuals have preferences and beliefs that are independent of endowments, but some feature of the empirical setting causes those endowed to prefer greater holdings of the asset. We provide a comprehensive discussion of these issues later in the paper, and focus here on two prominent examples of such features, namely, wealth effects and transactions costs/inertia. The winners in our experiment do receive a wealth shock due to the fact that Indian IPOs are typically underpriced relative to the first price they trade at on the market. However, we argue that the magnitude of the wealth effect the winners receive is unlikely to explain the divergence in holding patterns because (1) the actual gains in the treatment IPO (on average 62 USD) are small relative to the amount of cash participants are required to place in escrow in order to participate in these lotteries (on average 1,800 USD), suggesting it is unlikely that lottery winnings are relieving a major wealth constraint, (2) standard portfolio choice models do not predict that small increases in wealth should be allocated primarily to one stock, and (3) we find that endowment effects are similar for IPOs associated with low and high wealth shocks (where the wealth shock is defined as the listing return on the holding).

Transactions costs are another potential explanation that could drive the results even

if agents have preferences and/or beliefs that are independent of endowments. Brokerage commissions, transactions taxes, or just nuisance factors (e.g., the cost of accessing your account and placing a trade) could make winners appear less willing to sell the stock and make losers unlikely to buy it. A few pieces of evidence suggest that this is not the only causal mechanism at play. First, in related work we find that lottery winners are more likely to trade non-IPO stocks in their portfolio than lottery losers, suggesting that, if anything, winning the lottery textit reduces the inertia associated with making a trade (Anagol et al., 2015). It is unclear why winning the lottery would increase the probability of trading in non-IPO stocks, but simultaneously cause the winners to hold the IPO stock even though they might really prefer to sell it. Second, we find large divergences in behavior even for those who recently have made a large number of trades in their account. Thus, the results are not being driven by a subset of investors who are mainly inactive, and should assuage concerns that our results are being driven by “trade uncertainty” (Engelmann and Hollard, 2010). Third, the divergence between the holdings of winners and losers is much larger when the IPO stock performs poorly compared to when it performs well. If the only factor causing the divergence was transaction costs we would not expect such a strong relationship between IPO performance and the effect size. Fourth, a transactions cost explanation would be most appealing if the tendency of winners to hold the stock relative to losers did not cause the winners to have better or worse performance. However, on average, winners choosing to hold the IPO stock after it starts trading lose 36 percent (on average) over six months; so there is a reasonable incentive to overcome small transactions costs and sell the stock. While this evidence does not rule out the transactions costs explanation completely, we believe it is hard to argue that the only factor driving the results are transactions costs.

Given that it appears that simple features of our experimental setting, such as wealth effects or transactions costs, do not explain the majority of our findings, it is interesting to analyze the relationship between market experience and the endowment effect in this setting. We conduct two types of analyses along these lines. The first is a descriptive

analysis, where we compare the strength of this endowment effect across investors with different levels of experience in the IPO market. For each investor we observe the number of IPOs they have previously been allotted over the past 10 years. This measure of experience varies from 0 previous experiences up to 30 previous experiences at the 90th percentile of the distribution. Consistent with List et al. (2003), we find a strong negative correlation between this experience measure and the difference in holdings between lottery winners and losers (both with and without a set of portfolio and IPO level controls). However, while List et al. (2003) finds that endowment effects become negligible amongst his sample of experienced traders (i.e. sportscard dealers and very experienced non-dealers), we find substantial endowment effects even amongst investors who have participated in over thirty IPOs: on average these highly experienced winners still hold 35 percent of their allotments at the end of the month of randomly receiving the IPO, while the losers hold 7 percent of the initial allocation. These results are also interesting in light of Haigh and List (2005), who find that professional futures traders exhibit greater myopic loss aversion and raise the possibility that market experience might exacerbate behavioral anomalies; our evidence rejects the idea that more experienced market participants exhibit the endowment effect anomaly more strongly.

While the fact that such experienced lottery winners are so much more likely to hold the stock than similarly experienced losers is suggestive that experience does not eliminate this anomaly, it is possible that this correlation is confounded by selection effects. For example, our experience measure might be correlated with some unobserved factor that causes more experienced winners to hold the stock more than similarly experienced lottery losers (i.e., the negative effect of experience on the divergence of holdings between winners and losers are somewhat canceled out by this omitted factor when we estimate correlations). We note that this type of selection contradicts the most commonly assumed selection bias as discussed in List et al. (2003) and List (2011): those with more experience are typically thought to be *more* likely to trade in endowment experiments due to unobserved factors. Nonetheless, we

cannot rule out the presence of such unobserved factors based on correlations alone.

To make some progress on this issue, our second analysis exploits the random assignment of previous lotteries to provide a sharper comparison of whether the behavior of more experienced lottery players converges more than that of less experienced lottery players.² We find evidence consistent with such convergence: when we compare the behavior of randomly chosen winners and losers in future IPOs, we find that those who have previously won IPOs have smaller estimated endowment effects in the future. But, similar to the experience correlations discussed above, the rate of learning appears to be slow. Overall, the evidence from these two types of analyses suggests that while experience does substantially reduce this particular endowment effect, it seems unlikely that experience eliminates this anomaly completely.

We highlight three contributions of our work. First, our results suggest that endowment type effects may have important implications for asset markets, in addition to the consumer (mugs, pens) and durable goods (sportscards, collector's pins) markets where they have most commonly been studied. Second, our results shed light on a recent debate regarding the fundamental causes of the endowment effect. Kőszegi and Rabin (2006) present a model where reference points are set based on expectations, which predicts that endowment effects should decline in settings where subjects expect that there will be an opportunity to trade in the near future. Laboratory tests of this prediction, using similar designs, have yielded conflicting results (Heffetz and List, 2014; Ericson and Fuster, 2011; Knetsch and Wong, 2009). Our results appear to reject this prediction, because IPO applicants know that they will have the opportunity to trade the stock after it gets listed; at a minimum, endowment effects cannot be ruled out in real world settings even when participants expect to have the opportunity to trade their endowments. Third, our results lend credence to theoretical frameworks where an agents' decision process regarding an asset fundamentally changes once

²While our main comparison of lottery winners and losers constitutes a randomized experiment, our comparison of past winners and losers in future lotteries has one potentially important selection issue: the choice of whether to participate in future IPOs may depend on previous experience. We discuss how this type of selection might affect this set of experience estimates.

the asset enters their portfolio; lottery winners do not have any standard reasons to value the IPO stock more than lottery losers in our setting, but yet continue to hold it at much higher rates. Exploring the general equilibrium implications of this type of buyer \seller divergence appears to be a fruitful area for future research.³

2 The Experiment: India’s IPO Lotteries

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 a naturally occurring setting where some agents are randomly endowed with an asset, some agents are not endowed, and we can observe trading choices of that asset in the near future. In this section we describe why these lotteries occur, and in the next section describe why they can be used to estimate endowment effects.

To summarize, these IPO lotteries arise 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*. The outcome of the lottery is that some investors receive the minimum lot size (this is the treatment group) and others receive zero shares (the control group).

The fundamental reason for the lottery is that in India, regulations require that a firm must set aside 30% or 35% of its shares (depending on the type of issue) to be available for allocation to retail investors at the time of IPO.⁴ For the purposes of the regulation, “retail

³For example, in our setting it is possible that the reluctance of winners to sell is part of the explanation for why there is predictable long run underperformance of IPOs.

⁴We provide more details on these regulations in the online appendix.

investors” are defined as those with expressed share demands beneath a pre-set value.⁵ At the time of writing, this pre-set value is set by the regulator at Rs. 200,000 (roughly US \$3,400); this value has varied over time.⁶

The share allocation process in an Indian IPO 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 bids (and ultimately, share 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 of the price band multiplied by the desired number of shares. If the investor is allotted shares, and the final issue price is less than the ceiling price, the difference between the deposited and required amounts is refunded to the investor. In our sample 93 % of IPO applicants elect to submit cut-off bids.

Alternatively, retail investors have the option 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. As in the case of the cutoff bid, the investor once again deposits the maximum monetary amount consistent with their demand schedule at the time of submitting their bid.

⁵In practice, each brokerage account is counted as an individual retail investor for the purposes of the regulation, meaning that a single investor could in practice exceed this threshold by subscribing using multiple different brokerage accounts. However, this is not a concern for us as we can identify any such behavior in our data. This is because our data are aggregated across all brokerage accounts associated with the anonymized tax identification number of the investor.

⁶The Indian regulator, SEBI, introduced the definition of a retail investor on August 14, 2003 and capped the amount that retail investors could invest at Rs. 50,000 per brokerage account per IPO. This limit was increased to Rs. 100,000 on March 29, 2005, and once again increased to Rs. 200,000 on November 12, 2010. This regulatory definition technically permits institutions to be classified as retail when investing amounts smaller than the limit, but over our sample period, we verify using independent account classifications from the depositories that this hardly ever occurs, and accounts for a miniscule proportion of retail investment in IPOs. We simply remove these aberrations from our analysis.

If allotted shares, the investor’s order will be filled at the stated share demand associated with the final issue price, and a refund is processed for the difference between the final price and the amount placed in escrow. 7 % of our sample submits full demand schedules.

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 total supply of shares to retail investors, including any excess supply from other categories (for example, if employees and/or non-institutional investors participate in amounts less than they are offered, this can “overflow” into additional retail supply).⁷

We define “retail oversubscription” as the ratio of total retail demand for a firm’s shares to total supply of shares by the firm to retail investors, i.e., the total number of shares made available by the firm for retail investors to purchase.

There are then three possible cases:

1. Retail oversubscription is less than or equal to one. In this case, all retail investors are allotted shares according to their demand schedules.
2. Retail oversubscription is greater than one, and shares can be allocated to investors *in proportion to their stated demands without any violation of the minimum lot size constraint*. There is no lottery involved in this case.
3. Retail oversubscription is *far* greater than one (the issue is substantially oversubscribed), and a number of investors under a proportional allocation scheme would receive an allocation which is lower than the minimum lot size. This constraint cannot be violated by law, and therefore, all 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 and lottery applicants receive their proportional allotment in expectation.⁸

⁷Of course, total firm supply is restricted by the overall number of shares that the firm decides to issue, which is fixed prior to the commencement of the application process for the IPO.

⁸Appendix Section A shows a mathematical derivation of the probabilities of winning allotments based on the level of excess demand.

This third case, in which the lottery takes place, constitutes our endowment effect experiment. Far from being an unusual occurrence, in our sample alone (which does not even cover all IPOs in the Indian market over the sample period), roughly 1.5 million Indian investors participate in such lotteries over the 2007 to 2012 period in the set of 54 IPOs that we study.

The timeline of the application and allotment process is as follows. Applications are received over a two-day period termed the “subscription period.” Shares are allotted to the winner’s accounts approximately 12 days after the applications are received. The shares typically list approximately 21 days after the subscription period. Refunds of the escrow amounts begin to be processed after the allotments are made, usually 14 days after the allotments are made, so it is possible that the refunds are made after the shares are listed (we discuss how this might affect liquidity constraints later). All lottery losers receive a complete refund on their escrow amounts.

2.1 An Example: Barak Valley Cements IPO Allocation 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., $J = C$.

Appendix Table A.2 shows the official retail investor IPO allocation data for Barak Valley Cements.⁹ 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, the total number of share categories (C) equals 15,

⁹These data are obtained from
[urlhttp://www.chittorgarh.com/ipo/ipo_boa.asp?a=134](http://www.chittorgarh.com/ipo/ipo_boa.asp?a=134)

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. This maximum subscription amount would violate the prevailing (in 2007) regulatory maximum retail investor application constraint of Rs. 100,000 rupees per IPO. 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. 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 a 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 (2), and column (5). 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.¹⁰

As described briefly earlier, it is perhaps easiest to think of our data as comprising a large number of experiments, in which 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

¹⁰By regulation the firm allocates shares to investors rounded upwards to the nearest integer and will appropriately increase the total number of shares to accommodate the rounding off additions.

the causal effect of experiencing the IPO listing return on future investment behavior. We explain this more fully in the methodology section, following the data description below.

2.2 Data

When an individual investor applies to receive shares in an Indian IPO their application is routed through a registrar. In the event of heavy oversubscription leading to a randomized allotment of shares, the registrar will, in consultation with one of the stock exchanges, perform the randomization to determine which investors are allocated. We obtain data on the full set of applicants to 54 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 all IPOs 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,¹¹ 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.¹²

Our second major data source allows us to characterize the equity investing behavior of these IPO applicants (and in particular their trading decisions regarding the IPO stock). 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

¹¹Pin-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 they cover a larger geographical region in India.

¹²An 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. The alternative is paying by cheque, i.e., in either case, the money is placed in escrow prior to the allotment process, but in the case of ASBA, any refunds are processed a few days faster.

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 the IPO lottery only applies to this group of investors.

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 (or already have) an account with either of the two depositories (CDSL and NSDL), as the option to receive allocated shares in an IPO in physical form does not exist. For all applicants with accounts in CDSL, we observe accounts that applied for an IPO and were allotted in the lottery, i.e., the treatment group, as well as those that applied, but due to randomized allocation did not get allocated any share in an IPO. The latter group is the universe of counterfactuals in the IPO randomized lottery, i.e., the control group.

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.¹³ 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 in an attempt to increase the household's likelihood of treatment. While we do not have a direct way to control for this possibility, given our sample size, we do not believe that this is likely to affect our inferences materially.

Between March 2007 and March 2012, the common sample period for our total dataset, we observe 85 IPOs (of a total of roughly 240). Appendix Figure A.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. In our sample of 85 IPOs, 54 IPOs have at least one share category with a randomized lottery allocation, compared to the universe of 176 IPOs with randomized allocations over the period.¹⁴

Appendix Table A.1 presents summary statistics on the 54 IPOs with randomized allotments in our sample. The preponderance of IPOs in our sample, 31, are in the manufacturing sector, with 17 in the service sector, 4 in the technology sector, and 2 retail sector IPOs. The table shows that these IPOs account for 22% of all IPOs over this period by number, and US\$ 2.65 BN or roughly 8% of total IPO value over the period, varying from a low of 0.72% of total IPO capital in 2009 to a high of roughly 25% in 2011.

Between 32% and 35% of shares in these IPOs are allocated to retail investors who are

¹³In 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.

¹⁴We only consider IPOs that both undertake a randomized allocation and are mentioned in public sources such as www.chittorgarh.com in our analysis.

not employees of the IPO firm.¹⁵ The average IPO in our sample is 12 times oversubscribed, leading to an average of 8,691 treatment accounts and 20,248 control accounts per IPO, for a total of 1,562,706 accounts in our experiment. We observe a total of 383 randomized share categories (or experiments) across 54 IPOs, of which 323 randomized share categories experienced positive first-day listing gains in the stock market.

2.3 Using IPO Lotteries to Estimate Endowment Effects

We estimate endowment effects by comparing the tendency of lottery winners to continue to hold a stock that they were randomly assigned to the tendency of lottery losers to begin to hold a stock they randomly did not receive. It is important to clarify what exactly our lottery winners are endowed with relative to our lottery losers. Table 1 characterizes the application and allotment experience the investors in our analysis received upon being randomly chosen to receive IPO shares. Column (1) of the table shows the mean across all investors in the treatment groups or IPOs in our 323 share category experiments (for the positive first-day return IPOs) for each of the variables listed in the row headers.¹⁶ Columns (2) through (6) present the percentile of each variable in terms of the distribution across all of the experiments.¹⁷

On average, both our lottery winners and losers put 1,744 dollars into an escrow account to participate in the lottery (row 1, Table 1). Our lottery winners receive 150 dollars worth of the IPO stock from winning the IPO lottery (row 3). In addition, they also receive an instant gain of 62 dollars because on average the IPO stocks trade 42 percent higher when they list on the market relative to the issue price (row 5). Note that our lottery losers cannot purchase the stock at the issue price, so the total endowment that the winners receive, that

¹⁵This is slightly below the mandatory 35% allocation to retail investors because we do not include employees in this calculation as employees are not randomly assigned shares. For further details, refer to the online appendix.

¹⁶The weighting across the different share categories is done in exactly the same way as the regression framework we use weights the individual treatment and control groups.

¹⁷We first calculate the mean within each experiment, and then report the corresponding percentile across the experiments. For example, the median share category experiment had a mean application amount of 847 dollars (first row of Table 1).

the losers do not, is 212 dollars (150 +62) of the IPO stock. Once the stock has started trading, both winners and losers are free to trade to the stock. Also note that lottery winners and losers get a refund from their escrow accounts of approximately 1,690 and 1,743 dollars respectively.

Our approach has important similarities and differences with the two major laboratory methods of estimating endowment effects. The assignment of treatment is similar to the “valuation paradigm” methodology; experiments in the valuation paradigm randomly assign subjects to owning or not owning an object, and then survey the subjects on their valuation of that object. The key result is that sellers typically report higher valuations than buyers (the so called willingness to accept - willingness to pay gap).¹⁸ Our setting parallels the valuation paradigm method in the sense that winners of the IPO lotteries are randomly assigned to own the stock, whereas losers are randomly assigned to not own the stock. The randomization ensures that there should not be preference or belief based reasons why the winners should value the stock more than the losers. And similar to these valuation experiments, our winners receive a wealth effect when they are endowed with the IPO stock; in particular, our winners are allowed to purchase the IPO stock at the issue price and then sell it at the opening price; given that it is not possible for the lottery losers to purchase the stock at the issue price, this constitutes a wealth gain (or loss) that is not available to the lottery losers.¹⁹ This literature has typically argued that the wealth effect of assigning a consumer good (typically, a mug or a pen) are small relative to the subject’s wealth, and are therefore unlikely to drive their results. Our arguments regarding wealth effects (detailed later), are similar in that the gains on the lottery shares are most likely small relative to household wealth, and additionally, small relative to the escrow amount that households have made available to participate in the IPO allocation process in the first instance.

¹⁸For a review see Ericson and Fuster (2013), and for early papers using this method see Knetsch and Sinden (1984), Heberlein and Bishop (1986), and Kahneman, Knetsch and Thaler (1990).

¹⁹The wealth gain in our setting is not equal to the total amount of the endowment because our lottery losers receive a refund equal to the amount of the allotted stock (valued at the issue price). The differential wealth gain to winners versus losers is therefore just the gain in the value of the allotted stock based on the listing return of the stock.

A major difference between our design and experiments in the valuation paradigm is that we infer how subjects value the endowment based on their actual trading decisions, as opposed to surveying our investors regarding their valuation of the stock. If there is a large divergence in whether winners or losers hold the stock, this would suggest valuation differences across the two groups (assuming negligible wealth and transaction cost effects, which we investigate later). In this sense, our approach is similar to the “exchange paradigm” methodology (see, for example, Engelmann and Hollard (2010)), where subjects are randomly assigned good A or good B, and then later asked if they want to trade (after getting the opportunity to trade both groups should have the same proportion of good A and B owners - just as in our setting the winners and losers should have the same proportion of IPO stock holders). The extent to which the randomized groups holdings depart from equal proportions provides a quantitative estimate of the endowment effect.

3 Documenting the Endowment Effect

We begin by documenting the rate of ownership of the endowed IPO stock across our lottery losers and winners over time. We 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 only the randomized variation of winning 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 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 of winning the lottery stems solely from the random variation in treatment within each

experiment.²⁰ In particular, we estimate the causal effect of the experience of winning an IPO lottery on holdings of the IPO stock later by estimating the cross-sectional regression in each (event) month t :

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

Here, $y_{i,j,c,t}$ is an outcome variable of interest (for instance, an indicator for whether the account holds the IPO stock) for applicant i in IPO j , share category c , at event month t (we measure time in relation to the month 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 for IPO j in category c (investor is in the winner group), and 0 otherwise (investor is in the loser group). ρ_t are the estimated treatment effects in each event-month t . $\gamma_{j,c}$ 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. Conditional on the inclusion of these fixed effects, variation in treatment is random, meaning that the inclusion of controls should have no effect on our point estimates of ρ_t . Specification (3) identifies ρ_t as the causal impact of the experience of winning the IPO lottery on the outcome variable $y_{i,j,c,t}$.

Angrist (1998) shows that our estimated treatment effect ρ_t is a weighted average of the treatment effects from each separate share category experiment. In particular, the weights are constructed as:

$$w_c = \frac{r_c(1 - r_c)N_c}{\sum_{k=1}^{323} r_k(1 - r_k)N_k} \quad (2)$$

where r_c and N_c are the probability of treatment and sample sizes in share category c , and we have a total 323 share category experiments. Intuitively, the regression weights

²⁰See 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.

give more importance to experiments in which the probability of treatment is closer to $\frac{1}{2}$, and experiments with larger sample sizes. The basic idea is that the “good” experiments are ones in which there are many accounts in both treatment and control groups. This weighting scheme implies that our regression estimate only exploits purely random variation in treatment induced by the lotteries, since treatment versus control comparisons are only performed *within* share categories and given the fact that ρ_t is a weighted average of these share-category-specific effects.

Table 1 presents our main estimates. The columns correspond to the end of month portfolios after the IPO allotment. Each panel presents results from a different measure of the extent to which the account holds the IPO stock. Panel A uses an indicator for whether the account holds any of the IPO stock. We find that at the end of the first month the lottery winners are 55 percent more likely to hold the IPO stock than the lottery losers. This treatment effect declines to 45 percent at the end of six months. These differences are all significant at the 1 percent level. The numbers in brackets are standard errors and the numbers in parentheses are the lottery loser means of the dependent variable (using the weights in equation (7)). The loser group means show that it is relatively rare for lottery losers to own the stock; on average 1 percent of lottery losers own the IPO stock after the allotment, and six months after the allotment only 2 percent of losers own the stock.

Panel B defines the dependent variable as the fraction of the potential IPO allotment that the account holds. For example, if winners in a particular share category lottery won ten shares and a given account holds five shares, the dependent variable would be defined as .5. This definition of the dependent variable takes into account the number of shares an account holds (instead of just whether or not they hold the stock as in Panel A). Using this measure we see the treatment effect is slightly smaller at the end of month 1, but otherwise very similar to those in Panel A. However, comparing the lottery loser means across Panels A and B, we see that conditional on holding the IPO stock those in the lottery loser group choose to hold a substantially larger fraction than the lottery allotment. In particular, in Column (1) if

one percent of the lottery losers hold the stock at all, but their average fraction of allotment is five percent, this implies that lottery losers who choose to own the stock purchase five times the amount of lottery allotment.²¹

Panel C defines the dependent variable as an indicator for whether the account holds exactly the number of shares allotted to winners in the relevant share category. The results here are also very similar to those in Panel A, suggesting that most of the divergence in holding patterns between lottery winners and losers comes from the fact that the lottery winners continue to hold their exact allotments, while the losers are unlikely to purchase the exact allotment they would have received had they won in the lottery.

Panel D uses the value (in USD) of the IPO stock held in the portfolio at the end of the month. Lottery winners hold 87 dollars more of the stock on average at the end of the first month, 73 dollars more at the end of the second month, and 45 dollars at the end of the sixth month. This measure incorporates both the difference in the number of shares the winners and losers choose to hold as well as the returns earned on those shares. Some of the decline in value is due to the negative returns earned on these IPO stocks (the six month cumulative return is negative 10 percent). Panel E uses the weight of the portfolio in the IPO stock as the dependent variable; lottery winners own 10 percent more of their portfolio in the IPO stock, and this remains a substantial 6 percent more six months after allotment.

The final panel of the table calculates the average returns on holding the stock. The weights used to calculate these average returns are the same as those in equation (7), meaning that these return estimates combine information across our experiments in the same way as in our estimated treatment effects. On average the listing return (i.e., the percentage return the lottery winners gain as soon as the stock starts trading) is 40.9 percent. The next row shows the cumulative return a player earns by holding the stock to the end of the month given in the column. The last row gives the cumulative return the player earns relative to

²¹Suppose there are 10,000 lottery losers, the lottery allotment (to winners) was 10 shares, and 100 losers purchase the stock (1 percent). Also suppose that those 100 losers choose to purchase 50 shares. Then, the average fraction of the allotment held by lottery losers will be 5 percent ($.01*5+.99*0 = .05$).

the first price the stock traded at (the open price).

The returns data show that lottery winners on average lost money based on their choice to continue to hold the stock after it was initially listed. The returns relative to the open price are large and negative. So, in this sense, the lottery losers in our sample make a relatively good decision not to purchase these IPO stocks (on average) at the open price. We note, however, that this outcome depends on the average returns in our particular sample. Regardless of whether the lottery winners or losers made the right decision in our particular sample, what is important about these results is that these two groups choose to make substantially *different* decisions about whether they want to hold this stock. For example, if this pattern of negative post-issue returns is predictable, then we would expect *both* lottery winners and losers to choose not to hold the stock after listing; it would not be very difficult for lottery winners to have a decision rule where they always sell the stock after it lists. Instead, however, we see a strong divergence in behavior across the groups.

Table 4 extends the analysis to 24 months after the allotment of the lottery stock to winners. The sample here only includes IPOs that occurred in or before March 2010, as our data on portfolio holdings ends in March 2012. The results for this sample are similar to the results for the full sample in Table 3 in month 1, and the trend is similar through month 8. We find that even 24 months after allotment the lottery winners are 35 percent more likely to hold the IPO stock than the lottery losers. The lottery losers' propensity to hold the stock stays relatively constant at around 1.5 to 1.7 percent over these 24 months.

3.1 Explanations Involving Causal Effects of Endowments

1. Reference Dependent Loss Averse Preferences. While a number of potential causes of endowment effects have been proposed, the most prominent is based on the idea of reference dependent preferences with loss aversion as first proposed in Kahneman and Tversky (1979). In our setting, a prediction of the reference dependent utility with loss aversion explanation is that winners of the lottery should have a tendency to hold the stock

more if they experience a loss on the stock; this result is similar to the large literature documenting that investors are more likely to hold on to their gains than their losses (Shefrin and Statman, 1985; Odean, 1998). This theory, however, does not have a prediction for whether or not lottery losers should buy the stock, as it is unclear what reference point losers would judge their potential outcomes relative to.

Figure 1 presents evidence on the relationship between the propensity of lottery winners and losers to hold the IPO stock and the returns on the IPO stock in the first month. Each dot represents a different share category (i.e., experiment) in our sample. The blue dots show the fraction of the initial allotment that lottery winners in that experiment held at the end of the first month, and the red crosses show the corresponding fraction of initial allotment that the lottery losers held at the end of the first month. The x-axis is the percentage return on the stock (from the issue price) to the end of the first month. The figure shows a strong negative correlation between lottery winners tendency to hold the stock and the returns earned to the end of the first month. The figure also shows a slightly increased tendency for lottery losers to purchase the stock when the IPO stock's returns are lower.

The fact that the divergence is largest for IPOs that had a negative return at the end of the month is consistent with the reference dependent utility loss aversion explanation of the behavior of lottery winners. What is striking is that lottery losers, who, because of the randomization are very similar in terms of background characteristics, choose to hold very low levels of the stock. It is also interesting to note that the divergence is also positive even when the IPO stock is held at a gain (i.e., the blue dots are higher than the red dots for IPOs that had positive returns between 0 and 100 percent after a month).

While the reference dependent utility with loss aversion appears to be a good explanation for the tendency of our lottery winners to hold the stock, these factors are not useful for explaining why lottery losers purchase the stock at such a low rate. The simplest explanation is that lottery losers have rational expectations, in the sense that on average purchasing the IPO stocks in our sample after they start trading leads to substantial negative returns, and so

our lottery losers choose not to purchase the stock on the open market. Under this hypothesis, however, we would expect that lottery winners would sell the stock immediately after listing; given that the lottery winners and losers are chosen randomly, it is hard to understand why lottery losers would have rational expectations about the future performance of these IPO stocks but lottery winners would not.

2. Aversion to Bad Deals. Another plausible candidate explanation is that lottery losers have an “aversion to bad deals” as described in Weaver and Frederick (2012). These authors argue that endowment effects can often be explained not because endowments change valuations, but instead because endowments create reference prices, and some agents get utility directly from how the price they pay relates to a reference price. In our setting a natural reference price is the issue price of the IPO (i.e., the price at which lottery winners purchase the stock). Lottery losers might see purchasing the stock after the IPO as a “bad” deal because the stock typically trades higher than the issue price (even though the issue price is irrelevant for the future performance of the stock).

3. Quasi-Endowment Effects. In other settings, authors have argued that prospective buyers may get attached to objects before actually receiving the object through the process of applying or bidding for the object, and this attachment may cause the subject to eventually value the endowed object more; the attachment generated during the process of acquisition of an endowment object have been termed “quasi-endowment effects” (Heyman et al., 2004; Wolf et al., 2005) For example, Malmendier and Lee (2011) find that ebay auction bidders often bid more in an auction for an item than the fixed price the item is available for sale; one possible explanation is that the bidders become more attached to the item because they spend more time bidding in the auction. In our setting, both the lottery winners and losers went through the same application process, so the time spent considering the IPO, filing the application, and so on is naturally controlled for. However, lottery winners are credited the shares in their account a few days before the shares are actually listed. It is perhaps less clear why having the shares credited in the account a few days before listing would make

winners more attached to the shares. Nevertheless, we intend in future drafts to test if the endowment effect is stronger as greater time elapses between the allotment of shares and trading.

3.2 Explanations Where Endowments Do Not Change Valuations

1. Wealth Effects. Lottery winners in our sample get, on average, a wealth gain of approximately 67 dollars because IPOs begin trading at a price which 42 percent on average higher than the issue price. To what extent are our results driven by this wealth effect? Would we expect similar results if our lottery winners were given an extra 67 dollars in cash?

This pure wealth effects interpretation of our results is unlikely for the following reasons. First, both lottery winners and losers are required to deposit a much larger amount of cash to participate in these lotteries in the first place. Table 1 shows that, on average, in our sample our winners and losers both had to deposit 1800 dollars to participate in the lottery. Furthermore, both lottery winners and losers receive most of this cash back because the allotments are ultimately much smaller than the amount of shares applied for (due to the high over-subscription rates). Thus, it seems unlikely that the extra 67 dollars our winners have somehow causes them to be much more likely to hold the IPO stock as we know both lottery winners and losers have substantial cash on hand.

Second, Figure 2 shows that our estimated endowment effects do not vary strongly with the size of the wealth effect, which we measure as the dollar return the winners earn because the stock's opening trading price is higher than the issuing price. The figure shows that both the fraction of the allotment winners and losers hold does not vary much with the size of the wealth gained by lottery winners, again suggesting that the behavior is not driven by the wealth gains alone.

2. Transaction Costs. One possible explanation for the divergence we find between lottery winners and losers holding the IPO stock is that transaction costs make it unprofitable for lottery winners to sell the stock, and simultaneously make it unprofitable for lottery losers

to buy the stock. Note that under this explanation, both winners and losers have the same optimal holding levels, but the cost of getting to that optimal holding level outweighs the benefits of arriving at the optimal holding level. Transactions costs could include both purely monetary costs such as brokerage commissions, as well as non-monetary costs such as the time it takes to place a trade or just a general cost of taking any action (i.e., inertia, which we discuss below).

In terms of monetary transaction costs, there are two primary types of costs to consider: (1) brokerage commissions, and (2) securities transactions taxes.²² Our data does not include information on brokerage commissions costs, and we are not aware of any representative datasets on commissions for Indian equity accounts. However both the Bombay and National Stock Exchanges specify that brokers may not charge more than 2.5 percent of the valuation of a transaction as a brokerage fee. In our sample the average IPO allotment is worth 150 USD, so the commissions to buy or sell the full allotment are on average less than 3.75 USD. In reality commissions are typically much lower than the statutory maximums because of competition amongst brokers. We hand collected brokerage commissions from twenty major retail brokerage firms over our sample period (2007-2012) and found the commissions to vary between .3 to .9 percent of the transaction value, much less than the statutory maximum of 2.5 percent. Securities transaction taxes are an additional 14.5 basis points (Mohanty, 2011). Given these estimates it seems unlikely that monetary transactions costs would cause such a large divergence between the holdings of lottery winners and losers of the IPO stock.

3. Inertia. The non-monetary costs of taking action in the brokerage account could be potentially more important. For example, consider the 23 percent of lottery applicants who opened a new account to apply for the the IPO stock. These applicants might have less familiarity with the procedures of making a buy or sell order, and therefore may have

²²In addition to the direct securities transaction tax (12.5 basis points paid to central government during our sample period) there are three additional taxes charged at the time of transaction: a service tax on brokerage (10.3 percent of the brokerage commission paid to central government), a stamp duty (1 basis point of transaction value paid to state government), and a SEBI turnover fee (1 basis point of transaction value paid to stock market regulator).

to pay reasonably sized time costs to execute a trade. Further, there might just be inertia associated with making a trade on the IPO stock given the allocation is relatively small, especially given that for lottery winners holding the IPO stock is the status quo (Samuelson and Zeckhauser, 1988). Three empirical results make us hesitant to conclude that simple inertia is driving the divergence between lottery winners and losers in holding the IPO stock.

First, in related work we find that lottery winners are more likely to trade non-IPO stocks in their portfolio than lottery losers, suggesting that, if anything, winning the lottery reduces the inertia associated with making a trade (Anagol et al., 2015). It is unclear why winning the lottery would increase the probability of trading in non-IPO stocks, but simultaneously cause the winners to hold the IPO stock even though they would really prefer to sell it.

Second, under the inertia explanation, we would expect the divergence between lottery winners and losers to be largest for accounts that have made relatively few trades in past, and smaller for those that have made a large number of trades. Accounts that have made many recent trades are more likely to be able to overcome the inertia associated with the endowment. We do find that accounts that made more trades in the month prior to the lottery assignment do have a lower estimated divergence between winners and losers, but the relationship is weak; an increase of 10 trades made in the past month only reduces the estimated holding divergence between winners and losers by 2 percent (see next Section and Table 5.) These results suggest that even accounts that have been very recently active show large divergences between buyers and sellers.

Third, we find a strong relationship between the performance of the IPO stock and the divergence between lottery winner and loser holdings. If the results were mainly due to inertia, it is not clear why inertia would be so much more important when the lottery stock performed poorly versus well.

4. Flipping Incentives. In the United States, IPO shares are typically rationed to brokerage clients who have provided large value to the brokerage firm. There is substantial anecdotal and empirical evidence to suggest that brokerages discourage investors from quickly

selling their allotted shares, in particular by threatening that “flippers” will be denied future IPO allocations (Aggarwal, 2003). Thus, in the United States, it is possible that allottees of IPOs choose to hold the stock much longer than statistically similar non-allottees because they believe selling the stock will reduce their chances of being allotted future IPOs. A few factors make this explanation less plausible in the Indian setting. First, we show that those who sold their shares quickly in the past are equally likely to be allotted in the future (conditional on applying); this is another form of balance check that shows the allotments are truly random. Second, the lottery process is publicly advertised after allocations are made (i.e. the fraction of allottees randomly chosen appears in newspaper articles etc.), and it is generally common knowledge that winners are chosen at random; therefore, it is not clear why investors would assume that selling their shares quickly would hurt them in future allocations.

5. Tax Motivated Behavior. In India short term (less than 1 year) losses on stocks can be applied to short term gains on stocks to reduce capital gains tax liability.²³ Constantinides (1984) notes that under these types of tax incentives, and the presence of transactions costs, investors should slowly realize their losses with the volume of sales peaking right before the end of the fiscal year. This might give lottery winners an incentive to generally hold their shares that have experienced losses until the end of the Indian fiscal year (March 31), and then sell them soon after. While we intend to investigate this issue in greater detail, we simply note here that our finding that the endowment effect is strong 24 months following the initial allocation suggests that this cannot be a substantial part of the explanation, as this is a far longer horizon than the expected duration of any tax motivated trading deferral.

²³During the period of our study short term capital gains were taxed at 15 percent. There was no long term (greater than one year) capital gains tax, and therefore no opportunity for long term tax loss offsets.

4 Experience

4.1 Correlates of Endowment Effect Estimates

In this Section we study how the differential holding patterns of the IPO stock vary with plausible proxies of the investor's experience in the IPO market. To do this we interact our main treatment variable in equation (1) with a set of pre-determined variables that we believe are interesting proxies for the amount of experience in the IPO market. The purpose here is not to identify a causal relationship between these experience proxies and our estimated endowment effects, but to provide descriptive evidence on how the endowment effect estimates vary with various investor characteristics (similar to the analysis in List et al. (2003)). Table 5 presents the results.

Column (1) replicates our main finding that at the end of month 1, lottery winners own 54 percent more of the allotment than lottery losers. Column (2) interacts this treatment effect with a variable measuring the number of IPOs the account had been allotted in the month prior to the treatment IPO.²⁴ We find that having been allotted one additional IPO in the past reduces the estimated endowment effect by .7 percent (statistically significant at the one percent level). We also control directly for the number of IPOs having been allotted to the account in the past. The coefficient on this variable measures the correlation between how much holdings of the IPO stock amongst lottery losers increases with the number of IPOs they have been allotted in the past; the estimated coefficient is positive and statistically significant, suggesting that more experienced losers are (slightly) more likely to purchase the IPO stock in after it begins trading.

Columns (3) and (4) add two measures of portfolio size, the number of securities held and the value of the portfolio. Lottery winners with larger number of securities and larger portfolio sizes are more likely to hold the IPO. Both size variables are economically small; an additional security held in the portfolio increases the treatment effect by 1/10 of one

²⁴We define all independent variables in this regression in the month prior to the lottery allotment to avoid mechanical correlations between the dependent and independent variables due to the lottery allotment.

percent, and a 10 percent increase in portfolio size increases the treatment effect by 3/10 of one percent. These results are consistent with our previous evidence that wealth effects are unlikely to be driving the endowment effects in our setting (we would expect the effects to get smaller as wealth increases if wealth effects were driving the result).

Column (5) adds a variable measuring the months since the account was opened. The treatment effect increases by 1/10 of a percent for each additional month the account has been active. The direct effect of months since opened is small and not statistically significant at standard levels. Column (6) adds an interaction between the first day return of the IPO and the treatment effect. Consistent with the graphical evidence in Figure 1, a ten percent increase in first day return is correlated with a two percent decrease in the treatment effect. Column (7) adds the number of trades the account made in the month prior to allocation of the IPO, and an interaction of that variable with whether the account was a lottery winner. There is a negative and significant relationship between number of trades made in the past month, but the coefficients suggest that even for large numbers of trades the difference between winners and losers persists.

Column (8) replaces the linear measure of the number of IPOs allotted in the past with a dummy variable for accounts that have been allotted over thirty IPOs. The idea here is to see whether accounts with very large allotments in the past are particularly different in their holdings (i.e., potential non-linear effects in the number of previous allotments at the very high end). Accounts with over 30 allotment experiences have estimated endowment effects that are 20.7 percentage points smaller than the full sample (a reduction of 40 percent). On the other hand, it is interesting to note that even for these accounts the lottery winners are 27.6 percent more likely to hold the IPO stock at the end of the first month.

Overall, the results suggest that there is a correlation between accounts having greater experience (as proxied by number of IPOs allotted in the past) and smaller endowment effects, consistent with the findings in List et al. (2003). However, having experienced many allotments in the past does not appear to quickly lead to complete elimination of the

endowment anomaly in this setting.

4.2 Using Random Assignment to Estimate Experience Effects

The above analysis focused on testing whether the divergence between lottery winners and losers varies based on the total number of IPO lotteries experienced. The advantage of that approach is we have a more holistic measure of the investors total experience in the market. However, it is also possible to estimate the shorter run relationship between experience and endowment effects by testing whether winners in a recent, previous, lottery, show lower endowment effects in a current lottery.

Table 6 presents the results of ten such comparisons. We focus on the 10 pairs of lotteries in our data where there were the largest number of applicants that applied to both lotteries within the pair. For example, the first row analyzes the behavior of the 156,120 applicants who applied to both BGR Energy Systems and Future Capital IPOs. We term the first IPO as “IPO a ” (BGR in this case) and the second IPO as IPO b (Future Capital in this case). BGR Energy listed on January 3, 2008 and had a listing return of 66.9 percent. However, after listing BGR had a 29.9 percent loss up until the date that Future Capital listed (February 1, 2008). We are interested in whether the allotted BGR applicants show smaller endowment effects in their behavior regarding Future Capital.

To estimate whether BGR winners show a smaller endowment effect in decisions regarding Future Capital we estimate the following regression model, where the sample only includes accounts that applied to both BGR and Future Capital:

$$y_{i,c_a,c_b} = \alpha + \beta_1 Win_b_{i,c_a,c_b} + \beta_2 Win_b_{i,c_a,c_b} * Win_a_{i,c_a,c_b} + \beta_3 Win_a_{i,c_a,c_b} + \gamma_{c_a,c_b} + \epsilon_{i,c_a,c_b} \quad (3)$$

y_{i,c_a,c_b} is an indicator for whether account i in share category c_a of IPO a and in share category c_b of IPO b holds the IPO b stock at the end of the first month after IPO b was listed

(i.e. at the end of February 2008 in the case of the BGR/Future Capital pair represented in the first row. Note that a given account can only appear in exactly one share category in IPO a and one share category in IPO b because an account can only apply once to a given IPO. Win_{b,i,c_a,c_b} and Win_{a,i,c_a,c_b} are indicators for whether account i was allocated in IPO b and IPO a respectively. γ_{c_a,c_b} are fixed effects for each possible pair of share category combinations across IPOs a and b . We include this fixed effects to control for any factors that are common to people who chose to apply to given share categories in IPOs a and b .

We are primarily interested in the coefficient β_2 , which tells us the difference in the estimated endowment effect in IPO b based on whether the account was allocated in IPO a . Column (8) of Table 6 reports β_2 for the ten largest pairs of IPOs in terms of the number of applicants who applied to both. We would expect β_2 to typically be negative, because observing the performance of the IPO stock after listing should cause greater convergence in the behavior of winners and losers in the next IPO.²⁵ Consistent with this, we estimate negative coefficients in nine of the ten examples studied here. On the other hand, the estimated coefficients are small, suggesting that an account would require a large number of these experiences before the endowment effect was eliminated (similar to our conclusion in the previous analysis).

It is important to note that there are two potential mechanisms underlying our negative estimates of β_2 . The first is that winning shares in IPO a causes a given account to exhibit the endowment effect less in a future IPO (i.e. a causal effect of experience). The second is that the types of players who choose to apply to IPO b after winning shares in IPO a are differentially selected to be the type who have lower endowment effects (i.e. a selection effect). Previous studies, such as List (2011), focus on separating these two effects, but this is difficult in our setting as the choice to apply for a future IPO is endogenous.

However, we argue that in this particular analysis the joint effect is a primary object

²⁵For example, lottery winners who experience a negative open return should sell future allotments faster, thus reducing the convergence. Similarly, lottery losers who observe the IPO stock having a positive listing return should be more likely to purchase the stock on the open market.

of interest; it tells us whether the two forces of investors learning from experience as well as the force of experiences driving some investors out of the market, lead to lower market anomalies (such as the endowment effect) over time. If winning previous lotteries makes an account more likely to apply (which we show in Anagol et al. (2015)), then these results would suggest that there will be a modest reduction in endowment effects under the selection mechanism as well. For example, suppose the whole difference in behavior of past winners and losers in future IPOs is due to selection; this would mean that winning past lotteries induces a selection of investors who exhibit lower anomalies in the future.

Overall, our analysis of the relationship between experiences and estimated endowment effects suggest that (1) investors with greater experience due exhibit lower endowment effects (due to either a causal effect of experience or selection), and (2) that it is unlikely that experience quickly and absolutely eliminates the endowment anomaly.

5 Conclusion

In the absence of important wealth effects or transactions costs, standard economic theories predict a fundamental symmetry: the same person should not make different decisions about whether to hold an asset depending on whether he is endowed with the asset. Data on the behavior of applicants to Indian IPO lotteries refutes this prediction. We find that randomly receiving shares in an IPO increases the probability that an applicant holds these shares for many months after the allotment, and that standard factors such as wealth effects and transactions costs are unlikely to explain these effects.

References

- Aggarwal, Reena**, “Allocation of initial public offerings and flipping activity,” *Journal of Financial Economics*, 2003, 68 (1), 111–135.
- Anagol, Santosh and Hugh Hoikwang Kim**, “The Impact of Shrouded Fees: Evidence from a Natural Experiment in the Indian Mutual Funds Market,” *American Economic Review*, 2012, 102 (1), 576–93.
- , **Vimal Balasubramaniam, and Tarun Ramadorai**, “The Effects of Experience on Investor Behavior: Evidence from India’s IPO Lotteries,” *Available at SSRN 2568748*, 2015.
- Angrist, Joshua D and Jörn-Steffen Pischke**, *Mostly harmless econometrics: An empiricist’s companion*, Princeton university press, 2008.
- Angrist, Joshua D., Parag A. Pathak, and Christopher R. Walters**, “Explaining Charter School Effectiveness,” *American Economic Journal: Applied Economics*, 2013, 5 (4), 1–27.
- Black, Dan A, Jeffrey A Smith, Mark C Berger, and Brett J Noel**, “Is the threat of reemployment services more effective than the services themselves? Evidence from random assignment in the UI system,” *American Economic Review*, 2003, pp. 1313–1327.
- Campbell, John Y., Tarun Ramadorai, and Benjamin Ranish**, “Getting Better or Feeling Better? How Equity Investors Respond to Investment Experience,” Working Paper 20000, National Bureau of Economic Research March 2014.
- Campbell, John Y, Tarun Ramadorai, and Benjamin Ranish**, “The Impact of Regulation on Mortgage Risk: Evidence from India,” *American Economic Journal: Economic Policy*, *forthcoming*, 2015.
- Constantinides, George M**, “Optimal stock trading with personal taxes: Implications for prices and the abnormal January returns,” *Journal of Financial Economics*, 1984, 13 (1), 65–89.
- Ericson, Keith M Marzilli and Andreas Fuster**, “Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments,” *The Quarterly Journal of Economics*, 2011, 126 (4), 1879–1907.
- and – , “The endowment effect,” Technical Report, National Bureau of Economic Research 2013.
- Haigh, Michael S and John A List**, “Do professional traders exhibit myopic loss aversion? An experimental analysis,” *The Journal of Finance*, 2005, 60 (1), 523–534.
- Heffetz, Ori and John A List**, “Is the endowment effect an expectations effect?,” *Journal of the European Economic Association*, 2014, 12 (5), 1396–1422.

- Heyman, James E, Yesim Orhun, and Dan Ariely**, “Auction fever: The effect of opponents and quasi-endowment on product valuations,” *Journal of Interactive Marketing*, 2004, 18 (4), 7–21.
- Kahneman, Daniel and Amos Tversky**, “Prospect theory: An analysis of decision under risk,” *Econometrica: Journal of the Econometric Society*, 1979, pp. 263–291.
- Knetsch, Jack L and Wei-Kang Wong**, “The endowment effect and the reference state: Evidence and manipulations,” *Journal of Economic Behavior & Organization*, 2009, 71 (2), 407–413.
- Kőszegi, Botond and Matthew Rabin**, “A model of reference-dependent preferences,” *The Quarterly Journal of Economics*, 2006, pp. 1133–1165.
- List, John A**, “Does market experience eliminate market anomalies? The case of exogenous market experience,” Technical Report, National Bureau of Economic Research 2011.
- **et al.**, “Does market experience eliminate market anomalies?,” *QUARTERLY JOURNAL OF ECONOMICS-CAMBRIDGE MASSACHUSETTS-*, 2003, 118 (1), 41–72.
- Malmendier, Ulrike and Young Han Lee**, “The bidder’s curse,” *The American Economic Review*, 2011, pp. 749–787.
- Mohanty, Nirmal**, “Cost of Trading in Stock Exchanges: A Perspective,” *National Stock Exchange (NSE) Working Paper*, 2011.
- Odean, Terrance**, “Are investors reluctant to realize their losses?,” *The Journal of finance*, 1998, 53 (5), 1775–1798.
- Samuelson, William and Richard Zeckhauser**, “Status quo bias in decision making,” *Journal of risk and uncertainty*, 1988, 1 (1), 7–59.
- Shefrin, Hersh and Meir Statman**, “The disposition to sell winners too early and ride losers too long: Theory and evidence,” *journal of Finance*, 1985, pp. 777–790.
- Weaver, Ray and Shane Frederick**, “A reference price theory of the endowment effect,” *Journal of Marketing Research*, 2012, 49 (5), 696–707.
- Wolf, James R, Hal R Arkes, and Waleed A Muhanna**, “Is Overbidding in Online Auctions the Result of a Pseudo-Endowment Effect?,” *Available at SSRN 735464*, 2005.
- Zeiler, Kathryn and Charles R Plott**, “The willingness to pay/willingness to accept gap, the endowment effect, subject misconceptions and experimental procedures for eliciting valuations,” *American Economic Review*, 2004.

Table 1: CHARACTERIZING LOTTERY APPLICATION AND ALLOTMENT EXPERIENCE

Treatment Characteristics	Mean	Percentile Across Experiments				
		10	20	50	75	90
	(1)	(2)	(3)	(4)	(5)	(6)
Application Amount (\$)	1743.89	155.79	355.55	810.05	1402.18	2104.69
Probability of Treatment	0.35	0.09	0.20	0.37	0.64	0.84
Allotment Value (\$)	150.42	123.81	134.11	144.66	160.22	169.39
First Day Gain/Loss (%)	39.18	-7.56	6.104	17.13	37.08	87.77
First Day Gain/Loss (\$)	62.18	-11.27	8.68	24.94	53.59	143.49
Median Portfolio Value ($t - 1$, \$)	1868.27	810.40	1131.39	1635.56	2460.89	3217.51

Table 2: RANDOMIZATION CHECK

	Treatment	Control	Difference	% Experiments > 10% significance
	(1)	(2)	(3)	(4)
Applied/Allotted an IPO	0.379	0.379	0.000	8.97
Cutoff bid	0.928	0.928	0.000**	11.1
Application by Blocked Amount (ASBA)	0.048	0.048	-0.001	4.0
States of India				
Gujarat	0.350	0.348	0.002*	12.60
Maharashtra	0.216	0.216	0.000	9.30
Rajasthan	0.148	0.149	-0.001	8.70
Delhi	0.045	0.046	0.000	8.0
Portfolio Value > 0	0.778	0.778	0.000	9.28
IHS Portfolio Value	6.575	6.573	0.002	13.00
0	0.222	0.221	0.000	10.52
0 to 500\$	0.143	0.143	-0.001	8.66
500 to 1000\$	0.097	0.097	0.000	8.63
1000 to 5000\$	0.285	0.285	0.000	9.59
> 5000 \$	0.252	0.252	-0.001	10.21
IHS Gross Transaction Value	5.619	5.616	0.003	11.45
0	0.287	0.288	-0.001	8.97
0 to 500\$	0.183	0.183	-0.001	9.90
500 to 1000\$	0.127	0.127	0.000	9.59
1000 to 5000\$	0.287	0.285	0.002**	14.55
> 5000 \$	0.116	0.117	-0.001*	8.97
IHS Account Age	3.134	3.131	0.003*	12.38
New Account	0.041	0.041	0.000	5.00
1 Month old	0.067	0.066	0.001	6.80
2-6 Months old	0.201	0.203	-0.002**	8.35
7-13 Months old	0.143	0.143	0.000	9.00
14-25 Months old	0.151	0.152	-0.001	9.90
> 25 Months old	0.378	0.376	0.002	11.10

Table 3: EFFECT OF WINNING IPO LOTTERY ON OWNERSHIP OF IPO STOCK

Dependent Variable:	Month After Allotment					
	1	2	3	4	5	6
I(Holds IPO Stock)	0.55 [0.001] (0.014)	0.51 [0.001] (0.015)	0.49 [0.001] (0.015)	0.47 [0.001] (0.015)	0.46 [0.001] (0.016)	0.45 [0.001] (0.016)
Fraction of Allotment	0.53 [0.006] (0.055)	0.51 [0.006] (0.061)	0.49 [0.006] (0.060)	0.47 [0.006] (0.063)	0.47 [0.006] (0.062)	0.46 [0.006] (0.065)
I(Holds Exactly IPO Allotment)	0.52 [0.001] (0.002)	0.48 [0.001] (0.002)	0.46 [0.001] (0.002)	0.44 [0.001] (0.002)	0.43 [0.001] (0.002)	0.42 [0.001] (0.002)
Value of IPO Shares Held (U.S.D)	87.39 [1.016] (8.622)	73.30 [0.860] (8.425)	64.43 [0.864] (7.740)	61.64 [0.966] (8.310)	48.63 [0.919] (7.140)	44.99 [0.924] (7.100)
Portfolio Weight of IPO Stock	0.10 [0.000] (0.001)	0.08 [0.000] (0.001)	0.07 [0.000] (0.001)	0.07 [0.000] (0.001)	0.07 [0.000] (0.001)	0.06 [0.000] (0.001)
Mean Listing Return	40.9					
Mean Return Over Issue Price	13	2	-1	2	-12	-10
Mean Return Over Open Price	-19	-27	-28	-26	-36	-36

The sample size is 1,561,497 accounts in each month. Standard errors in brackets and the mean of the dependent variable for lottery losers is in parentheses. All treatment effects are significant at the 1 percent level.

Table 4: LONG RUN EFFECT OF WINNING IPO LOTTERY ON OWNERSHIP OF IPO STOCK

Dependent Variable:	Month After Allotment								
	1	4	8	11	12	13	16	20	24
I(Holds IPO Stock)	0.574	0.494	0.459	0.442	0.435	0.428	0.410	0.382	0.349
	[0.0006]	[0.0006]	[0.0006]	[0.0006]	[0.0006]	[0.0006]	[0.0006]	[0.0006]	[0.0006]
	(0.013)	(0.014)	(0.016)	(0.017)	(0.017)	(0.017)	(0.017)	(0.016)	(0.015)

The sample includes all IPOs that occurred 24 months before the end of our portfolio data in March 2012. The sample size is 1,090,346 accounts in each month. Standard errors in brackets and the mean of the dependent variable for lottery losers is in parentheses. All treatment effects are significant at the 1 percent level.

Table 5: HETEROGENEOUS WINNER EFFECTS BY PRE-EXISTING ACCOUNT CHARACTERISTIC

	1	2	3	4	5	6	7	8	9
Winner	0.552*** [0.008]	0.605*** [0.009]	0.591*** [0.010]	0.449*** [0.015]	0.445*** [0.015]	0.544*** [0.017]	.445*** [.0126]	0.562*** [0.008]	0.541*** [0.017]
Winner*# of IPOs Allotted (t-1)		-0.007*** [0.001]	-0.009*** [0.001]	-0.010*** [0.001]	-0.011*** [0.001]	-0.011*** [0.001]	-.010*** [.001]		
# IPOs Allotted (t-1)		0.001** [0.001]	0.001** [0.001]	0.001 [0.001]	0.001 [0.001]	0.001 [0.001]	.000 [.000]		
Winner*Number Securities (t-1)			0.002*** [0.000]	0.001** [0.000]	0.001* [0.000]	0.001* [0.000]	.002*** [.000]		0.000 [0.000]
Number Securities (t-1)			0.002*** [0.000]	0.001*** [0.000]	0.001*** [0.000]	0.002*** [0.000]	.000 [.001]		0.002 [0.000]
Winner*IHS(Portfolio Value (t-1))				0.026*** [0.002]	0.024*** [0.002]	0.025*** [0.000]	.019 [.002]		0.023*** [0.002]
IHS(Portfolio Value (t-1))				0.004*** [0.001]	0.004*** [0.001]	0.004*** [0.001]	.005 [.001]		0.004*** [0.001]
Winner*Months Since Opened (t-1)					0.001** [0.000]	0.001** [0.000]	.000 [.000]		-0.001*** [0.000]
Account Months Since Opened (t-1)					0.000 [0.000]	0.000** [0.000]	.000 [.000]		0.000 [0.000]
Winner*IPO First Day Return (%)						-0.002*** [0.000]	-.001*** [.001]		-0.002*** [0.000]
Winner*# Trades (t-1)							-.002*** [.000]		-.002*** [.000]
# Trades (t-1)							.002 [.000]		.002*** [.000]
Winner*(> 30 IPOs Allotted (t-1))								-0.207*** [0.034]	-0.221*** [0.036]
> 30 IPOs Allotted (t-1)								0.076*** [0.025]	0.039 [0.026]

The dependent variable is the fraction of the winning allotment held. The sample size is 1,561,497 accounts. All specifications include IPO Share Category fixed effects.

Table 6: EFFECT OF WINNING PREVIOUS LOTTERIES ON PROPENSITY TO HOLD FUTURE IPO ALLOCATIONS

IPO A		IPO B			Differential		
Name	Listing Date	Listing Return (%)	Open Return (%)	Name	Listing Date	Observations	Winner Effect
BGR	1/3/2008	66.88	-29.86	Future Capital	2/1/2008	156120	-0.024*** [0.003]
Career Point	10/6/2010	48.71	-15.92	P&S Bank	12/30/2010	34488	-0.026*** [0.009]
Omaxe	8/9/2007	29.03	-27.54	BGR	1/3/2008	34574	-0.010 [0.006]
Vishal Retail	7/4/2007	75.01	187.50	BGR	1/3/2008	49150	-0.022*** [0.007]
Omaxe	8/9/2007	29.03	-35.48	Future Capital	2/1/2008	34418	0.011* [0.006]
Vishal Retail	7/4/2007	75.01	187.50	Future Capital	2/1/2008	46816	-0.025*** [0.007]
Meghmani	6/28/2007	75.00	-14.44	BGR	1/3/2008	29304	-0.080*** [0.008]
Omnitech	8/14/2007	75.00	-3.67	BGR	1/3/2008	29276	-0.038*** [0.010]
BGR	1/3/2008	66.88	-10.32	P&S Bank	12/30/2010	48469	-0.001 [0.007]
Future Capital	2/1/2008	36.47	-81.82	P&S Bank	12/30/2010	54337	-0.008

The dependent variable is the fraction of the winning allotment held. The full sample includes all applicants for 196 shares in BGR that also applied for 128 shares in Future Capital. Standard errors in brackets and mean of the dependent variable for lottery losers in the parentheses.

Figure 1: Percent Allotment Held Versus Holding Returns at End of Month One

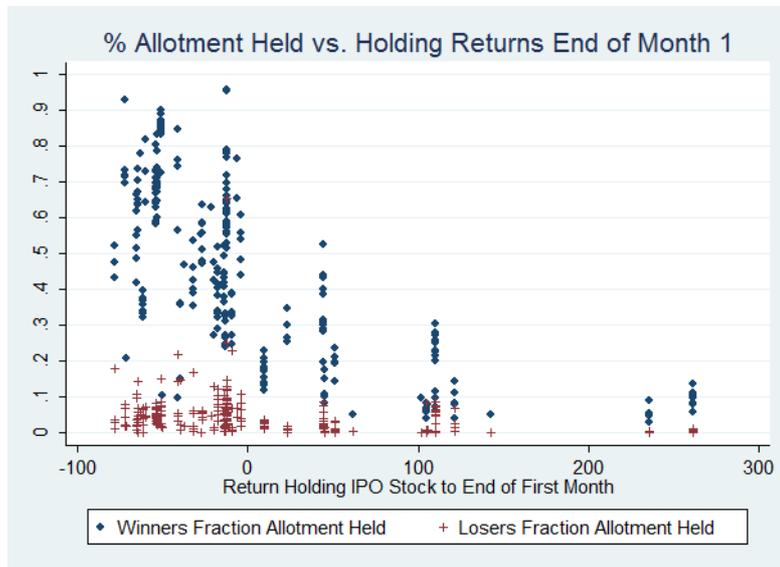
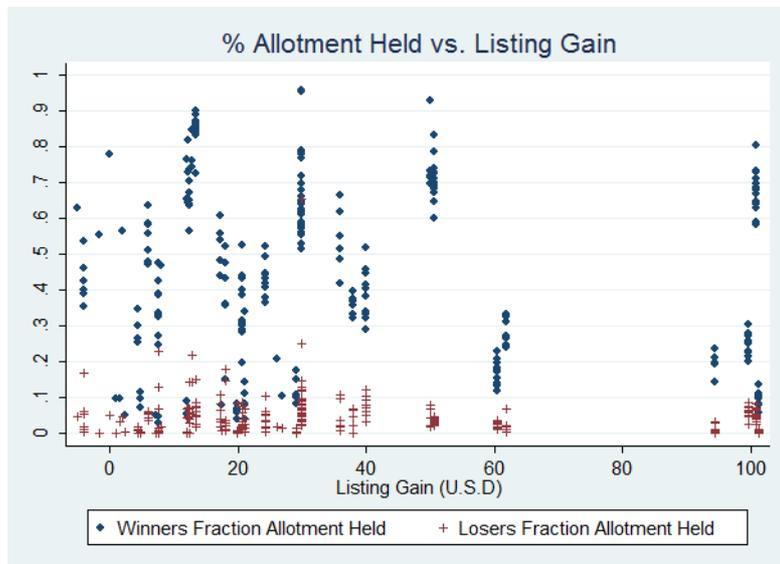


Figure 2: Percent Allotment at End of Month One Versus Listing Return



Supplementary Appendix

A The Probability of Treatment

Let S be the total supply of shares that the firm decides to allocate to retail investors. Let $c = 1, \dots, C$ index “share categories,” which are integer multiples of the minimum lot size x for which investors can bid. The set of possible numbers of shares for which investors can bid is therefore: $x, 2x, \dots, Cx$.²⁶ 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:

$$D = \sum_{c=1}^C cxa_c. \quad (4)$$

Retail oversubscription v is then defined as:

$$v = \frac{D}{S}. \quad (5)$$

As described in case (1) above, if $v \leq 1$ at the ceiling price, then all investors get the shares for which they applied, and if $v > 1$, one of cases (2) or (3) will apply.

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 $c \in [J, \dots, C]$ receive proportional allocations which are greater than or equal to x , and share categories $c' \in [1, \dots, J]$ receive proportional allocations which are less than x . If $J = 1$ then we are in case (2), otherwise we are in case (3).

In either case, investors in share categories $c \geq J$ receive a proportional allotment $\frac{cx}{v}$, and a total number of shares equalling $\sum_{c=J}^C \frac{cx}{v} a_c$. However, investors in share categories

²⁶Note that the minimum lot size is also the mandatory lot size increment.

$c' \in [1, \dots, J)$ 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.,²⁷

$$Z = S - \sum_{c=J}^C \lfloor \frac{c}{v} x a_c \rfloor. \quad (6)$$

These are the shares allocated by lottery in case (3). 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 $c' \in [1, \dots, J)$ 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 $c' \in [1, \dots, J)$:

$$\frac{p_{c'}}{p_{c'-1}} = \frac{c'x}{(c'-1)x} = \frac{c'}{c'-1}. \quad (7)$$

The combination of equation (7) and the fact that the total remaining shares are described by equation (6) gives us:

$$\sum_{c'=1}^{J-1} (p_{c'}) x a_{c'} + \sum_{c'=1}^{J-1} (1 - p_{c'}) \times 0 = Z. \quad (8)$$

Solving (8), we get that $p_{c'} = \frac{c'}{v}$ of winning exactly x shares in share categories $c' \in [1, \dots, J)$. We show the solution in an appendix to the paper.

In general, the probability of winning increases proportionally with the number of share lots bid for c , and decreases with the overall level of over-subscription v . This implies that the probability of winning will vary across share categories within IPOs, as well as across IPOs. In other words, there may be some self-selection of investors into share categories – that is, by applying for more share lots, they increase the probability of winning. 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. In other words, the relevant

²⁷By regulation, the shares to be allotted $\sum_{c=J}^C \frac{c}{v} x a_c$ is rounded to the nearest integer.

control group is the set of investors *within* the same share category who were unsuccessful in the lottery. As we explain more fully below, it is precisely this within-share-category experimental variation that we exploit in estimating the effects of winning (or losing) the IPO lottery on subsequent portfolio decisions.

Table A.1: IPO CHARACTERISTICS

	2007	2008	2009	2010	2011	All
IPOs in sample						
Number of IPOs in sample	12	10	2	22	8	54
Percentage of all IPOs in India	12.04	31.58	11.76	32.84	20.51	22.13
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	3.00	8.77	0.72	11.01	24.62	7.71
Percentage issued (Retail investors excl. employees)	33.01	34.33	34.88	32.71	35.00	33.50
Over-subscription ratio	21.95	12.63	2.11	10.10	6.72	12.06
No. of randomized share categories (“Experiments”)	109	55	2	177	40	383
Total no. of share categories	178	152	28	398	227	983
No. of IPOs from different sectors						
Technology	1	1	0	2	0	4
Manufacturing	8	6	2	12	3	31
Other Services	2	3	0	8	4	17
Retail	1	0	0	0	1	2

Table A.2: EXAMPLE IPO ALLOCATION PROCESS: BARAK VALLEY CEMENT IPO ALLOCATION

Share Category (c) (0)	Shares Bid For ($c \times x$) (1)	# Applications a_c (2)	Total Shares $a_c \times c \times x$ (3)	Proportional Allocation $\frac{cx}{v}$ (4)	Win Probability $\frac{c}{v}$ (5)	Shares Allocated (6)	# Treatment group $\frac{c}{v} \times a_c$ (7)	# Control group $(1 - \frac{c}{v}) \times a_c$ (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

Note: Columns (7) and (8) are obtained after applying the regulation defined rounding off methodology as described in section 2.1.

Figure A.1: IPO FREQUENCY

