Online Appendix
Learning from noise: Evidence from India's IPO lotteries

Santosh Anagol   Vimal Balasubramaniam   Tarun Ramadorai

A Appendix Tables and Figures

Figure A.1: Comparison of the lottery sample to India and the United States

Panel A presents the empirical kernel density plot of the distribution of the logarithmic value of all equity investments in US dollars in the United States (black dashed line) from the Survey of Consumer Finances (SCF), 2013, in the Indian depository (green dashed line), and our lottery sample (green-blue dotted line) as of March 2012. The Indian portfolio value distribution is scaled by the ratio of per capita GDP in India to the United States. Panel B shows the histogram of the number of trades placed in a month for the three samples. We divide the annual number of trades by 12 to obtain the average monthly number of trades for the SCF.

Panel A: Portfolio value distribution

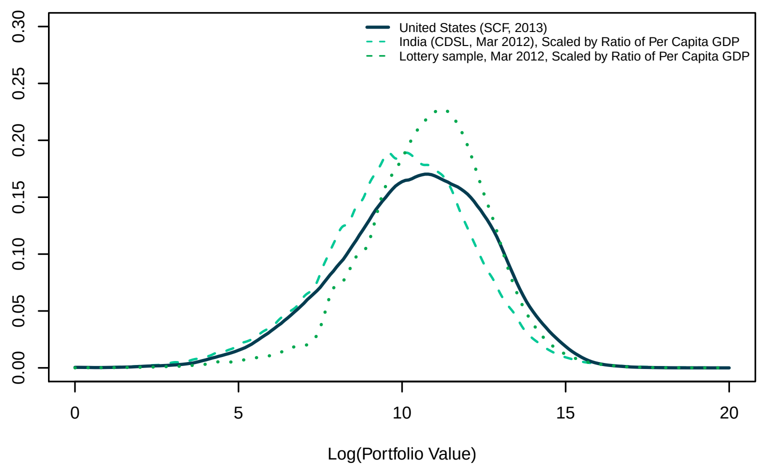

Panel B: Histogram of number of trades

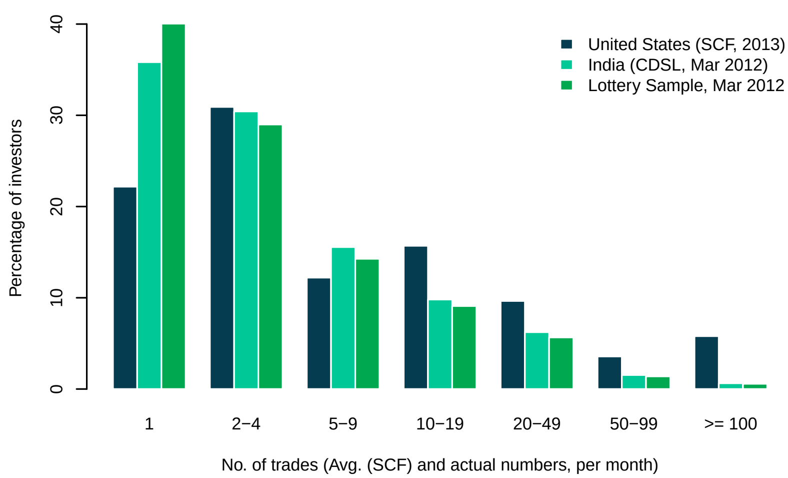

Figure A.2: IPO frequency

This figure plots the monthly frequency of IPOs in India between January 2007 and November 2011. The number of treatment IPOs in sample are plotted in blue, and the rest in maroon color.

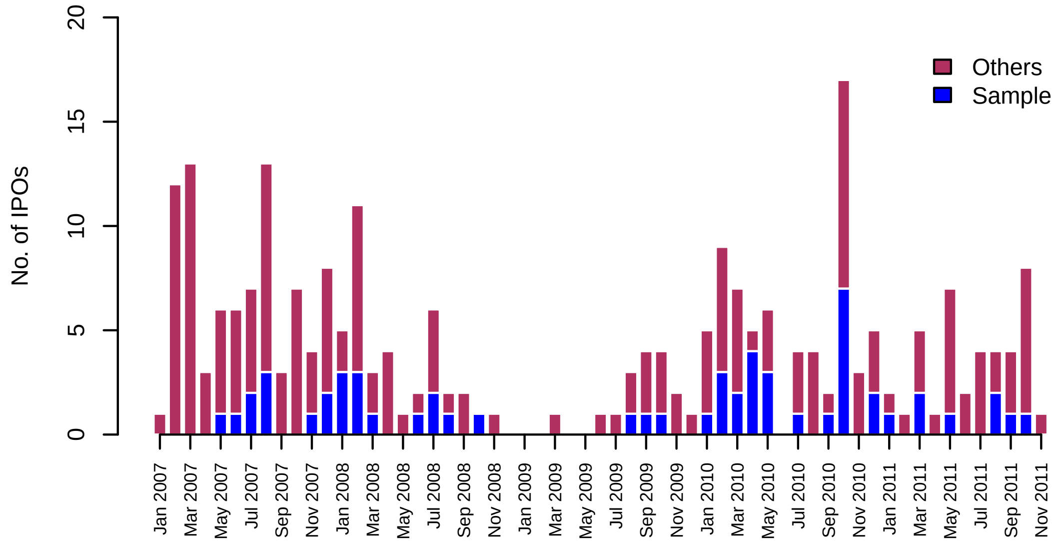

Figure A.3: IPO return experience

Panel A plots the first-day returns computed as the percentage change in the first-day closing price of the stock from the issue price of the IPO stock. Each dot represents a treatment IPO in the sample in each year, and the dashed blue line plots the year average first-day return of our treatment IPOs. Panel B plots the first-day return variability computed as the percentage difference in the highest first-day return and the lowest first-day return for the treatment IPO. Each dot represents a treatment IPO in the sample, and the dashed blue line plots the yearly average return variability of our treatment IPOs.

Panel A: First-day returns

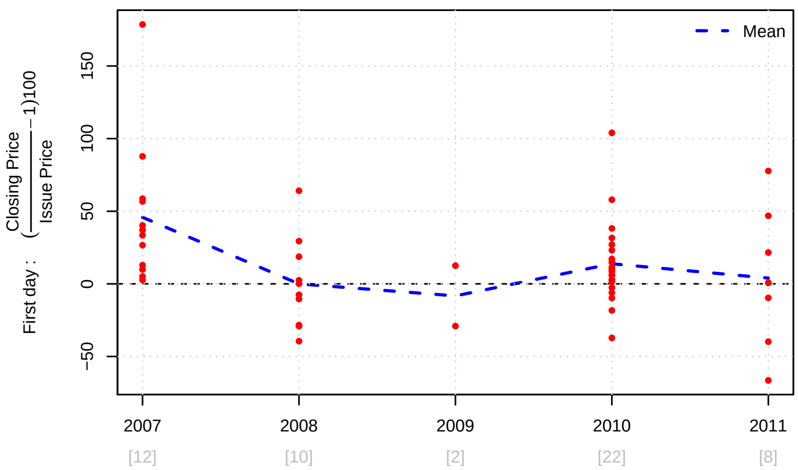

Panel B: First-day returns variability

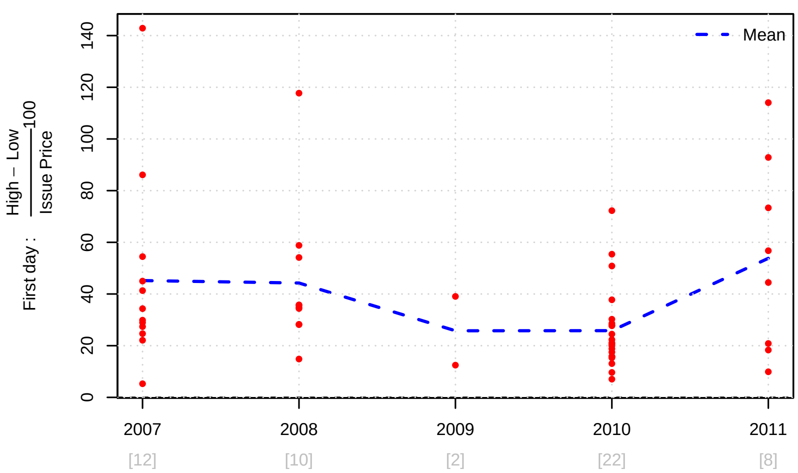

Figure A.4: Mental accounting

Panel A plots the total IPO listing gain (value of IPO shares at the end of the first trading day minus the amount paid for IPO shares) in US dollars on the $x-$axis (log-scale), and the treatment effect on a single purchase ticket size in each investor account in US dollars (log-scale) for each positive listing gain IPO in our sample. Similarly, Panel B, C and D, plot the estimated treatment effect on purchase ticket size averaging across all purchases made by an account, single sale ticket size in each investor account, and the average across all sales made by an account, respectively.

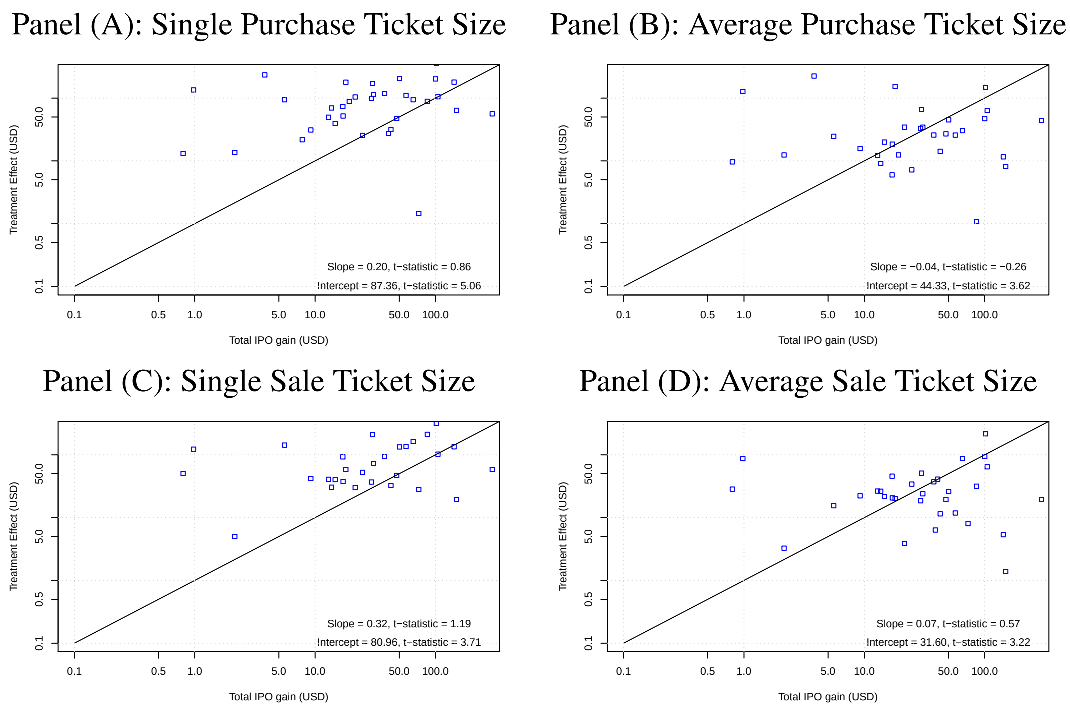

Figure A.5: Daily returns distribution:
Indian stocks 2002 − 2012

This figure plots the pooled daily returns distribution for all Indian stocks between January 2002 and March 2012.

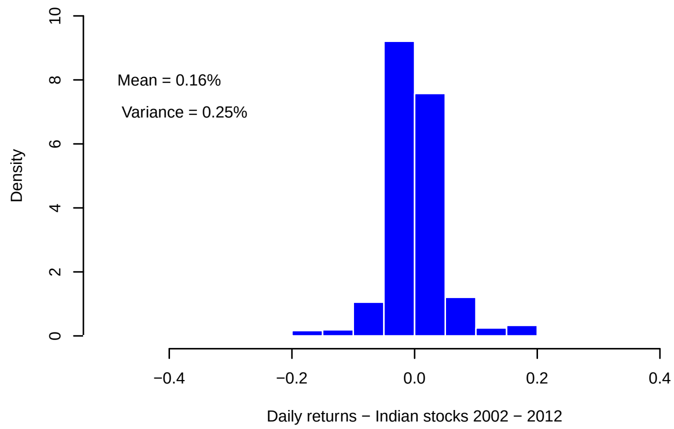

Figure A.6: Treatment effect by past trading activity

This figure plots the estimated treatment effect on Log(1+ Gross Transactions Value) in the first full month after the lottery for investors with different levels of trading activity. The first bar on the figure presents the treatment effect for the first full month after the treatment IPO, for investors between 0 and 5 trades in the month before the treatment IPO. Similarly, the remainder of the three bars in the figure plot the estimated treatment effect for investors with (5,10], (10,30], and with 31 or more trades respectively. The red lines mark the confidence intervals for the treatment effect at 5% confidence level. The sample size for each group is presented in (.) under each bar plot.

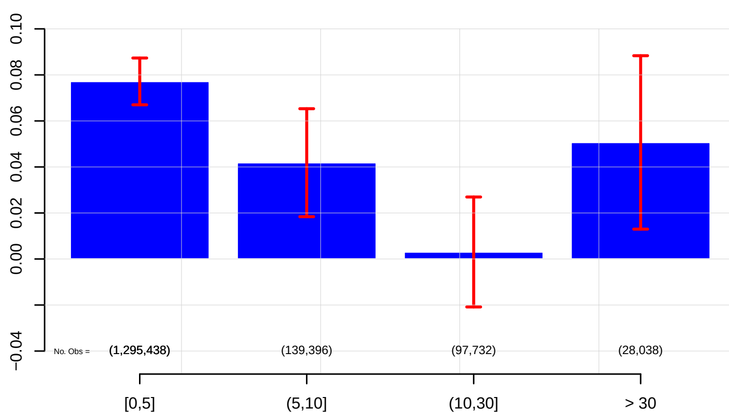

Figure A.7: Return experience across experience bins

Panel A plots the lottery-weighted returns and Panel B, equal-weighted returns experienced in each experience bin, [0,1], [2,3] and >= 4 past random wins, respectively. Each of these experience bins is represented by blue, red and green, respectively. Each group of bar plots in both panels represent different time-horizons at which past random wins are considered. The first set of bars are for random wins within six months before the treatment IPO (Random Wins (6M)), the second set for random wins within 12 months before the treatment IPO (Random Wins (12M)), and the last for all past random wins (Random Wins (All)).

Panel A: Lottery-weighted returns

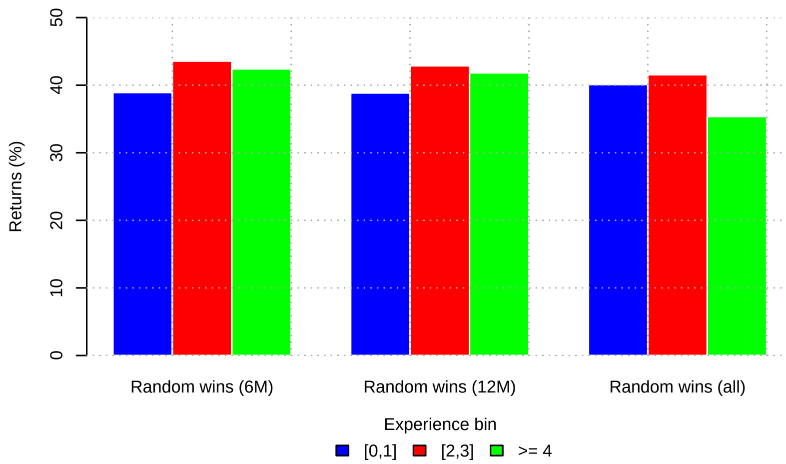

Panel B: Equal-weighted returns

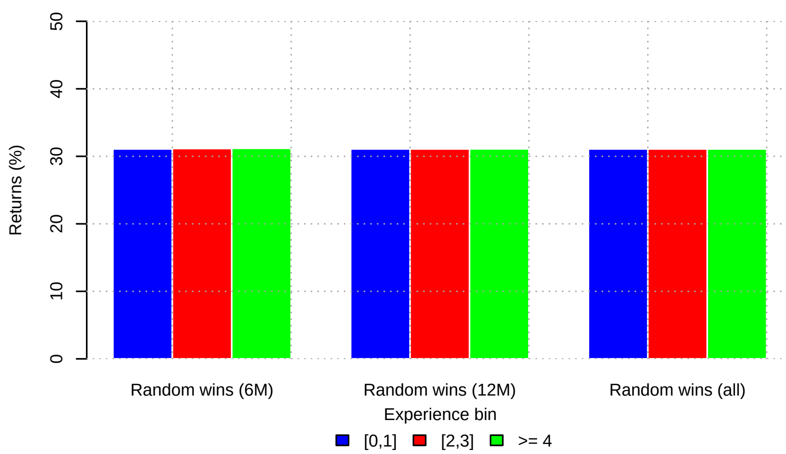

Figure A.8: Heterogeneous Treatment Effect by Net Investment

This figure presents the estimated treatment effect on Log(1+ Gross Transactions Value) in the first full month after the lottery decomposed by quintiles of net investment $NI_{i,t}$, defined to be $PV_{i,t} - (1 + r_{i,t})PV_{i,t-1}$, where $r_{i,t}$ is the account's portfolio return at time t, and PV is the total account portfolio value at time t and $t-1$, in US dollars. The mean net investment within each quintile is presented on the $x-$axis, and the treatment effect in percent on the $y-$axis. The confidence interval marks statistical significance at the 5% confidence level.

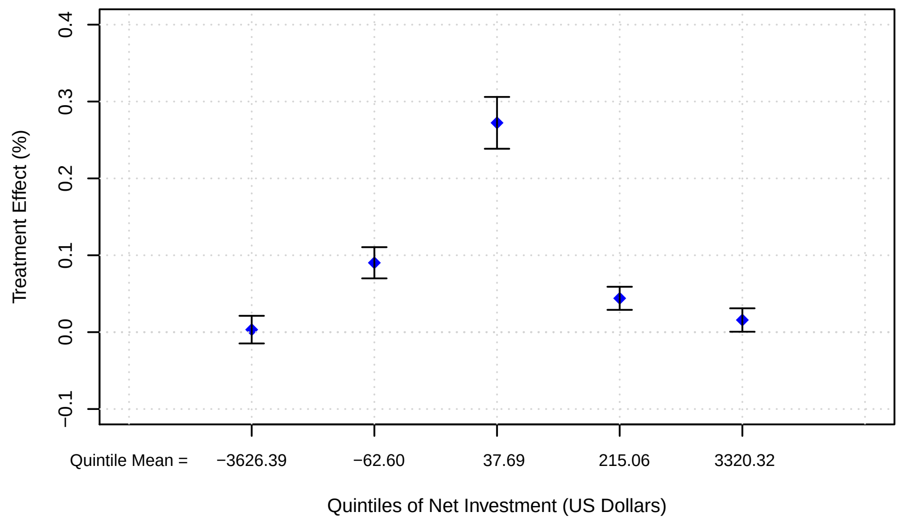

Table A.1: Experience effects for Negative Return IPOs

This table presents the regression estimates of experience effects on investor behavior, for negative return treatment IPOs in sample, mirroring the estimates for positive return treatment IPOs in Panel A, Table 4, in the paper. ***,**,* indicate statistically significant differences at the 1%, 5% and 10% levels respectively.

	Event-time							
	-1	0	1	2	3	4	5	6
1. IPO participation	0.0041 (0.0034)	0.0010 (0.0030)	-0.0055* (0.0030)	-0.0071** (0.0034)	-0.0018 (0.0021)	-0.0045** (0.0020)	-0.0068*** (0.0024)	-0.0023* (0.0012)
2. Gross transaction value	0.0001 (0.0249)	-0.0050 (0.0238)	-0.0048 (0.0225)	-0.0585** (0.0266)	-0.0469* (0.0276)	0.0044 (0.0278)	-0.0493** (0.0250)	-0.0640** (0.0265)
3. Gross transaction value / Portfolio value (t-1)	0.0008 (0.0064)	-0.0039 (0.0267)	-0.0037 (0.0285)	-0.0574* (0.0309)	-0.0458** (0.0215)	0.0054 (0.0319)	-0.0482* (0.0281)	-0.0628** (0.0306)
4. Gross no. of Transactions	-0.0032 (0.0084)	-0.0013 (0.0083)	-0.0075 (0.0078)	-0.0179** (0.0081)	-0.0079 (0.0080)	-0.0009 (0.0081)	-0.0135* (0.0076)	-0.0098 (0.0078)
5. Weight in IPO sector	-0.0001 (0.0009)	0.0007 (0.0012)	-0.0001 (0.0009)	-0.0004 (0.0009)	-0.0006 (0.0009)	-0.0006 (0.0009)	-0.0009 (0.0009)	-0.0013 (0.0009)
6. Portfolio value	-0.0011 (0.0236)	-0.0135 (0.0216)	0.0062 (0.0209)	-0.0048 (0.0214)	-0.0063 (0.0220)	0.0112 (0.0224)	0.0093 (0.0225)	-0.0128 (0.0232)

Table A.2: Treatment effect by past flipping activity

This table tests whether there is a difference in magnitude of trading volume (Log(1+ Gross Transactions Value) in the first full month after the lottery) between those who realize their gains and those who do not, from IPOs. Panel A presents results from a logistic regression for two types of "flipper" variables. Type 1 refers to a dummy variable that takes the value 1 when an investor has ever flipped a previously allotted IPO. Type 2 refers to a dummy variable that takes the value 1 when an investor has flipped their most recent past IPO that was allotted. The two columns present the coefficient estimate, and the odds ratio. Panel B estimates a heterogeneous treatment effect by Type 1 flippers, and non-flippers. Similarly, Panel C by Type 2 flippers and non-flippers. ***,**,* indicate statistically significant differences at the 1%, 5% and 10% levels respectively.

Panel A: Predicting flipping activity by lottery winners		
	β	Odds ratio
Type 1: Ever flipped in the past	0.859*** (0.007)	2.36*** (0.017)
Type 2: Flipped in most recent past IPO	0.304*** (0.011)	1.35*** (0.015)
Panel B: Treatment effect by flipping activity (Type 1)		
Flippers	0.065*** (0.008)	
Non-flippers	0.088*** (0.008)	
Panel C: Treatment effect by flipping activity (Type 2)		
Flippers	0.069*** (0.017)	
Non-flippers	0.078*** (0.006)	

Table A.3: Gross transaction value decomposition By Purchase and Sale value, and IPO sector

This table presents treatment effects separately for purchases and sales (log(1+purchase value), and log(1+sales value)), in the first two columns, and the total trading value (log(1+gross transactions value)) by trading volume in the Non-IPO sectors and the IPO sector in the last two columns, respecitvely.

			Gross Transactions Value	
Event time	Purchase value	Sale value	Non-IPO Sector	IPO Sector
	(1)	(2)	(3)	(4)
-1	0.007	0.01	0.009	0.004
0	0.011*	0.027***	0.021***	0.015***
1	0.063***	0.069***	0.057***	0.068***
2	0.057***	0.058***	0.048***	0.047***
3	0.037***	0.032***	0.025***	0.033***
4	0.023***	0.029***	0.028***	0.020***
5	0.022***	0.030***	0.019***	0.023***
6	0.023***	0.026***	0.024***	0.017***

Table A.4: Heterogeneous treatment effects: Age, Portfolio Value, Trading Activity

This table presents formal tests between randomly experienced returns, past random wins, and different portfolio characteristics, and whether these portfolio characteristics moderate such a relationship. Column 1 presents a regression with account age (Age) as the characteristic, Column 2 with portfolio value (PV) and Column 3 with gross transaction value (TV), both measured the month before the treatment IPO. ***,**,* indicate statistically significant differences at the 1%, 5% and 10% levels respectively.

	Age	**PV**	**TV**
	(1)	(2)	(3)
I(Allot)	0.1224***	0.1554***	0.1468***
	(0.0088)	(0.0115)	(0.0103)
I(Allot) × RandWins (All)	-0.0458***	0.0266***	-0.0347***
	(0.0045)	(0.0077)	(0.0059)
I(Allot) × Age	-0.0008***		
	(0.0003)		
I(Allot) × PV		-0.0080***	
		(0.0011)	
I(Allot) × TV			-0.0081***
			(0.0011)
I(Allot) × Age × RandWins (All)	0.0007***		
	(0.0001)		
I(Allot) × PV × RandWins(All)		-0.0028***	
		(0.0006)	
I(Allot) × TV × RandWins (All)			0.0030***
			(0.0006)
Age × RandWins (All)	-0.0026***		
	(0.0001)		
PV × RandWins (All)		0.0066***	
		(0.0004)	
TV × RandWins (All)			0.0056***
			(0.0003)
PV		0.1477***	
		(0.0006)	
TV			0.1907***
			(0.0006)
Age	0.0060***		
	(0.0002)		
RandWins (All)	0.2549***	-0.0068	-0.0006
	(0.0026)	(0.0047)	(0.0035)
N	1,561,496	1,561,496	1,561,496

Table A.5: Evaluation of heterogeneous treatment effects

Panel A documents the cut-off values at the 25th, 50th and 75th percentile of the cross-sectional distribution of three portfolio characteristics, and the number of all past random wins. Panel B.1 presents formal estimates of the changes in the treatment effect by account age to experience. Panel B.2, and B.3, present the estimates for portfolio value and transaction value, measured at the end of the month before treatment IPO, respectively.

Panel A: Percentile of cross-sectional distribution			
	25th	50th	75th
Age (in months)	4 (Young)	15 (Middle-aged)	34 (Old)
Log(1+Portfolio value)	8.3 (Small)	10.8 (Medium)	12.18 (Large)
Log(1+Gross transactions value)	0 (Small)	9.9 (Medium)	11.42 (Large)
No. of past random wins (all)	0		2

Panel B.1: Age			
Coefficient Estimate		**Difference**	
(1)		(2)	
Young(2-0)	-0.0887***	Young(2-0) – Old(2-0)	-0.0659***
	(0.0080)		(0.0078)
Middle-aged(2-0)	-0.0677***	Young(2-0) – Middle-aged(2-0)	-0.0210***
	(0.0062)		(0.0025)
Old(2-0)	-0.0228***	Middle-aged(2-0) – Old(2-0)	-0.0449***
	(0.0048)		(0.0053)
Panel B.2: Log(1+Portfolio Value)			
Coefficient Estimate		**Difference**	
Small(2-0)	0.0531***	Small(2-0) – Large(2-0)	0.0701***
	(0.0159)		(0.0170)
Medium(2-0)	-0.0061***	Small(2-0) – Medium(2-0)	0.0592***
	(0.0044)		(0.0144)
Large(2-0)	-0.0170***	Medium(2-0) – Large(2-0)	0.0109***
	(0.0046)		(0.0026)
Panel B.3: Log(1+Gross Transactions Value)			
Coefficient Estimate		**Difference**	
Small(2-0)	-0.0693***	Small(2-0) – Large(2-0)	-0.0697***
	(0.0121)		(0.0134)
Medium(2-0)	-0.0103***	Small(2-0) – Medium(2-0)	-0.0590***
	(0.0041)		(0.0114)
Large(2-0)	0.0003	Medium(2-0) – Large(2-0)	-0.0106
	(0.0046)		(0.0021)

Table A.6: Treatment Effect on Net Investment into Trading Account

This table presents the estimated treatment effect of randomized return experience on net investment $NI_{i,t}$, defined to be $PV_{i,t} - (1 + r_{i,t})PV_{i,t-1}$, where $r_{i,t}$ is the account's portfolio return at time t, and PV is the total account portfolio value at time t and $t - 1$, in US dollars. The first row estimates the treatment effect without considering investment into the IPO stock, and the second, including the IPO stock.

	-1	0	1	2	3	4	5	6
Net Investment	-1.910	-4.756	-6.1468*	1.851	3.623	-0.740	0.476	0.186
(w/o IPO stock)	(3.496)	(3.672)	(3.395)	(3.210)	(2.988)	(2.980)	(2.838)	(2.695)
	15.675	362.543	168.861	263.742	36.784	-69.355	76.718	-5.035
Net Investment	-1.910	117.8628***	-36.8046***	-7.7383**	-5.7662*	-5.0442*	-5.5884*	-2.106
(with IPO stock)	(3.496)	(3.675)	(3.411)	(3.216)	(2.993)	(2.983)	(2.841)	(2.698)
	15.675	398.799	162.545	261.744	34.232	-70.549	74.893	-5.575

Table A.7: Return experience and trading volume: Randomized vs. Non-Randomized Estimates

This table presents a comparison of experience effect estimates using randomized variation in returns experience from our main empirical specification (equation 1) (column 1), with the past returns experience for a sample of investors who purchased exactly one stock in month $t-1$ ("Purchase returns", columns 2 and 3). Panel A presents the estimates with no additional control variables, without (columns 1–2) and with (column 3) calendar month fixed-effects. Panel (B) presents the same specification as in Panel (A) with additional controls, including individual fixed-effects. The coefficient for randomized returns (columns 1) is normalized by mean IPO returns of 39.18% for comparison, at the end of first full month after listing ($\frac{0.07}{39.18\%} = 0.0018$). We benchmark to the positive returns sample to be consistent with Table 4 of the paper. Columns 2 and 3 are estimated on a 10% random sample from the universe of CDSL investors, observed over 120 months. Robust standard errors in parentheses, and ***,**,* represent statistical significance at 1%, 5% and 10% respectively.

	(1) Randomized returns	(2) Purchase returns	(3) Purchase returns
Panel A: "Raw" estimates			
Experience effect	0.0018***	0.0158***	0.0024***
	(0.006)	(0.000)	(0.000)
IPO × Share category Fixed Effects	Yes	No	No
Calendar-Month Fixed Effects	No	No	Yes
Panel B: Estimates with additional controls			
Experience effect	0.0018***	0.0035***	0.0033***
	(0.006)	(0.000)	(0.000)
IPO × Share category Fixed Effects	Yes	No	No
Time fixed effects	No	Yes	Yes
Individual fixed effects	No	Yes	Yes
Control variables (lagged)			
Portfolio size	No	No	Yes
Trading volume	No	No	Yes
Age	No	No	Yes
No. of securities held	No	No	Yes
N	1,473,073	2,337,002	2,337,002

B Indian IPO Regulations

Indian IPO regulations are relatively complex and vary by whether an unlisted company fulfil specific eligibility rules. Although the allocation procedure for retail investors, as described in the main text of the paper, does not vary across IPOs, regulations surrounding the allocation procedure can be different. We document the nature of IPOs in our sample alongside the Indian IPO regulations below.

The Securities Exchange Board of India (SEBI) Disclosure and Investor Protection Guidelines (till 2009), henceforth "DIP guidelines", SEBI Issue of Capital and Disclosure Requirements Regulation (since 2009), henceforth "ICDR regulations", and Section (19) (b) (2) of the Securities Contract Regulation Rules ("SCRR") made under the Securities Contract Regulation Act, 1956, alongside the Companies Act, 1956 govern the IPO process in India.

Eligibility criteria

An unlisted company may make an initial public offering (IPO) of equity shares if it meets the following conditions alongside at least 1000 investors participate in the IPO process (Rule-set 1): [1]

1. The company has net tangible assets of at least ₹30 million in each of the preceding three full years (calendar years), of which not more than 50% held in monetary assets. If more than 50% held in monetary assets, the company has firm commitments to deploy excess monetary assets in its business.

2. The company has a track record of distributable profits (as defined in the Companies Act, 1956), for at least three years out of the immediately preceding five years.

3. The company has a net worth of at least ₹10 million in each of the preceding three full years (calendar years).

4. The aggregate of the proposed issue and all previous issues in the same financial year in terms

[1]See Page 15-16, Section 2.2.1 of DIP guidelines, which is similar to Chapter II of the ICDR regulations, accessed on April 20 2015. They can are at http://www.sebi.gov.in/guide/sebiidcrreg.pdf and http://www.sebi.gov.in/guide/DipGuidelines2009.pdf

of size does not exceed five times its pre-issue net worth as per the audited balance sheet of the last financial year.

When a company does not fulfil these requirements, it can still undertake an IPO provided the following conditions are fulfilled (Rule-set 2):[2]

1. The issue is made through the book-building process, with *at least 50% of the net offer to public* allotted to Qualified Institutional Buyers (QIBs), failing which all subscription amount will have to be refunded.[3]

2. The minimum post-issue face value of capital will be ₹100 million.

All 54 IPOs in our sample are book-built IPOs, where the net offer to the public are allocated according to the same procedure.[4] All book-built IPOs need to mandatorily achieve a minimum of 90% of the initial intended issue.[5] When a company undertakes a 100% book-built issue, the following percentage of issue will have to be initially set aside for the following investor categories:[6]

1. *Not less than* 35% of the net offer to public will be made available to *retail investors*

2. *Not less than* 15% of the net offer to the public will be made available to *non-institutional investors*

3. *Not more than* 50% of the net offer to the public shall be made available for allocation to QIBs.

---

[2]See Page 18, Section 2.2.2 (i) - (iv) of the DIP guidelines, identical to the conditions in ICDR regulations, accessed on April 20, 2015, at `http://www.sebi.gov.in/guide/sebiidcrreg.pdf` and `http://www.sebi.gov.in/guide/DipGuidelines2009.pdf`

[3]QIBs are defined under Chapter I, definition (zd) of the ICDR regulations (Page 6). This includes mutual funds, venture capital funds (domestic and foreign), a public financial institution, banks, insurance companies and so on.

[4]See Section 11.3.5 (i) of DIP guidelines accessed on April 20 2015.

[5]See ICDR (2009), Chapter I (14) (1), page 13

[6]The Indian regulator, SEBI, introduced the definition of a retail investor on August 14, 2003, and capped the amount that retail investors could invest at 50,000 rupees per brokerage account per IPO. This limit was increased to 100,000 rupees on March 29, 2005, and again increased to 200,000 rupees on November 12, 2010. See Section 11.3.5 (i), footnotes 480,481,482,483 on Page 216 of the DIP guidelines. "Non-institutional buyers" are all those who are not QIBs and Retail Investors - see Chapter I, definition (w) on Page 5 of ICDR regulations.

When the company does not fulfil the criteria set in Rule-set 1, then condition (3) above is *mandatory*. Further, when the company undertakes an IPO under the SCRR, the percentage requirements become 30% (retail investors), 10% (non-institutional investors) and a *mandatory* 60% to QIBs. Any shares set-aside for employees of the company is also considered to be under the "retail investor" category.[7]

For purposes of the Indian regulation, "retail investors" are defined as those with expressed share demands beneath a pre-set value.[8] At the end of the sample period that we consider, this pre-set value was set by the regulator at ₹200,000 (roughly US $3,400); this value has varied over time.[9]

Once the bidding is complete if any of the investor categories are under-subscribed (subject to the allocation rule above), then, with full disclosure and in conjunction with the stock exchange, a company can reallocate the shares to the other investor categories.[10] However, the QIB category cannot be under-subscribed if the IPO is undertaken under Rule-set 2 or Section (19) (2) (b) of the SCRR.

While the regulation provides for an alternative in the event of under-subscription, in reality, this occurs more frequently with non-institutional investors. Data from our sample of 54 IPOs show that non-institutional buyers are almost always under-subscribed. Retail investors are therefore significant to achieve the minimum of 90% of the initial intended issue, without which the IPO will fail.

In our sample of 54 IPOs, firms issue under both the SCRR and the DIP/ICRR paths. Further,

[7]Note that this has been inferred from Section 11.3.5 (i), read with footnotes 480-483 on Page 216 of the DIP guidelines.

[8]In practice, each brokerage account is counted as an individual retail investor by 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 as our data are aggregated across all brokerage accounts associated with the anonymized tax identification number of the investor.

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

[10]See DIP guidelines (2009), Section 11.3.2 (v) read with 11.3.5 (i) and 11.3.5 (iv) (Pages 217-219).

the *ex-post* percentage of the total final public issue to retail investors can be higher than the values mentioned above. This will have to be explicitly disclosed at the time of allotment of an issue. In our sample, nearly one-third of the total (final) issue size is always allotted to retail investors.

The share allocation process in an Indian IPO begins with the lead investment bank, along with the issuer, setting 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. Throughout the main paper, we refer to the potential available integer multiples of the minimum lot size as separate "share categories". For example, the Barak Valley Cement IPO, described in the paper, had a minimum lot size of 150 shares, and investors could bid for share categories of 150 up to 2,250 shares (i.e. 150, 300, 450, ..., 2250). These share categories are critical because the lottery randomization is always conducted *within categories*, i.e., only across investors who applied for a specific share category.

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 cutoff 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. If the bid is successful and the investor is allotted shares, the order will be filled at the investor's 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—the remaining 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 investor types (for example, if employees and/or non-institutional investors participate in amounts less than they are offered, this can "overflow" into additional retail supply).[11] The remainder of the allocation procedure is explained in the main paper, as it forms the core of our identification stratey.

C Balance Tests

Table 3 shows that virtually identical fractions (38%) of both treatment and control investors applied to an IPO with our registrar, or were allotted shares in an IPO not covered by our registrar, in the month *prior* to treatment. Nearly identical fractions (92.6%) of the applicants placed a cutoff bid in IPOs and the remainder a full demand schedule.

The next set of variables describe the trading behavior of our treatment and control samples. We focus on the total dollar trading (purchase and sales) volume, calculated as the sum of the value of stocks bought and sold in a month. We find that the average monthly trading volume is roughly US$ 275, including zeros. These values are highly skewed, so we transform this variable using $log(1+x)$ transformation. While 29% of accounts made no trades in the month before treatment, nearly half of the accounts observed traded more than US $1,000 in the month before treatment. Overall, many investors in the sample trade substantial amounts.

The next block of rows of Table 3 shows statistics about the distribution of investor portfolio values, the number of securities held in the portfolio the month before the IPO and the account age of investors.

The table shows that 78.5% of treatment and control investors had an account value greater than zero in the month before the IPO. Portfolio value amounts are also highly skewed, so we once again

[11] Of 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.

transform this variable with a log(1+x) transformation. Portfolio values average about US$ 783, including zeros and are not significantly different across treatment and control accounts.

The next few rows show the fractions of treatment and control accounts that fall into the range of portfolio values described in the row headers. The distribution of portfolio values is roughly U-shaped in both treatment and control accounts, with a relatively large number of accounts with zero value (some of these correspond to new market entrants (rookies), as we identify below), few accounts with portfolio value between US$ 500 and 1,000, and roughly a quarter of the accounts with portfolio values over US$ 5,000.

On average, both treatment and control accounts hold about nine different equities in their portfolio the month before treatment. In terms of account age at the time of the treatment IPO, 31% of accounts are less than six months old, 31% are between 7 and 25 months old, and about 38% are over 25 months old.

Overall, we find that the differences across treatment and control groups are small, and importantly, not statistically significant. The fraction of experiments with greater than ten percent significance is around ten percent. Given the similarity of treatment and control groups across this extensive set of background characteristics, the IPO shares assignment appear to random within share categories, to investors.

D "Depository Shifting" Behavior from Investors

When an individual has multiple accounts with CDSL, we aggregate their market activity by the unique anonymized tax account number before we conduct our analysis, thus capturing trading behaviour across multiple accounts within CDSL. When an individual has accounts across both NSDL and CDSL and moves their trading activity from one to another depository, they can do it in two ways. First, they could keep money in both accounts, but move their trading activity from one account to another. If depository shifting happens in this manner, there is no way for us to detect it given we do not possess the data required to connect trading behavior across depositories for investors that hold accounts in both NSDL and CDSL.

A second possible way is that investors move money from NSDL accounts to CDSL accounts after winning an IPO lottery in a CDSL account (and vice-versa after losing an IPO lottery in a CDSL account). In principle, an investor could trade and close out positions in one depository account, and repurchase positions in the other depository account. Alternatively, they could provide a written request for transfer of their shares held in one depository account to another with proof of identity and a lengthy process for the transfer of shares held in their account.

We can evaluate if this form of the "depository shifting" hypothesis can explain our results. The key observation is that this form of the depository shifting hypothesis predicts that our results should mainly be driven by net new inflows into accounts, in addition to the increases in trading volume we observe. For example, suppose lottery winners moving their investments from their NDSL account to their CDSL account can drive our results; then, we should observe an increase in net new investment into the CDSL account right after the lottery. Our test has two parts. The first is to observe whether net new investment increases substantially more for lottery winners than losers in the CDSL accounts, which would be consistent with the shifting of money from NSDL to CDSL accounts (on average). The second is to test whether our treatment effects are concentrated in accounts that make large new net investments.

Borrowing from the mutual funds literature measuring net AUM adjusting for returns accrual as in Sirri and Tufano (1998), we measure net new investment for an account i in month t as follows:

$$NI_{i,t} = PV_{i,t} - (1 + r_{i,t})PV_{i,t-1}$$

where $r_{i,t}$ is the account's portfolio return at time t, and PV (in USD) is the total account portfolio value at time t and $t-1$ respectively. We measure net investment without the IPO stock in question as in all other outcome variables, and also estimate the total net new investment including the IPO stock in our analysis.[12]

Online Appendix Table A.6 presents results on the treatment effect of winning the lottery on net investment. The first panel excludes the IPO stock (both the IPO stock allocation, as well as any

[12]We also verify the robustness of our tests to assuming that flows come in at the beginning of the month, before returns are accrued (as in Sirri and Tufano (1998)), i.e., $NI_{i,t} = PV_{i,t} - (1 + r_{i,t-1})PV_{i,t-1}$ and our tests yield consistent results and do not depend on this assumption.

trading in that particular stock). We find no statistically significant increase in net-new investment in the months after the lottery when we exclude the IPO stock, and when we include the IPO stock we find a negative treatment effect (primarily driven by lottery winners selling their IPO allocations for cash). We note here that the increase in net investment in the second panel, for event month zero, is a mechanical increase in portfolio size with the IPO allotment for lottery winners. These results are inconsistent with the idea that lottery winners are transferring large amounts from their NSDL to CDSL accounts after winning the lottery.

While on average net investment is not increasing for treatment accounts, it is theoretically possible that the increase in trading volume we observe (the main result in our paper) is concentrated in accounts where net investment did increase. To delineate whether our treatment effects are a result of new money, we take the quintile of $NI_{i,t}$ at the first full month after the IPO (event time = 1), and estimate the treatment effect within each quintile. The middling bin, i.e., 3rd quintile, is formed of investors with the least amount of outflow/inflow of money. If the treatment effect is valid for this sample, one may argue that this is not due to any inflow/outflow for either the lottery winners or losers, thus providing supporting evidence to rule out the depository shifting explanation.

We note here that $NI_{i,t}$ in and after the first full month is itself an outcome variable affected by treatment so within each quintile it is technically possible that the treatment and control groups are imbalanced. However, given that we find no treatment effects on net investment overall (when excluding the IPO stock), we do not think this is a significant concern.

Online Appendix Figure A.8 decomposes the treatment effect by quintiles of net investment. Our results primarily emanate from accounts with little or no additional investment, thus providing additional evidence that the estimated treatment effects are not significantly driven by net new investment in the market.

E Randomized and Non-Randomized Return Experiences

To conduct a comparison between randomized and non-randomized variation in return experiences, we study the return experience on the purchase of the most recently selected stock for investors who made exactly one purchase. To facilitate comparison with the randomized setting, we compute the total gross transaction value as the sum of purchase and sale transactions values for all stocks other than the most recently purchased stock in a given month t. We then use log of (1+gross transaction value) computed in this fashion.

We first draw a 10% random sample from the universe of CDSL investor accounts, observed over 120 months. To explain the outcome of interest y, we then estimate:

$$y_{i,t} = \mu_i + \omega_t + \gamma \text{RetExp}_{i,t-1} + X_{i,t-1}{}'\beta + \varepsilon_{i,t} \tag{1}$$

Here, μ_i are individual fixed effects, ω_t are calendar month fixed effects, $\text{RetExp}_{i,t-1}$ is a measure of the experienced experienced return at the end of month $t-1$, and $X_{i,t-1}$ are investor level controls. The coefficient γ is the experience effect estimated using traditional, non-randomized return experience.

Column 2 of Online Appendix Table A.7 presents results when the sample is restricted to investors who purchased exactly one stock in month $t-1$, meaning that the measure of experience is the returns on this purchased stock.

Panel A of Online Appendix Table A.7 presents the estimates of the above regression without including controls X on the right-hand side, while Panel B adds controls measured with a one-month lag. These controls include portfolio size, trading value, account age, and the number of securities in the portfolio. In some specifications for non-randomized returns, we also add individual-specific fixed effects.[13]

In both panels of the table, Column 1 reproduces our main results using randomized experienced returns. In order to be consistent about the sample referred to throughout the paper, we use the positive returns sample of 1,473,073 for this analysis. The difference in magnitudes here arises from a scaling factor. To make this specification comparable to the one with a continuous return variable, we normalize the treatment effect of winning the lottery by 39.18%, as this is the weighted average return on our IPO lotteries.

[13]We do not add these for the randomized returns, given that many individuals participate in only one lottery, and including them will eliminate the treatment-control balance induced by the experiment.

In Panel A, Column 2 do not include calendar-month fixed-effects, while Column 3 of Panel A add calendar-month fixed-effects to the right-hand side. The inclusion of calendar-month fixed effects causes a considerable difference when experience is measured using non-randomized returns. The estimated effects reduce by a substantial margin with calendar-month fixed effects.

It is perhaps not surprising that the portfolio spillover effects of non-experimental returns are highly sensitive to the inclusion of time fixed effects; any common factors driving returns will confound the estimate of experienced return effects. The advantage of the experimental variation in returns delivered by the IPO lotteries is that it enables estimation of portfolio spillover effects without such confounds affecting inferences.

To compare the point estimates from the randomized and non-randomized return experiences, we turn to Columns 3 of Panel (B). The experience effect looks quite different, depending on the method of estimation. It is perhaps fairest to compare the magnitudes of the single stock purchase results (Column 3) with the randomized return experience (Column 1). Here, we can see that the point estimate for the experience is roughly two times higher for the non-randomized experience measure than that generated by the experimental variation. In percentage point terms, a 10 percent increase in a purchased security return leads to a 3.3% increase in trading volume, versus 1.8 % increase in trading volume in the case of randomized return experiences.

Both estimates are positive and statistically significant and quantitatively similar. We hesitate to draw firm conclusions from this comparison, as there are other differences between winning and IPO lottery and purchasing a stock.

F Bayesian Learning From Noise: A Model

We consider an agent maximizing exponential utility $-\exp(-\gamma W)$ of terminal wealth W, with risk aversion coefficient γ.

$$\max_{q} -\exp(-\gamma W) \tag{2}$$

$$\text{s.t.} W = r_f(W_0 - q) + qr \tag{3}$$

The agent has to allocate initial wealth W_0 between a risk-free asset with return r_f, and a risky asset with price normalized to 1, delivering a random gross payoff r. The true distribution of r is $N(\bar{r}, \sigma_r^2)$ and the choice variable q, is the amount invested in the risky asset.

In this setting, the agent does not know the true distribution of the risky asset payoff r. She begins with a prior distribution and then updates to a posterior distribution based on signals (return experiences, winning the IPO lottery, etc.). We call the prior distribution of the risky asset payoff $\hat{r}$ and assume that this prior distribution is normal with mean payoff vector $\bar{r}_0$ and variance σ_0^2 (i.e. $\hat{r} \sim N(\bar{r}_0, \sigma_0^2)$). For our purposes it is useful to define the precision of the return as $\rho_0 = \frac{1}{\sigma_0^2}$.

Maximizing utility subject to the prior distribution of asset returns, the optimal portfolio allocation is:

$$q^* = \frac{(\bar{r}_0 - r_f)\rho_0}{\gamma} \tag{4}$$

As in standard mean-variance portfolio analysis, the agent allocates greater amounts to the risky asset when 1) the mean excess return $(\bar{r}_0 - r_f)$ increases, and 2) the precision of her signal regarding returns increases.

F.0.1 Learning about the distribution of market returns

The first case that we consider is that lottery winners experiencing gains or losses interpret the positive or negative noise shock as an informative signal about the distribution of future returns, whereas lottery losing agents act as if they receive no such signal. Of course, investors that process information completely would react to returns on the IPO stock regardless of whether or not they won the lottery. If this were true, we would find no differences in trading behavior between winners and losers, which is clearly not borne out in the data.[14] We denote the signal by s, and define it as $s = r + \varepsilon$ where r can be interpreted as the realized listing return on the randomly assigned IPO stock, and the noise in the signal ε is assumed to be normally distributed with mean zero and variance σ_ε^2. As in Chamley (2004), with Bayesian updating in response to this signal, the agent's posterior distribution is $N(\bar{r}_1, \frac{1}{\rho_1})$, with:

[14]Camerer and Ho (1999) present a model where agents' learning involves differential weighting actual versus foregone returns. In the context of that model we are assuming that agents place weights of one and zero respectively on actual versus forgone payoffs.

$$\rho_1 = \rho_0 + \rho_\varepsilon \quad (5)$$

$$\bar{r}_1 = \alpha s + (1-\alpha)\bar{r}_0 \quad (6)$$

$$\alpha = \frac{\rho_\varepsilon}{\rho_1} \quad (7)$$

Posterior signal precision increases as long as the precision of the incoming signal is greater than zero (i.e., if $\rho_\varepsilon > 0$).[15] The posterior mean return is a weighted average of the signal and the prior mean return, and the signal gets more weight as the precision of the signal increases relative to the precision of the prior distribution. Taking the ratio of the portfolio allocation to the market subsequent to receiving the signal to the portfolio allocation amount prior to receiving the signal, the allocation to the market increases if and only if:

$$\frac{q_1}{q_0} = \frac{\rho_1(\bar{r}_1 - \bar{r}_f)}{\rho_0(\bar{r}_0 - \bar{r}_f)} = \frac{(\rho_0 + \rho_\varepsilon)(\alpha s + (1-\alpha)\bar{r}_0 - \bar{r}_f))}{\rho_0(\bar{r}_0 - \bar{r}_f)} > 1. \quad (8)$$

We know that $\rho_0 + \rho_\varepsilon > \rho_0$ given the updating rule. For a signal that is larger than the mean of the prior distribution (i.e., $s > \bar{r}_0$) we also know that $\bar{r}_1 > \bar{r}_0$. This implies that agents will increase their allocation to the risky asset.

Given the average one-day return on IPOs is 39 percent, it seems reasonable to assume that average signal from winning the IPO is above the investors' prior expected return. It is clear that if agents' learn from their random IPO experience about future IPO returns, then this learning can explain the increased investment in future IPOs. Under this assumption that the learning from the IPO returns pertains to the return on the stock market as a whole, the prediction is that lottery winners in positive return IPOs should increase their allocation to the market; we do not observe this in the data, suggesting that learning about the level of returns available in the market is not the

[15]Note that the precision increasing is a consequence of the normal assumption; if returns and signals were assumed binary, it is possible that precision can decrease (Chamley, 2004). Given that the distribution of stock returns is closer to normal than binary we focus on the more natural case for this context of normal returns.

main mechanism driving our broader results on the impact of this experience in the IPO market.

We are also interested in the predictions of this model about the agent's trading volume. In this simple model, we interpret trading volume as the magnitude of the change in q^* in response to the arrival of future signals (e.g., any signals the investor might receive in subsequent periods). We interpret signals here broadly to mean any kind of information or news on the performance of stocks.

The key result here is that the larger the precision of the prior distribution of the agent, the less her risky asset share will respond to any given *new* signal on average. To see this, note that the relative change in q^* can be rewritten as:

$$\frac{q_1}{q_0} = \frac{(\rho_0+\rho_\varepsilon)}{\rho_0}\frac{\left(\frac{\rho_\varepsilon}{\rho_0+\rho_\varepsilon}s+\frac{\rho_0}{\rho_0+\rho_\varepsilon}\bar{r}_0-\bar{r}_f\right)}{(\bar{r}_0-\bar{r}_f)} \tag{9}$$

Two forces make the agent less likely to respond to signals when their prior precision is larger. First, the difference between the posterior and prior expected return ($\bar{r}_0$) will be smaller, as a greater precision ρ_0 means the prior is weighted more heavily in forming the posterior. Second, given $\frac{\rho_0+\rho_\varepsilon}{\rho_0}$ is decreasing in ρ_0, for any given change in posterior return, the agent changes her allocation less.

The point here is that treated agents update their beliefs relative to the control group, meaning that the posterior signal precision of the treated is ρ_1. This becomes their new prior, and so comparing the treated agents to the control agents, who continue to have signal precision ρ_0, we would expect *future* trading volume of the treated to be lower than that of control agents, given the discussion in the previous paragraph. This model therefore predicts that lottery winners should trade *less* than lottery losers after the noise shock; intuitively, winning the lottery provides information which (at least weakly) reduces the value of future information in forming trading decisions. This prediction, however, is not supported in the data.

F.0.2 Learning about own ability

We next consider the case in which the agent interprets the randomly experienced gains or losses from the IPO as a positive or negative shock about their own investment ability. To do so, we assume that a sensible measure of the agent's ability is the precision of the signal distribution ρ_ε; the higher the agent's signal precision, the greater is the agent's ability. This is a common assumption in the literature on stock trading (see, for e.g., Gervais and Odean (2001); Linnainmaa (2011)).[16]

The key prediction is that an agent who believes that the signals he receives are of higher precision will respond more to any given future signal, i.e., such an agent will trade more.

To see this, we present equation (10) again below:

$$\frac{q_1}{q_0} = \frac{(\rho_0+\rho_\varepsilon)}{\rho_0}\frac{\left(\frac{\rho_\varepsilon}{\rho_0+\rho_\varepsilon}s+\frac{\rho_0}{\rho_0+\rho_\varepsilon}\bar{r}_0-\bar{r}_f)\right)}{(\bar{r}_0-\bar{r}_f)}$$

There are two effects of a larger ρ_ε on the change in optimal portfolio allocation. First, for any given signal s, the agent updates the mean return of the distribution more when ρ_ε is larger, because the posterior of the mean of the return distribution $(\frac{\rho_\varepsilon}{\rho_0+\rho_\varepsilon}s+\frac{\rho_0}{\rho_0+\rho_\varepsilon}\bar{r}_0)$ is a weighted average of the signal and the prior mean distribution. Second, the agent responds more to any changes in the mean of the return distribution because the perceived precision about the risky asset's return is higher for any signal (this is conveyed through the first term $\frac{(\rho_0+\rho_\varepsilon)}{\rho_0}$ of this equation). These two effects imply that responses to signals via trading activity will be larger for an agent who believes that their signal precision is higher.

This prediction is consistent with our findings of both greater buying and selling activity amongst positive return IPO lottery winners. Furthermore, if (as seems plausible) winning a negative return IPO lottery *reduces* the agent's perceived precision, then the model is also able to explain our finding that lottery winners of negative return IPOs reduce their buying and selling.

The model can also explain the attenuation in the marginal response to a lottery in an intuitive

[16]This framework is related to the "hot hands" fallacy (see, e.g., (Gilovich, Vallone, and Tversky, 1985)), where agents mistakenly interpret randomly experienced successes as time-varying skill. Our model here provides a specific structure on what it would mean for an agent to update their beliefs about their skill based on the noise shock of winning the positive return lottery.

fashion–additional perceived signals about the agent's own ability/signal precision following the initial one will factor less into their inferences of their signal precision as they learn more in the usual Bayesian manner. Another positive feature of this model, in contrast with the model in which agents learn about returns, is that it does not necessarily predict that a winner of a positive IPO lottery should increase their allocation to the market overall. If we interpret the model strictly as an alternative to the model in which agents learn about returns, there is no current signal about market returns conveyed by the IPO return; so the implications are solely about trading volume, since ρ_ε will differentially affect the responses only to future signals. Even if we were to relax this strict assumption and allow that winning an IPO lottery with a positive return signal might also increase perceived returns, the agent updating about her own ability is exposed to many new signals over time, and because of their greater precision, will respond to those as well, dampening the effect on the portfolio allocation overall.

We note that this model is not completely able to explain our results, however, as it cannot deliver the tilt towards the sector of the treatment IPO. To explain this feature of the data, we would need an auxiliary assumption such as a warm glow (see, for example, Morewedge et al., 2009; Bordalo et al., 2012) towards this sector following random gains on the IPO, and the reverse following losses. Another important consideration is that in the data, the increases in trading volume caused by the treatment are temporary. If signal precision is perceived to be permanently higher, we would see a long-run level of trading volume that is higher for the treated than for the control group–we do not have evidence of this in the data. To explain this feature of our results, in the model of learning about ability, perceived shocks to signal precision must also be temporary rather than permanent. Put differently, while investors appear to be learning about their own ability from noise, within a few months, they appear to learn that they have made a mistaken inference about their ability.

F.1 Multiple risky assets

Consider an agent maximizing exponential utility $-\exp(-\gamma W)$ of terminal wealth W, with risk aversion coefficient γ.

$$\max_{q} - \exp(-\gamma W) \tag{10}$$

$$\text{s.t.} W = r_f(W_0 - \iota q) + q\tilde{z} \tag{11}$$

The agent has to allocate initial wealth W_0 between a risk-free asset with return r_f, and a risky asset vector with prices normalized to 1, delivering a random payoff vector of $\tilde{z}$. The choice variable is q, the vector of amounts invested in the risky asset. ι is a row vector of ones.

The agent does not know the true distribution of the risk asset payoff vector $\tilde{z}$; he begins with a prior distribution and then updates to a posterior distribution based on signals (return experiences, winning the IPO lottery, etc.). We call the prior distribution of the risky asset payoff $\hat{r}$ and assume that this prior distribution is jointly normal with mean payoff vector $\bar{r}_0$ and variance-covariance matrix Σ_0 (i.e. $\hat{r} \sim N(\bar{r}_0, \Sigma_0)$). Maximizing utility subject to the prior distribution of asset returns, the optimal portfolio allocation vector $q^* = \frac{1}{\gamma}\Sigma_0^{-1}(\bar{r}_0 - r_f) = \frac{1}{\gamma}\Lambda_0(\bar{r}_0 - r_f)$, where $\Lambda_0 = \Sigma_0^{-1}$ is the precision matrix of the agent's priors about risky asset returns.

F.1.1 Learning about the distribution of market returns

The first case that we consider is that agents randomly experiencing gains or losses interpret this as a signal about returns, whereas lottery losing agents that do not experience these returns act as if they receive no signal. We denote the signal by $\tilde{s} \sim N(\tilde{r}, \Sigma_\varepsilon)$, where $\Sigma_\varepsilon = \Lambda_\varepsilon^{-1}$ is a matrix of noise variances and covariances associated with incoming signals.

If this were the case, following the IPO win, treated agents would use Bayes' rule to calculate the posterior distribution of returns. Σ_S denotes the posterior variance-covariance matrix of returns, and $\tilde{r}_S$ the posterior mean returns:

$$\Sigma_S = \Sigma_0(\Sigma_0+\Sigma_\varepsilon)^{-1}\Sigma_\varepsilon, \Lambda_S = \Sigma_S^{-1}. \tag{12}$$

$$\tilde{r}_S = \Sigma_0(\Sigma_0+\Sigma_\varepsilon)^{-1}\tilde{s} + (\Sigma_0+\Sigma_\varepsilon)^{-1}\Sigma_\varepsilon\bar{r}_0. \tag{13}$$

Agents would adjust their portfolio to $q_S^* = \frac{1}{\gamma}\Lambda_S(\tilde{r}_S - r_f)$. We can generically consider two potential types of signals. The first is that agents perceive the randomly experienced gains as a sign of returns on the market. Equivalently, this is a rise in the expected returns of all stocks, i.e., all elements of $\tilde{r}_S - r_f$ are greater than the corresponding element in $\bar{r}_0 - r_f$. The second type of perceived signal is more restricted, to particular stocks (say those in the sector of the treatment IPO). In this case, some elements of $\tilde{r}_S - r_f$ are greater than the corresponding element in $\bar{r}_S - r_f$, while other elements are equivalent in the two excess return vectors.

In the first case, i.e. a perceived rise in the expected return of all stocks, all elements of q_S^* should be greater than the corresponding elements of q^* as a result of the mean return expectation on risky assets increasing. Moreover, since the agents' perceived posterior total precision Λ_S is larger following the signal, this provides an additional upward kick to q_S^*.

In the second case, i.e., a rise in the perceived expected return of some stocks, some elements of q_S^* should be greater than the corresponding elements of q^*, and the agent's perceived posterior precision Λ_S is also greater, albeit for only some elements of the matrix.

In both cases, the model yields counterfactual predictions. This is for two reasons. First, as we know, the average allocation to risky assets for the treated does not rise — and this is predicted by both types of signals about returns. It is true that a more restricted perceived signal does help to explain the tilt towards particular sectors, but without a corresponding reduction in the means of other sectors, or indeed an auxiliary assumption such as mental accounting, it is difficult to rationalize the observed lack of an increase in the overall portfolio weight to risky assets. The second counterfactual prediction is perhaps more troubling. Since the agent's posterior precision Λ_S is higher following the perceived signal, we should observe that treated agents' portfolio weights

q_S^* react *less* to any *future* incoming signals of either direction, assuming that the post-IPO signal distribution is balanced across treated and control agents.[17] This would predict a *drop* in future trading volume for the treated agents, regardless of whether the initial signal was perceived to be positive or negative. This runs strongly counter to our empirical results.[18]

F.1.2 Learning about own ability

We next consider the case in which the agent perceives the randomly experienced gains or losses from the IPO as a positive or negative shock to their own investment ability.

To do so, we assume that a sensible measure of the agent's ability is the precision of the signal distribution, i.e., Λ_ε. The greater the signal precision, the greater the agent's ability. This is a common assumption in the literature on stock trading (see, for e.g., Gervais and Odean 2001, and Linnainmaa, 2011).

To model agents drawing inferences about their own ability, we consider agents with priors over $\tilde{\Lambda}_\varepsilon \sim \text{Wishart}(\nu, \Lambda_\varepsilon^0)$, where ν is the number of observations about own ability that the agent has previously experienced. If the treated agent updates their prior about their own ability in response to a randomly experienced gain (essentially, treating experienced luck as a signal of their own skill), then the agent will update perceived signal precision as $\tilde{\Lambda}_\varepsilon^S$, with some of its elements increasing or decreasing with the randomly experienced IPO return, and posterior ν^S incremented as well.

The model yields a set of predictions that are closer to what we observe in the data. First, in response to an update to perceived signal precision, the model-predicted average allocation to risky assets for the treated does not rise, because there is no change in this model to the parameters of the currently perceived return distribution. Second, since the treated agent's posterior precision of incoming signals is higher, we should observe that treated agents' portfolio weights q_S^* react *more* to future incoming signals of either direction, once again assuming that the post-IPO signal distribution

[17] Intuition is provided by the single risky asset case. If the treated agent's posterior precision is ρ_S, and the control agent's is ρ, and $\rho_S > \rho$, any incoming signal with precision ρ_ε about returns will be weighted $\frac{\rho_\varepsilon}{\rho_S} < \frac{\rho_\varepsilon}{\rho}$ by the treated agent, affecting the distribution of portfolio weights less. This is the standard result about additional informational updates affecting posteriors marginally less.

[18] We note also that any interpretation of the IPO lottery as a shock to the precision of the market return distribution also delivers the same (counterfactual) implication that the winners will be less responsive to all future signals.

is balanced across treated and control agents.[19] This would predict a *rise* in future trading volume for the treated agents, with randomly experienced gains, and a *drop* for randomly experienced losses. This explains an important feature of our empirical results. Finally, the model can also explain the attenuation in the marginal response to a lottery in an intuitive fashion — additional perceived signals about the agent's own ability/signal precision following the initial one will factor less into their inferences of their signal precision as they learn more in the usual Bayesian manner.[20]

That said, however, this model does not completely explain our results, as it cannot deliver the tilt towards the sector of the treatment sector IPO. To explain this result, we will require recourse to an auxiliary assumption such as a warm glow towards this sector following random gains on the IPO and the reverse following losses.

Another important consideration here is that in the data, the increases in trading volume caused by the treatment are temporary. If signal precision is perceived to be permanently higher, we would see a long-run level of trading volume that is higher for the treated than for the control group, which we do not. To explain this feature of our results, in the model of learning about ability, perceived shocks to signal precision must also be temporary. Put differently, while investors appear to be learning about their own ability from noise, within a few months, they appear to learn that they have made a mistaken inference about their ability.

G Heterogenous Treatment Effects: Repeated Experiences and Responses to Noise

Each column of Table A.4 focuses on a different portfolio characteristic that we believe might modulate how previous experiences affect the treatment effect of winning the lottery. To explain this table we focus on Column (1); the analyses in Columns (2) and (3) are analogous.

[19]Again, intuition is provided by the single risky asset case. If the treated agent's prior precision is ρ, same as that of the control agent, and *incoming* signal precision ρ_ε^S is greater for treated agents than the ρ_ε of control agents new signals about returns will be weighted $\frac{\rho_\varepsilon^S}{\rho} > \frac{\rho_\varepsilon}{\rho}$ by the treated agent, affecting the distribution of weights more.

[20]Intuition for this is available by considering the univariate case with a Gamma distributed signal precision—the conjugate distribution is Gamma, with the posterior shape parameter driving the nonlinear rate of convergence to the truth in the presence of incoming shocks with constant effects on the scale parameter. Put differently, increments to ν drive the nonlinear rate of learning about signal precision in the model.

To explain this analysis it is useful to work our way up to the heterogeneous treatment we are interested in. The main result of our paper is the effect of winning the lottery on future trading volume; this is the coefficient on I(Allot) in Column (1) of Online Appendix Table A.4. The current section of the paper focuses on how past experiences of winning IPO lotteries change how an investor responds to winning a *subsequent* lottery. This relationship is indicated by the coefficient on I(Allot) $\times$ RandWins(All) in Column (1); this coefficient tells us how much the effect of winning the lottery decreases for investors with higher numbers of lottery wins in the past.

We now want to test whether this decrease in the I(Allot) effect that comes from increasing RandWins(All) is larger for young accounts versus old accounts. To assess this, we need to look at the triple interaction term I(Allot) $\times$ Age $\times$ Randwins(All) in a specification that also has all of the relevant double interaction terms. The coefficient on the triple interaction term can be used to calculate two versions of the I(Allot) $\times$ Randwins(All) interaction; one version for young accounts and one version for old accounts. We then test whether these two interactions are statistically different from each other.

To do this calculation (in the same way as in Figure 3), we define "young" accounts at the mean of the first tercile bin of the account age distribution (measured in months), "middle-aged" accounts as the mean of the second tercile bin of the account age distribution, and "old" accounts as the mean of the third tercile bin of the account age distribution, the month before the IPO lottery, also reported in Column 1, Figure 3. We define high experience accounts as those at the 75th percentile of number of Randwins(All) (equivalent to two prior IPO experiences) and low experience accounts as those at the 25th percentile (equivalent to zero previous IPO experiences. Online Appendix Table A.5, Panel A, gives these values for the relevant calculations for Columns (1)-(3) of Online Appendix Table A.4.

Panel B.1 of Appendix Table A.5 shows how account age moderates the relationship between the treatment effect of winning the IPO lottery and past IPO experiences. Take the first column "Coefficient Estimate". Here, Young(2-0) indicates the change in the treatment effect for a young account that moves from having 2 previous IPO experiences to having 0 previous IPO experiences;

this turns out to be −8.8 percentage points (in other words, a young account with 2 previous experiences has a winning the lottery effect on trading volume that is 8.8 percentage points lower than young account with 0 previous experiences). Similarly, the Old(2-0) term is the difference in treatment effects for old accounts with 2 experiences versus old accounts with 0 experiences; the difference here is −0.02 percentage points.

The "Difference" column gives Young(2-0)-Old(2-0), which in this case is −0.0659.[21] So in the case of account age, young accounts' responses are affected more by the number of past experiences than old accounts, in the sense that the treatment effect of winning the lottery reduces with experience more for young accounts than old accounts. The remaining two rows of Panel B.1 do the same comparisons for Young vs. Middle-aged and Middle-aged vs. Old accounts. These rows also show that younger accounts, in general, are affected more by the number of past experiences than older accounts in terms of reductions in treatment effects.

Panels B.2 and B.3 do the same calculations, but substitute portfolio value and trading volume (respectively) for account age. The results suggest that the treatment effect of winning the lottery for investors with smaller portfolios is affected less by previous experiences than the corresponding treatment effect for investors with larger portfolios. Panel B.3 suggests that the treatment effects for low trading volume accounts are affected more by previous IPO experiences than treatment effects for high trading volume accounts.

[21] The coefficients are obtained as a linear combination of coefficients reported in Online Appendix Table A.4. The standard error of the linear combination is calculated as the square root of the variance of the linear combination using the robust variance-covariance-matrix from the regression reported in Appendix Table A.4

References

Bordalo, P., Gennaioli, N., Shleifer, A., 2012. Salience in experimental tests of the endowment effect. The American Economic Review 102 (3), 47–52.

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

Chamley, C., 2004. Rational herds: Economic models of social learning. Cambridge University Press.

Gervais, S., Odean, T., 2001. Learning to be overconfident. Review of Financial Studies 14 (1), 1–27.

Gilovich, T., Vallone, R., Tversky, A., 1985. The hot hand in basketball: On the misperception of random sequences. Cognitive Psychology 17 (3), 295–314.

Linnainmaa, J. T., 2011. Why do (some) households trade so much? The Review of Financial Studies 24 (5), 1630–1666.

Morewedge, C. K., Shu, L. L., Gilbert, D. T., Wilson, T. D., 2009. Bad riddance or good rubbish? Ownership and not loss aversion causes the endowment effect. Journal of Experimental Social Psychology 45 (4), 947–951.

Sirri, E. R., Tufano, P., 1998. Costly search and mutual fund flows. The Journal of Finance 53 (5), 1589–1622.