Navigating Through Fear and Greed: The Experience-Driven Disposition Effect *

Tai Lo Yeung Rong Liu Jessica Wachter Michael Kahana

Yongjie Zhang†

August 9, 2026

Click here for the latest version

Abstract

Using transaction-level data on 189,530 Chinese retail investors over the 2013–2016 boom and bust, we measure experience as counts of large realized gains and losses and find that it reshapes the disposition effect, the tendency to sell winners and hold losers, in an asymmetric way. A two-standard-deviation increase in accumulated large losses raises the propensity to sell winners rather than losers by about 23% of the baseline effect, while an equivalent increase in gains lowers it by about 21%. Composition-matched simulations show the asymmetry is not a mechanical artifact of building experience from past trades. To explain both the baseline effect and this asymmetry, we build a memory-based recall model in which similarity-weighted retrieval of past outcomes guides selling. A single retrieval channel reproduces all three patterns without

*Presented at the Memory, Beliefs, and Choice workshop, the 2026 SFS Cavalcade, and the 2026 Workshop on Cognitive Economics. We thank our discussants Stefano Cassella and Ian Krajbich for their thoughtful comments, and Alex Imas, Lorenz Kueng, Sean Myers, Cameron Peng, and Nikolai Roussanov for helpful suggestions. Liu and Zhang acknowledge funding from the National Natural Science Foundation of China (72342022 and 72271184). Kahana and Wachter gratefully acknowledge research support from National Institutes of Health Grant MH55687 for the modeling work.

†Yeung: USI Lugano, tai.lo.yeung@usi.ch, Liu: Tianjin University, lrong@tju.edu.cn, Wachter: University of Pennsylvania, jwachter@wharton.upenn.edu, Kahana: University of Pennsylvania, kahana@psych.upenn.edu, Zhang: Tianjin University, yjz@tju.edu.cn.

shifting preferences.

Keywords: Retail trading, disposition effect, investor experience, memory, cued recall

JEL Codes: D14, D91, G11, G14, G41

1 Introduction

Between July 2013 and February 2016, the Chinese stock market more than doubled and then shed over 40% of its value in less than a year. We use this boom and bust, together with complete within-brokerage trading records for 189,530 retail investors at one of China's largest brokerages, to ask how the gains and losses an investor accumulates reshape the disposition effect, the tendency to sell winners and hold losers (Shefrin and Statman, 1985; Odean, 1998). A long line of field studies reads this bias as fading with experience and takes that fading as evidence that investors grow more sophisticated as they trade (Feng and Seasholes, 2005; Seru, Shumway and Stoffman, 2010). We find this is only half true: experience attenuates the disposition effect when it arrives as gains but amplifies it when it arrives as losses. Investors who have absorbed large losses do not learn to cut losses; they hold losers and realize winners more aggressively than before.

The asymmetry is economically large. We measure experience as an investor's cumulative count of realized outcomes, trades closed more than ±5% from the average purchase price, and estimate how the propensity to sell a winner rather than a loser moves with the counts of past gains and past losses. A two-standard-deviation increase in accumulated losses raises that propensity by 0.57 percentage points, about 23% of the 2.48-percentage-point baseline disposition effect; an equal increase in gains lowers it by 0.51 points, about 21%. Both estimates are precise ($t = 3.3$ and -3.3) and hold under investor, stock, and time fixed effects, with standard errors clustered on all three dimensions.

To organize these facts, we develop a memory-based model of the selling decision that adapts the temporal-context framework of Howard and Kahana (2002) to an investor's portfolio. The investor stores each holding as an episode, a stock paired with the outcome she experiences while owning it; when she decides what to sell, the cue "stocks I own" retrieves these episodes in proportion to their similarity and frequency, each competing with the others for recall. Large moves occur almost only in stocks and, for stocks, are positively skewed (Oh and Wachter, 2022), so a large gain is what most readily comes to mind, and the investor sells the winner that does. This produces the baseline disposition effect. Realizing a gain re-tags that outcome from owned to sold and weakens the link, so the bias fades after a run of gains; realizing a loss reinforces the association between ownership and surviving gains, so the bias sharpens after losses. One retrieval channel thus generates the baseline effect, its decay after gains, and its amplification after losses, with no change in preferences.

Because experience is built from an investor's own past sales, one might worry that the interactions are a mechanical by-product of how a portfolio turns over. They are not. On the loss side, the net mechanical force implied by how portfolios actually evolve is negative—the opposite sign from our estimate—so a positive loss-side coefficient cannot be an accounting artifact (Section 5). The effect also survives a range of competing explanations, from unobserved skill and rebalancing to return-extrapolation and overtrading accounts.

Three contributions follow. First, the content of experience, not its quantity, governs the disposition effect. A field literature proxies experience by the number of trades and reads the fading bias as growing sophistication (Feng and Seasholes, 2005; Seru, Shumway and Stoffman, 2010); once we split realized outcomes by valence, gains and losses move the bias in opposite directions, and losses amplify it rather than tame it. The finding also runs against a belief-pessimism reading, under which painful losses would make investors quicker to realize losses and so weaken the bias; we find the reverse, which is difficult to reconcile with a simple pessimistic-belief account and consistent with memory.

Second, we give the disposition effect a memory microfoundation. Adapting the temporal-context model of Howard and Kahana (2002) and the memory-and-decisions framework of Malmendier and Wachter (2021), we show that one cued-recall channel reproduces the baseline effect and both experience patterns.

Third, we take the mechanism beyond the facts that motivated it. The same channel delivers the V-shaped selling propensity of Ben-David and Hirshleifer (2012): large returns of either sign are the most retrievable, so the most extreme positions are the ones most likely to be sold. The rank effect of Hartzmark (2015) points the same way: the best- and worst-performing positions in a portfolio are the ones most likely to be sold, and they are also the most retrievable. Because diversification strips out the right tail on which retrieval feeds, the mechanism predicts a weaker disposition effect for mutual funds; Chang, Solomon and Westerfield (2016*a*) show that the same investors who exhibit a positive disposition effect in individual stocks exhibit a negative one in mutual funds. The channel also fits direct evidence on investor memory: recalled returns move with market conditions (Jiang et al., 2024), and the causal result of Frydman and Wang (2020) shows that making the purchase price more prominent on a brokerage screen raises the disposition effect.

The paper proceeds as follows. Section 2 describes the brokerage data, the 2013–2016 market cycle, and how we measure winning and losing experiences. Section 3 presents the core result, the opposite-signed response of the disposition effect to gains and losses, and shows it survives alternative thresholds and a ratio specification. Section 4 develops the memory-based recall model and maps it to the evidence. Section 5 rules out alternative mechanisms: a mechanical artifact of portfolio composition, unobserved skill and portfolio-level effects, mean reversion, return extrapolation, portfolio rebalancing, and overconfident overtrading. Section 6 takes the mechanism to independent variation: the experience effect is stronger among younger investors and weaker for mutual funds. Section 7 concludes.

2 Data and Methodology

2.1 Retail Investment Data

We draw on transaction-level records from a major Chinese securities brokerage spanning early July 2013 to late February 2016. The unit of observation is a single trade: over 40 million transactions by 189,530 individual retail investors trading roughly 4,000 stocks. Table 1 reports investor descriptive statistics. Each record carries the number of shares bought and sold, the purchase and sale prices, and the prevailing market price. Linking records to account IDs adds investor demographics—age, gender, and education.[1] The brokerage also records each investor's mandatory risk-tolerance assessment, completed at account opening.

[1]Under China's real-name registration system, each securities account is tied to a single national identity card, and an investor cannot hold more than one account at the same brokerage. The account we observe therefore reflects the investor's complete holdings at this broker—the relevant unit for measuring the disposition effect and portfolio composition. An investor may in principle maintain a separate account at another brokerage. Unobserved outside positions and trades introduce measurement error both in account-wide portfolio variables and in our experience counts, which capture within-brokerage rather than total personal trading history; if this error is approximately classical, it attenuates the estimated experience effects toward zero.

Figure 1: Chinese Bull and Bear Market, 2013–2016

Notes. This figure illustrates the performance of the SSE Composite Index over a specified period. The green area represents the period covered by our data, while the gray area marks the period identified as a market crash. The SSE Composite Index, which includes all A-shares and B-shares listed on the Shanghai Stock Exchange, reflects the overall performance of the Chinese stock market. The horizontal axis is the calendar date (July 2013 to February 2016) and the vertical axis is the daily closing index level. Source: Shanghai Stock Exchange.

2.2 Bull and Bear Markets

Figure 1 places our sample period in market context, plotting the Chinese stock market over the span of our data. Unlike datasets confined to bull markets, our window covers bull, bear, and low-volatility regimes. Investors in our sample therefore face widely varying odds of realizing gains, losses, or a mix of the two—the cross-investor variation our design needs.[2]

[2]Many of the most relevant studies draw on predominantly bull-market periods—for example, Feng and Seasholes (2005) in China and Barber and Odean (2000*a*) in the United States both cover sustained upswings—whereas our 2013–2016 window spans a full boom and bust, with the index rising from about 2,000 to 5,000 before falling back below 3,000.

2.3 Constructing the Key Variables

We build our panel at the investor–stock–day level: on each date an investor sells at least one security, we record the position sold along with every other position the investor holds that day. The panel runs to 43,649,867 investor–stock–day observations, the sample we take to the regressions. Dropping accounts opened before our first observation date, July 1, 2013, restricts the sample to accounts whose complete trading history we observe; average account tenure is 340 days, measured from account opening through the last observable date, February 29, 2016. We describe each step below.

2.3.1 Portfolio Positions and Investment Gains and Losses

Positions. We sort the data by account ID, then stock ticker, then date. Following Odean (1998) and subsequent disposition-effect studies (Feng and Seasholes, 2005; Seru, Shumway and Stoffman, 2010), we define a position as beginning when an investor first purchases a given stock and ending at the first sale of that stock, and we treat any shares remaining after a partial sale as the start of a new position.[3] An investor can therefore reinvest in the same stock many times within a position; when this happens, we set the position's cost basis to the volume-weighted average of all purchase prices, so the realized return at sale is measured against that average. Table 1 reports descriptive statistics for the holdings and trading experiences. The average investor holds 7.6 stocks, and portfolio values are strongly right-skewed.

Gains and Losses. Following Odean (1998), we evaluate a portfolio's gains and losses only on dates when the investor sells some security, because on no-sale days we cannot tell deliberate holding from inattention (Chang, Solomon and Westerfield, 2016*b*). For each investor–stock–day we compare the relevant price to the cost basis and classify the position as a gain (Gain = 1) when the return is positive and a loss (Gain = 0) otherwise. For the stock actually sold, we use the realized return computed at the sell order price; for every other stock held that day, we use the unrealized ('paper') return computed at the day's

[3] We find similar results if partial sales are excluded from our sample.

Table 1: Descriptive Statistics of Investors, Holdings, and Trading Experiences

	Mean	SD	Percentiles 25th	50th	75th
Investors and Accounts					
Number of investors	189,530	-	-	-	-
Account tenure (days)	339.90	162.00	266	308	362
Number of account-stock-days	43,649,867	-	-	-	-
Percentage of male investors (%)	55%	-	-	-	-
Average investor age (years)	37.88	11.06	29	36	45
Risk-taking level (1-5)	3.05	0.74	3	3	3
Education level (1-7)	4.11	1.28	3	4	5
Holdings					
Number of stock holdings	7.59	9.24	2.27	4.83	9.52
Portfolio value (RMB)	358,875	3,357,000	19,467	63,472	206,929
Trading Experiences					
Losing experiences	12.40	28.30	0	4	13
Losing experiences (5%)	6.45	13.39	0	2	8
Losing experiences (30%)	0.68	1.92	0	0	1
Winning experiences	14.79	35.95	1	5	16
Winning experiences (5%)	5.66	12.90	0	2	7
Winning experiences (30%)	0.28	1.49	0	0	0

Notes. Risk-taking level 1 represents the most risk-averse and 5 the most risk-seeking investors. Education level 1 represents the highest level of education (PhD degree) and 7 represents the lowest level of education (elementary school level). The number of stocks held is calculated at the trading day level, reflecting the average number of stocks an investor has during selling transactions. Portfolio value is calculated using the same approach. Losing and winning experiences represent the number of respective experiences an investor had. RMB stands for RenMinBi, the currency of mainland People's Republic of China (PRC). Trading experiences are measured by the number of realized returns, where each negative realized return is counted as one losing experience. The 5% and 30% experience measures count only prior trades whose realized returns (in absolute value) exceed 5% and 30%, respectively, relative to the investor's average purchase price. The sample period is July 2013 to February 2016, covering 189,530 investor accounts; the number of account-stock-days reported here is the same quantity used as the number of observations in the regression analysis. Demographic variables (gender and education) are missing for some accounts (N = 128,563 and 122,452, respectively), so heterogeneity analyses involving these characteristics run on the corresponding sub-samples. The trading-experience and portfolio-value distributions are heavily right-skewed; standard deviations exceed means due to a small hyperactive tail (e.g., the maximum Losing Experience is 2,563).

opening price.[4,5]

A losing experience is a completed sale that realizes a return below a threshold relative

[4] We use the opening price because it is predetermined relative to the day's sale decisions; the closing price is not.

[5] Details of our per-transaction return calculations are provided in Appendix A1.

Figure 2: Distribution of Trading Experience

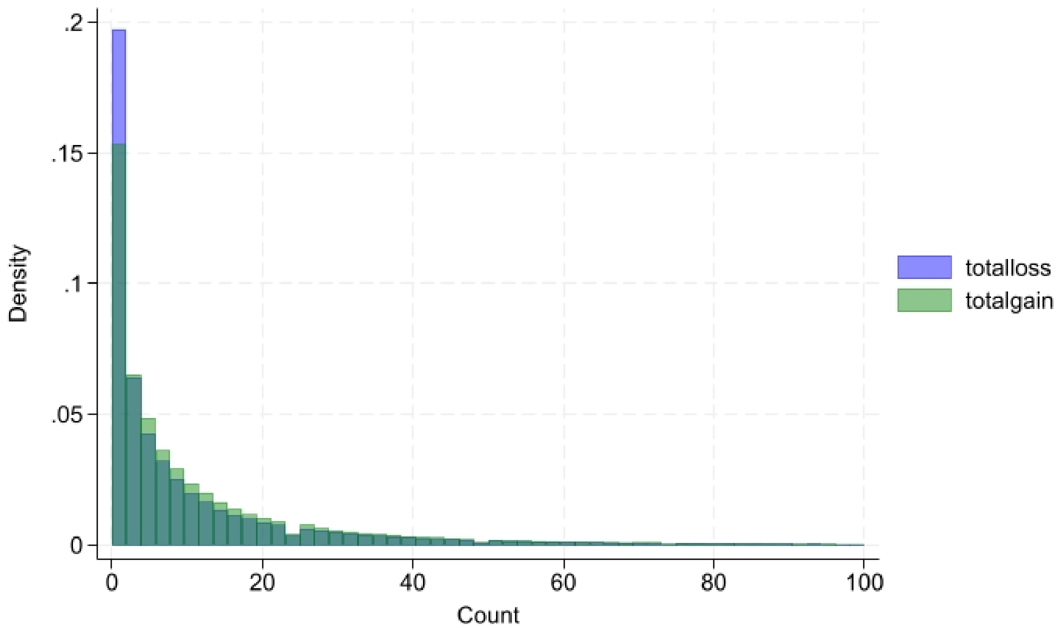

Notes. This figure shows the distribution of trading experiences across accounts in our sample. The green bars represent each account's total number of winning experiences ("totalgain"), while the purple bars reflect the total number of losing experiences ("totalloss"). These counts are measured cumulatively as of the end of the observation period (i.e., at time T). Although trading experience is modeled as a time-varying variable in our empirical analysis, the figure presents a static snapshot of each account's experience as of the final period. The horizontal axis is the number of trading experiences and the vertical axis is the number of accounts.

to the investor's average purchase price, and a winning experience one above it; an investor's losing and winning experiences are the running counts of each. We vary the threshold from 5% to 30%; results are robust across these cutoffs. Table 1 shows that investors record 12 losing and 15 winning experiences on average, the gap reflecting a sample that spans somewhat more bullish than bearish months.

Figure 2 plots the account-level distribution of winning and losing experiences. At exactly one experience, accounts with a single loss far outnumber those with a single win—many investors exit after a first loss, leaving accounts stranded with one losing experience and nothing more. Above that point, the counts of accounts track each other closely across comparable levels of winning and losing experience.

2.3.2 Measuring Trading Experience

Measuring trading experience raises a measurement problem of its own. Account-level data are precise and free of the self-report error that survey measures carry, but they do not tell us how investors themselves perceive their experience. Each measure we use therefore embodies an assumption about that perception: the trade-count measure equates experience with activity; our valence-split measures equate it with realized outcomes.

We begin from the trade-count measure of Dhar and Zhu (2006), which proxies experience by accumulated trading activity.[6] Dhar and Zhu (2006) reads this measure as learning: List (2003) find experimentally that repetition breeds familiarity, so traders who act more often grow more sophisticated and less prone to bias, and Dhar and Zhu (2006) carries the same logic to investors. Both report that a higher cumulative trade count predicts weaker bias.

We treat this measure as a foil. Neither List (2003) nor Dhar and Zhu (2006) can separate a winning trade from a losing one—a distinction the sports-card setting of List (2003) cannot support—so each count pools outcomes of opposite valence. We show instead that past wins and past losses act on the current sell decision through distinct, opposing channels. We first replicate the pooled trade-count result, then split experience by outcome to show how past realized outcomes bear on the current decision.

Distinguishing Win and Loss Experiences. We start from the pooled count above, then separate experiences by valence, defining:

$$\text{Loss Experience}_{i,t} = \sum_{T=0}^{t-1} \mathbf{1}\{\text{return}(s,j)_T < -X\}$$

$$\text{Win Experience}_{i,t} = \sum_{T=0}^{t-1} \mathbf{1}\{\text{return}(s,j)_T > X\}$$

where $\mathbf{1}\{\cdot\}$ denotes a loss (win) experienced by investor i at time T for stock j, and X is the threshold above which a realized gain or loss counts as an experience. We vary X for robustness and adopt $X = 5\%$ and $X = 30\%$ in our primary analysis.

[6] The construction of trading episodes and the associated experience measures is described in detail in Appendix A2.

Because more active investors mechanically accumulate more losses in absolute terms, we also scale by total activity, using the share of losing experiences among all experiences:

$$\text{Loss Experience Share}_{i,t} = \frac{\sum_{T=0}^{t-1} \mathbf{1}\{\text{return}(s,j)_T < -X\}}{\text{Number of prior sales}_{i,t}} \tag{1}$$

A simple example: when $X = 0$, if an investor has 15 selling transactions before making the current decision and 5 of them have positive realized returns, then Loss Experience Share$_{i,t}$ = 2/3.

2.4 Measuring the Disposition Effect

We measure the disposition effect two ways: an investor-level proportion and a position-level regression coefficient. Following Dhar and Zhu (2006), we start with the simplest investor-level measure. Following Odean (1998), we define the Proportion of Gains Realized (PGR) and the Proportion of Losses Realized (PLR) as

$$\text{PGR} = \frac{\text{Realized Gains}}{\text{Realized Gains} + \text{Paper Gains}}$$

$$\text{PLR} = \frac{\text{Realized Losses}}{\text{Realized Losses} + \text{Paper Losses}}$$

$$\text{DE} = \text{PGR} - \text{PLR} \tag{2}$$

Realized Gain (Loss) counts the stocks sold in a portfolio at a positive (negative) realized return; Paper Gain (Loss) counts the stocks held at a positive (negative) unrealized return at the time of calculation. Equation 2 is the standard measure of the disposition effect's magnitude: DE > 0 signals the disposition effect, a higher propensity to realize gains than losses.

We next quantify that magnitude with the regression approach of Odean (1998), regressing a stock-position sell indicator (1 = Sell; 0 = Hold) on a *Gain* indicator (1 = the stock trades at an unrealized gain; 0 otherwise) among other controls. Its coefficient measures how much more likely investors are to sell a winner.

3 Experience-Driven Disposition Effect

This section takes the memory model's three qualitative predictions to the data: a disposition effect already present at zero measured large-outcome experience, its attenuation as realized gains accumulate, and its amplification as realized losses accumulate. Section 4 derives all three from a single retrieval channel. The first is plain in the raw data: across all paper returns recorded on sale dates, investors realize 12.4% of gains but only 8.3% of losses (9.7% overall), in line with Odean (1998) and subsequent studies.

We estimate a linear probability model (LPM) with high-dimensional fixed effects, following An et al. (2024). Investor, stock, and time fixed effects absorb unobserved heterogeneity across investors, persistent differences across stocks, and common time variation such as the boom–bust cycle. We cluster standard errors three ways—by investor, day, and stock (Cameron, Gelbach and Miller, 2011)—letting the sale residuals correlate arbitrarily within an investor, within a day, and within a stock. Because experience varies only at the investor level, investor clustering binds for the experience interactions: collapsing to two-way (investor and day) or one-way (investor) clustering leaves their standard errors essentially unchanged (untabulated).

3.1 A Benchmark: Trading Experience and the Disposition Effect

We begin by confirming that our data reproduce the benchmark pattern in Feng and Seasholes (2005), estimating a regression close to theirs[7]:

$$\text{Sale}_{i,j,t} = \alpha + \beta_1 \cdot \text{Gain}_{i,j,t} + \beta_2 \cdot \text{Experience}_{i,t} + \beta_3 \cdot \text{Gain}_{i,j,t} \times \text{Experience}_{i,t} + \epsilon_{i,j,t} \qquad (3)$$

$\text{Sale}_{i,j,t}$ equals 1 if stock j in account i is sold, in whole or in part, at time t, and 0 otherwise, so the fitted value is investor i's willingness to sell asset j at time t. $\text{Gain}_{i,j,t}$ equals 1 if the market price of asset j exceeds its average purchase price at time t (a positive unrealized return), and 0 otherwise, so β_1 measures the propensity to sell a winner. We measure

[7] While they employ a hazard model, we use linear probability model (LPM) in our analysis to demonstrate that we observe the same pattern in our results.

Table 2: Trading Experience and Disposition Effect

Dependent Variable:	**Sale × 100** (1)	(2)
Gain	2.5037***	2.7506***
	(4.74)	(4.77)
Total Experience	0.0025*	0.0029**
	(1.83)	(2.06)
Gain × Total Experience		-0.0020**
		(-2.58)
Constant	8.5853***	8.5345***
	(34.88)	(33.44)
Observations	43,649,867	43,649,867
R-squared	0.139	0.139

Notes. The table reports results for the linear probability model shown in equation 3 on our sample. *Total Experience* is the cumulative number of prior trades (wins or losses) executed by the investor before the current sell decision. Column (1) omits the *Gain* × *Total Experience* interaction; Column (2) adds it, capturing how the disposition effect varies with overall trading experience. The outcome dummy is multiplied by 100 to improve the readability of the coefficients. Both columns include investor, stock, and time fixed effects, and standard errors are clustered by investor, time, and stock. The sample period is July 2013 to February 2016. Numbers in parentheses are t-statistics. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

experience as

$$\text{Experience}_{i,t} = \text{Number of sales taken by investor } i \text{ up until date } t.$$

The coefficient of interest is β_3 on the interaction $\text{Gain}_{i,j,t} \times \text{Experience}_{i,t}$. Because *Gain* is binary, experience shifts the sale probability by β_2 for a losing position and by $\beta_2 + \beta_3$ for a winning one, so β_3 is the effect of experience on the disposition effect itself: a positive β_3 means experience amplifies the bias, a negative β_3 that it tames it.

We report the estimates in Table 2. Column (2) shows that the coefficient on *Gain* is positive and significant, confirming the disposition effect among retail investors, while the interaction between *Gain* and *Total Experience* is negative and significant, so the disposition effect attenuates as investors accumulate experience.

The disposition effect thus weakens as investors accumulate trading experience, consistent with the pattern documented by Feng and Seasholes (2005).[8]

[8] Feng and Seasholes (2005) additionally allow this effect to be non-linear, using dummies for bins of accumulated trades (fewer than five prior trades, fewer than ten, and so on). Because we separate winning from losing experiences and require each realization to clear a return threshold, our experience counts are far lower than theirs, leaving a binned specification little power to detect non-linearity; we therefore do not pursue it.

Table 3: The Experience-Driven Disposition Effect

Dependent Variable:	**Sale × 100**				
$\|X\| > 5\%$	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**
	Main Effects		**Interactions**		**Full Spec.**
Gain	2.4936***	2.4963***	2.4765***	2.6628***	2.4781***
	(4.74)	(4.74)	(4.34)	(4.79)	(4.35)
Loss Experience	0.0163***		0.0162***		0.0109
	(3.17)		(3.12)		(1.13)
Win Experience		0.0087		0.0103	0.0034
		(1.31)		(1.50)	(0.36)
Gain × Loss Experience			0.0007		0.0212***
			(0.30)		(3.28)
Gain × Win Experience				-0.0068***	-0.0198***
				(-2.66)	(-3.30)
Constant	8.4347***	8.6229***	8.4384***	8.5822***	8.4974***
	(37.46)	(35.03)	(36.15)	(33.86)	(36.63)
Observations	43,649,867	43,649,867	43,649,867	43,649,867	43,649,867
R-squared	0.138	0.138	0.138	0.138	0.138

Notes. Estimates from the linear probability model in equation 3; the dependent variable is an indicator for sale, multiplied by 100. *Loss Experience* and *Win Experience* count the investor's prior trades that realized a loss or a gain exceeding **5%** relative to the average purchase price. Columns (1)–(2) include only the main effects for *Loss Experience* and *Win Experience*; columns (3)–(4) add interactions between *Gain* and each count; column (5) reports the full specification with both counts and both interactions. *Gain* equals 1 if the position has an unrealized gain at the decision time. All specifications include investor, stock, and time fixed effects, and standard errors are clustered by investor, time, and stock. The sample period is July 2013 to February 2016. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

3.2 Trading Experience with Both Wins and Losses

Total trade count, however, masks the valence of experience. The same number of trades can reflect a history of mostly gains, mostly losses, or a balance of the two, and there is little reason these should enter the current sell decision in the same way; we therefore separate winning from losing experience. We define a *losing (winning) experience* as a sale that realizes a return below -5% (above $+5\%$) relative to the investor's average purchase price. Our results are robust to alternative cutoffs[9].

We split prior experience by valence in Table 3. The baseline disposition effect holds across every column: the coefficient on *Gain* is positive and significant at the 1% level, so investors sell winners more readily than losers. Winning and losing experience move that tendency in opposite directions.

Column (5) reports the full specification. The interaction of *Gain* with *Loss Experience* is

[9] We test the robustness of our results using alternative thresholds of 5%, 10%, 15%, 20%, 25%, and 30%. For clarity of presentation, we primarily report results based on the 5% threshold, while also reporting the 30% threshold to capture only the largest outcomes.

positive and significant at the 1% level (0.0212, $t = 3.28$): the disposition effect strengthens as losing experiences accumulate. A two-standard-deviation increase in *Loss Experience* (26.78) raises the propensity to sell a winner rather than a loser by 0.57 percentage points—about 23% of the 2.48-percentage-point baseline *Gain* effect.

The interaction of *Gain* with *Win Experience* runs the other way, negative and significant at the 1% level (-0.0198, $t = -3.30$): the disposition effect weakens as winning experiences accumulate. An equal increase in *Win Experience* (25.80) cuts the relative propensity to realize winners by 0.51 percentage points, roughly 21% of the baseline. In per-trade terms, each additional large realized loss widens the winner–loser sale gap by about 0.02 percentage points and each additional large gain narrows it by the same amount, just under 1% of the baseline gap apiece.

Two features of the column progression matter. First, the attenuation is already visible when the win interaction enters alone (column (4): -0.0068, $t = -2.66$), while the loss interaction is small and insignificant on its own (column (3): 0.0007, $t = 0.30$); entering both jointly sharpens the valence-specific asymmetry. Second, either channel is an order of magnitude larger than the attenuation the pooled count delivered in Table 2 (-0.0020): pooling nets two opposing forces against each other and understates both.

Past outcomes thus enter the current sell decision asymmetrically—losing experiences reinforce the urge to lock in gains, winning experiences temper it.

3.3 Varying the Threshold for Wins and Losses

We count only outcomes beyond a 5% threshold as experiences; the asymmetry does not hinge on that choice. Figure 3 plots the *Gain* × *Loss Experience* and *Gain* × *Win Experience* coefficients as the classification threshold varies from 5% to 30%. Both interactions keep their signs across the full range, the win side negative and the loss side positive. Their per-event magnitudes grow as the threshold rises, though part of that growth is mechanical—a higher cutoff counts rarer events, so each counted experience carries more weight—and we therefore read the figure as evidence of sign stability rather than as a salience gradient. Table 4 reports the full specification at the extreme 30% cutoff.

We label the two patterns for reference. The *fear effect* is the loss-side amplification:

Figure 3: Experience–Gain Interactions Across Return Thresholds

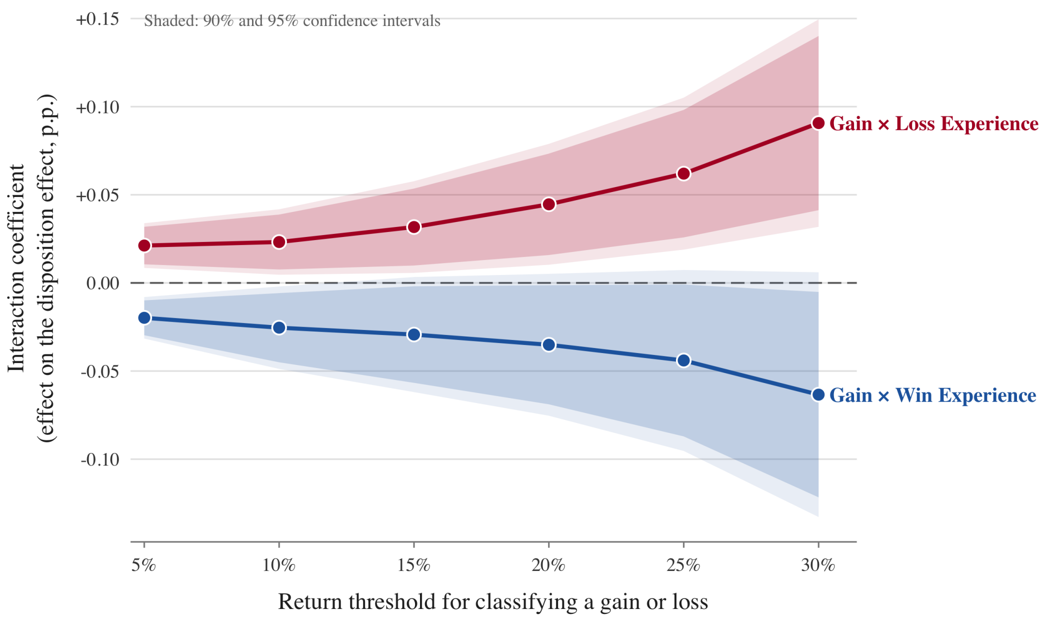

Notes. The figure plots the two interaction coefficients—*Gain* × *Loss Experience* (red, fear) and *Gain* × *Win Experience* (blue, greed)—estimated jointly from the full specification in Column (5) of Table 3. The level effects are not plotted. The horizontal axis is the threshold $|X|$ above which a realized gain or loss counts as an experience, varied from 5% to 30%; a higher threshold counts only more extreme realizations. Shaded regions are 90% (darker) and 95% (lighter) confidence intervals, computed from standard errors clustered by investor, time, and stock.

investors who have realized large losses sell winners relative to losers more readily. The *greed effect* is the win-side attenuation: investors who have realized large gains hold their winners longer. The labels are descriptive, not explanations. Section 4 derives both patterns from cued recall, and Section 5 rules out the extrapolative optimism and overconfidence that the greed label might invite.

3.4 Trading Experience as a Ratio

We next measure experience as a share rather than a count. Counts conflate valence with volume: more active investors accumulate more winning and more losing experiences alike, so each count partly tracks overall activity rather than the balance of past outcomes. The share of prior trades that ended in a large loss (gain) holds activity fixed and asks only about the mix (Equation 1).

Table 4: **Experience-Driven Disposition Effect (30% Threshold)**

Dependent Variable:	**Sale × 100**	
$\|X\| > 30\%$	**(1)**	**(2)**
Gain	2.5063***	2.4275***
	(4.74)	(4.42)
Loss Experience	0.0912*	0.0657
	(1.90)	(1.19)
Win Experience	-0.0035	0.0147
	(-0.04)	(0.14)
Gain × Loss Experience		0.0907***
		(3.02)
Gain × Win Experience		-0.0634*
		(-1.79)
Constant	8.7203***	8.7541***
	(44.35)	(43.88)
Observations	43,649,867	43,649,867
R-squared	0.139	0.139

Notes. Linear probability model where the dependent variable equals 1 if the position is sold (multiplied by 100). *Loss Experience* and *Win Experience* count prior trades with realized returns below −30% or above +30%, respectively, relative to the investor's average purchase price. Coefficients are reported with t-statistics in parentheses. All specifications include investor, stock, and time fixed effects, and standard errors are clustered by investor, time, and stock. The sample period is July 2013 to February 2016. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 5 re-estimates the same linear probability model, replacing the experience counts with the losing- and winning-experience shares (Equation 1). The winning-versus-losing contrast sharpens. In Column (2), the *Gain* × *Loss Experience Share* interaction enters at 32.37 ($t = 11.63$), significant at the 1% level: moving from a loss-free history to an all-loss one—the extremes of the share—widens the gain–loss gap in selling by roughly 32 percentage points, against a baseline *Gain* coefficient of 1.80. The winning share pulls the other way. In Column (4), the *Gain* × *Win Experience Share* interaction is −5.93 ($t = -3.72$), significant at the 1% level, so a loss-heavy past strengthens the disposition effect while a win-heavy past dampens it[10]. Column (5) enters both shares jointly and reproduces the pattern, with the two interactions at 32.49 ($t = 11.66$) and −6.80 ($t = -4.34$). The asymmetry survives the move from counts to shares, so it tracks the relative mix of winning and losing experiences, not the number of trades behind them.

[10] At low thresholds nearly every realized trade clears the cutoff and so counts as a win or loss, so the two shares sum to approximately one, $s^{\text{win}} + s^{\text{loss}} \approx 1$ (exactly one when no sub-threshold trades remain). The shares—and their *Gain* interactions—are then mutually near-collinear (and each near-collinear with the constant), leaving the joint specification in Column (5) ill-conditioned. We therefore measure experience shares at the 30% threshold, where a substantial fraction of trades fall below the cutoff: this restores independent variation in the two shares and lets both enter jointly without inflating their standard errors. The contrast is not specific to this threshold—the count-based specification of Table 3, in which the winning- and losing-experience counts are not collinear, identifies the two channels separately at the 5% threshold.

Table 5: The Experience-Driven Disposition Effect: Share of Losing vs. Winning Experiences

Dependent Variable:	**Sale × 100**				
$\|X\| > 30\%$	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**
Gain	2.4502***	1.8020***	2.4504***	2.5127***	1.8712***
	(4.64)	(3.36)	(4.64)	(4.71)	(3.45)
Loss Experience Share	3.9275***	-3.0334***			-2.9903***
	(5.93)	(-2.89)			(-2.85)
Gain × Loss Experience Share		32.3716***			32.4871***
		(11.63)			(11.66)
Win Experience Share			2.7827***	6.0898***	6.8336***
			(3.45)	(4.67)	(5.29)
Gain × Win Experience Share				-5.9286***	-6.7951***
				(-3.72)	(-4.34)
Constant	8.4407***	8.6346***	8.4973***	8.4649***	8.5648***
	(44.39)	(45.11)	(45.73)	(44.90)	(44.05)
Observations	42,969,625	42,969,625	42,969,625	42,969,625	42,969,625
R-squared	0.127	0.127	0.127	0.127	0.127

Notes. This table reports estimates from the linear probability model described in equation 3, using an alternative measure of trading experience based on the share of large losing or winning trades. The dependent variable is a dummy indicating whether a stock is sold, multiplied by 100. Columns (1)–(2) use *Loss Experience Share* – the proportion of prior trades that ended in a large loss; Columns (3)–(4) use *Win Experience Share*; Column (5) includes both shares jointly. The interaction terms *Gain × Loss Experience Share* and *Gain × Win Experience Share* capture how the disposition effect varies with the composition of past experiences. Experience shares are measured at the 30% threshold: at lower thresholds nearly all trades clear the cutoff, so $s^{\text{win}} + s^{\text{loss}} \approx 1$ and the two shares cannot be separately identified in Column (5), whereas the count-based results of Table 3 establish the same winning-versus-losing contrast at the 5% threshold. All specifications include investor, stock, and time fixed effects. Standard errors are clustered by investor, time, and stock. The sample period is July 2013 to February 2016. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

4 A Memory-Based Model

In this section, we develop a model of cued recall in the context of stock market decision-making, drawing on the frameworks of Howard and Kahana (2002) and Kahana, Paron and Wachter (2025). We then extend the model to align with and explain the key empirical patterns documented in our analysis.

The standard account of the disposition effect relies on prospect theory. The literature has pointed out numerous problems with this account, including the requirement of specifying a reference point,[11] and explaining why investors would have a utility function over returns, and in particular, returns that are converted into cash (sometimes called "realization utility").

[11] Prospect theory has its origin in explaining subject's responses to lottery-type questions in the laboratory, where it is necessary methodologically to define a reference point (Luce, 2000), utility over gains and losses. At some point, this methodological point became a theory about the world, though on what basis was not made clear.

The more recent evidence on the diminishment of the disposition effect when accounting for experience increases these problems. While it is plausible that investors may exhibit more rational behavior after learning from their mistakes, the premise of prospect theory is that the behavior is rational because it originates in the utility function. Learning about having a different utility function is at present outside of economics.

Recent models have gotten around this difficulty by asserting that utility parameters change, and through experience (e.g., Gennaioli and Shleifer (2010), Barberis (2012), Barberis and Xiong (2012)). This account does not explain why the parameters would change in a particular way (nor, of course, does it explain why "realization utility" over losses is how investors think about their portfolios).

The evidence in the present paper strains this approach even further, for it fails to explain why winning experiences versus losing experiences would affect the utility function differently.

Here we present a single mechanism, based on representativeness combined with skewness in underlying stock returns, that accounts for the baseline disposition effect, its diminution under the experience of gains, and its increase under the experience of large losses. Our focus is, however, on explaining the new evidence in this paper, acknowledging that there might be other explanations for the disposition effect itself, particularly when skewness is taken into account.

The decision of which stock to sell resembles a cued recall task, with "stocks I own" forming the cue. We use the version of the Howard and Kahana (2002) temporal context model in Kahana, Paron and Wachter (2025), applied to cued recall rather than associative recognition.

Agents view features f_i: basis vectors in a high-dimensional vector space. These features become embedded in memory together with a temporal context, an internal mental state that is persistent and evolves through time. Context, through the memory matrix, governs what comes to mind. Specifically, given an initial condition M_0, memory is comprised of outer products of features and contexts $\{f_i, x_i\}_{i=1}^n$:

$$M_n \equiv M_{n-1} + x_n f_n^\top \tag{4}$$
$$= M_0 + \sum_{i=1}^{n} x_i f_i^\top \tag{5}$$

In what follows, we will be interested in a process known as cued recall, which has long been studied in the literature on human memory. Context is retrieved in response to a features cue:

$$\boldsymbol{x}_{\text{CUE}} \propto M_n f_{\text{CUE}}, \tag{6}$$

where context is scaled to have a norm of 1. Substituting (6) into (5) yields:

$$\boldsymbol{x}_{\text{CUE}} \propto M_0 f_{\text{CUE}} + \sum_{i=1}^{n} x_i (f_i^\top f_{\text{CUE}}). \tag{7}$$

Equation 7 shows that features evoke the past contexts under which they are experienced. A context x appears in the sum in (7) if the corresponding f_i equals the current feature; otherwise it does not. Thus context is a weighted average of past contexts under which that feature was experienced. In what follows, we assume M_0 is the zero matrix and we suppress the n subscript on M. Context x_i is endogenous and persistent, and will be described in more detail in what follows.

We use well-established principles of association to determine what pops to mind. As mentioned above, we think of the decision as a cued recall task. In a cued recall task, the first step is to retrieve the context associated with the physical features that form the cue as in (6). Unlike in a recognition task, where it is sufficient to judge the similarity of the cued context with a target context,[12] here features are retrieved based on the cue as in Howard and Kahana (2002):

$$\boldsymbol{f}_{\text{CUE}} \propto M^\top \boldsymbol{x}_{\text{CUE}}. \tag{8}$$

These are features that are experienced under contexts similar to $\boldsymbol{x}_{\text{CUE}}$. There will certainly be nonzero weight on the element corresponding to f_{CUE}, namely $f_{\text{CUE}} \cdot \boldsymbol{f}_{\text{CUE}} \neq 0$. However,

[12] See Kahana, Paron and Wachter (2025) for a discussion of recall versus recognition

any feature experienced at the same time as, or close in time to, f_{CUE} will also be recalled. This is the core mechanism of the temporal context model of Howard and Kahana (2002).

Each candidate feature f_i likewise retrieves a context

$$\boldsymbol{x}_i \propto M f_i. \tag{9}$$

Like $\boldsymbol{x}_{\text{CUE}}$, each retrieved context is scaled to have a norm of one; throughout, the norm is the L^2 norm. This per-candidate normalization is the source of the $1/\sqrt{n_i}$ factors in Theorem 2 below, where n_i denotes the number of observations of f_i.

The vector of retrieved features, $\boldsymbol{f}_{\text{CUE}}$, has nonzero weight on a number of elements. For one of them to rise above the others, it first needs to have a nonzero weight in $\boldsymbol{f}_{\text{CUE}}$, and relatively speaking, a similar context to $\boldsymbol{x}_{\text{CUE}}$. Because recall is competitive, if the context is sufficiently different from $\boldsymbol{x}_{\text{CUE}}$, recall of that feature will be suppressed. The following model captures these properties. The competition runs over feature types rather than individual episodes: multiple episodes may share the same feature, and each of the N distinct feature types stored in memory enters the competition once. For $i = 1, \ldots, N$,

$$\mathbb{P}(\text{recall of } f_i) \propto (\boldsymbol{x}_i \cdot \boldsymbol{x}_{\text{CUE}}) \|f_i \cdot \boldsymbol{f}_{\text{CUE}}\|_0.$$

Equivalently,

$$\mathbb{P}(\text{recall of } f_i) = \frac{(\boldsymbol{x}_i \cdot \boldsymbol{x}_{\text{CUE}}) \|f_i \cdot \boldsymbol{f}_{\text{CUE}}\|_0}{\sum_{j=1}^{N} (\boldsymbol{x}_j \cdot \boldsymbol{x}_{\text{CUE}}) \|f_j \cdot \boldsymbol{f}_{\text{CUE}}\|_0}. \tag{10}$$

Note that $\|f_i \cdot \boldsymbol{f}_{\text{CUE}}\|_0$ is the zero norm; it is either equal to zero or one, zero if f_i is not recalled in (8). This specification is a simplified version of the retrieval process in Lohnas, Polyn and Kahana (2015).

We use the model in Kahana, Paron and Wachter (2025) for context evolution:

$$x_{i+1} \begin{cases} = & x_i \text{ if no distractor} \\ \in & \{x_1, \ldots, x_i\}^{\perp} \text{ otherwise} \end{cases} \tag{11}$$

See Kahana, Paron and Wachter for further details. This formula for context allows for basis vector features that are observed instantaneously to be coded together in context, for example, a stock that is owned is a feature, as is whether the stock has a gain. This allows for the creation of a composite feature, the sum of two feature vectors observed at (nearly) the same point in time. However, (11) is a highly simplified version of what is believed to be true contextual evolution; therefore, we abstract from well-established principles of recall, such as recency and contiguity.

We now link this memory model to the specific context of the decision to sell gains or losses. The agent's information set contains owned stocks alongside other commodities. Among these other commodities are stocks the agent has previously sold, which we track separately because selling moves a stock from the owned to the sold category, and it is this reallocation that generates the experience effects below. We therefore distinguish three categories: owned stocks, sold stocks, and other, where sold stocks are a sub-category of the non-owned ("other") commodities. Before the agent has any selling experience there are no sold stocks, so $n_{\text{SOLD}} = 0$ and the baseline analysis involves only owned stocks and other commodities. Each commodity could have either a large or a small return. Define

$$n_k \equiv \text{number of observations of } k$$

and

$$n_{k,\ell} \equiv \text{number of joint observations of } k \text{ and } \ell$$

for $k, \ell \in \{\textsc{Owned}, \textsc{Sold}, \textsc{Other}, \textsc{Large}, \textsc{Small}\}$. The cue is "stocks I own." We use the shorthand notation $\mathbb{P}(\textsc{Large})$ for $\mathbb{P}($ recall of $f_{\text{LARGE}})$, and similarly for $\mathbb{P}(\textsc{Small})$.

Assumption 1 *Agents perceive returns as belonging to categories: small and large. Whereas commodities of any type can have small returns, only stocks can have large returns.*

Theorem 1 *Assume that the agent has no winning or losing experiences. The agent competitively recalls either large or small returns with*

$$\begin{aligned}
\mathbb{P}(\textsc{Large}) &\propto (\boldsymbol{x}_{\text{LARGE}} \cdot \boldsymbol{x}_{\text{OWNED}})\|f_{\text{LARGE}} \cdot \boldsymbol{f}_{\text{OWNED}}\|_0 \\
\mathbb{P}(\textsc{Small}) &\propto (\boldsymbol{x}_{\text{SMALL}} \cdot \boldsymbol{x}_{\text{OWNED}})\|f_{\text{SMALL}} \cdot \boldsymbol{f}_{\text{OWNED}}\|_0
\end{aligned}$$

with $\mathbb{P}(\textsc{Large}) + \mathbb{P}(\textsc{Small}) = 1$.

Note that we assume that either LARGE or SMALL is recalled. Of course, the agent will also recall OWNED. As is standard in the memory literature on recall, we assume that the agent suppresses the recall of merely repetitive features.

We apply results in Kahana, Paron and Wachter (2025) to connect the right-hand side quantities to the numbers of observations:

Theorem 2

$$\mathbb{P}(\textsc{Large}) \propto \frac{n_{\textsc{Owned},\textsc{Large}}}{\sqrt{n_{\textsc{Owned}} n_{\textsc{Large}}}} \tag{12}$$

$$\mathbb{P}(\textsc{Small}) \propto \frac{n_{\textsc{Owned},\textsc{Small}}}{\sqrt{n_{\textsc{Owned}} n_{\textsc{Small}}}} \tag{13}$$

Proof. See Kahana, Paron and Wachter (2025). ■

Corollary 1 *If the number of observations of assets other than stocks is sufficiently large,* $\mathbb{P}(\textsc{Large}) > \mathbb{P}(\textsc{Small})$.

Proof. It follows from Theorem 2 that $\mathbb{P}(\textsc{Large}) > \mathbb{P}(\textsc{Small})$ if and only if

$$\frac{n_{\textsc{Owned},\textsc{Large}}}{n_{\textsc{Owned},\textsc{Small}}} > \sqrt{\frac{n_{\textsc{Large}}}{n_{\textsc{Small}}}} \tag{14}$$

By Assumption 1, $n_{\textsc{Small}} = n_{\textsc{Owned},\textsc{Small}} + n_{\textsc{Other},\textsc{Small}} = n_{\textsc{Owned},\textsc{Small}} + n_{\textsc{Other}}$, whereas $n_{\textsc{Large}} = n_{\textsc{Owned},\textsc{Large}}$. Increasing $n_{\textsc{Other}}$ increases $n_{\textsc{Small}}$, and hence decreases the right-hand side of (14), while making no other change. The inequality therefore holds for $n_{\textsc{Other}}$ sufficiently large. ■

We now assume that the number of commodities with small price changes is far larger than the number of stocks. This, combined with the fact that stocks exhibit positive skewness (and the memory dynamics) delivers the baseline disposition effect.

Assumption 2 *The agent sells the stock whose returns pop to mind.*

Assumption 3 *Stock returns are positively skewed.*

In our sample, the skewness of paper returns—unrealized gains and losses at the account level—is 23.1, which confirms a highly right-skewed distribution. Consistent with this, Oh and Wachter (2022) provide broader evidence that stock returns generally exhibit positive skewness.

Assumption 4 *The agent creates memories according to (4) and (11).*

Assumption 5 *Cued recall works as follows. The cue retrieves context $\boldsymbol{x}_{\text{CUE}} \propto M f_{\text{CUE}}$. This context in turn retrieves the features $\boldsymbol{f}_{\text{CUE}} \propto M^{\top} \boldsymbol{x}_{\text{CUE}}$. Which of these features gets recalled is competitive with probability determined by nonzero weight in $\boldsymbol{f}_{\text{CUE}}$ and similarity of context associated with the features and the cue as in (10)*

Theorem 3 *The agent is more likely to sell gains than losses.*

Proof. By Theorem 2, the agent is more likely to recall large rather than small returns. It follows from Assumption 3 that these large returns are more likely to be gains. It follows from Assumption 5 that the agent recalls these gains and, finally, from Assumption 2 that the stock with the gain is the stock the agent chooses to sell. ■

The key mechanism is that small returns have contexts associated with "other." They are not as good a match, in terms of context, to "stock," and hence less likely to be recalled as part of the cue.

Given this basic structure, the experience effects follow easily. Let $\mathbb{P}_k(\text{LARGE})$ denote the probability of recall of large returns, given k winning experiences, and similarly for $\mathbb{P}_k(\text{SMALL})$.

Theorem 4 *Winning experiences decrease the tendency to sell gains.*

Proof. It follows from Assumption 3 that the large returns are predominantly gains. It suffices to show that winning experiences reduce the ratio $\mathbb{P}_k(\text{LARGE})/\mathbb{P}_k(\text{SMALL})$. After k winning experiences $n_{\text{OWNED,LARGE}}$ is replaced by $n_{\text{OWNED,LARGE}} - k$ whereas $n_{\text{SOLD,LARGE}}$ is replaced with $n_{\text{SOLD,LARGE}} + k$. As a consequence, the total number of owned stocks n_{OWNED} also falls by k, whereas the total number of sold stocks rises by k. It suffices to show that $\mathbb{P}_k(\text{LARGE})/\mathbb{P}_k(\text{SMALL})$ is decreasing in k:

$$\frac{\mathbb{P}_k(\text{LARGE})}{\mathbb{P}_k(\text{SMALL})} = \frac{n_{\text{OWNED,LARGE}} - k}{n_{\text{OWNED,SMALL}}} \sqrt{\frac{n_{\text{SMALL}}}{n_{\text{LARGE}}}}$$

The right-hand side is clearly decreasing in k. ■

Likewise, if the agent sells losses, because these losses tend to be small, this will cement the association between owned stocks and large gains. A losing experience means $n_{\text{Owned,Small}}$ is replaced by $n_{\text{Owned,Small}} - k$. The proof is analogous to Theorem 4.

Theorem 5 *Losing experiences increase the tendency to sell gains.*

The results in the previous sections also indicate that large losses increase the disposition effect more than small losses. While large losses are less likely than large gains, should they occur, they will have the stated effect. We introduce additional feature vectors $f_{\text{Large loss}}, f_{\text{Large gain}}$, instead of f_{Large}. We modify the notation accordingly.

Like large gains, large losses only occur with stocks. We now need to consider three-way comparisons in the theorems above, but given that large gains are more common than large losses, the baseline disposition effect (Theorem 3), the effect of winning experiences (Theorem 4) and the effect of small losing experiences (Theorem 5) still work under this extended model.

As with winning experiences, k large losing experiences implies $n_{\text{Owned,Large loss}}$ is replaced by $n_{\text{Owned,Large loss}} - k$, whereas $n_{\text{Sold,Large loss}}$ is replaced by $n_{\text{Sold,Large loss}} + k$. This will decrease the probability $\mathbb{P}(\text{Large loss})$ by the same reasoning as Theorem 4.

More broadly, because large returns of either sign are the most retrievable—and hence the most likely to be sold—the model implies a V-shaped relation between a position's realized return and the propensity to sell it, the regularity documented by Ben-David and Hirshleifer (2012).

The model's results correspond to the three empirical patterns of Section 3. Theorem 3 delivers the baseline disposition effect with no winning or losing experiences: stock returns are positively skewed, so the large returns in memory are predominantly gains, and the investor recalls, and sells, a gain. Theorem 4 shows that winning experiences decrease the tendency to sell gains; this is the attenuation in *Gain* × *Win Experience*. Theorem 5 shows that losing experiences increase the tendency to sell gains; this is the amplification in *Gain* × *Loss Experience*.

5 Ruling Out Other Mechanisms

Having documented the experience-driven asymmetry and shown that a single cued-recall channel reproduces it, we now rule out the leading alternative mechanisms. The most basic is that the interactions are a mechanical by-product of how experience is built from past trades. Beyond it, the pattern might reflect unobserved trading skill or a portfolio-level rather than a position-level disposition effect (An et al., 2024), an artifact of recent returns, portfolio rebalancing, regression to the mean, or the overtrading that overconfidence would bring. We take these in turn; none accounts for the asymmetry.

5.1 Is the Experience Effect Mechanical?

Investors who have been trading for longer tend to accumulate more positions. If trading activity does not scale with positions, then the probability of selling any one position will fall with the number of positions, and thus also with trading experience. Because disposition is measured by taking the difference between the propensity to sell winners and the propensity to sell losers, the overall size of the effect might diminish with experience due to this effect. Furthermore, if investors who have been trading for longer tend to realize more winners, then the disposition effect would fall, mechanically, with winning experiences. To rule out this confound, we turn to simulations.[13]

We conduct a Monte Carlo simulation in which investors sell positions independently of their past experience and rerun our statistical analyses on the simulated data. We calibrate portfolio dynamics to observed moments in the data. First, as the reasoning in the previous paragraph suggests, within-investor experience and portfolio size are strongly positively correlated.[14] It is roughly the same across the two types of experiences: $+0.61$ for winning experience, $+0.58$ for losing. We also calibrate to the decline in the held winner share as experience accumulates. Not surprisingly, the more winning experiences, the lower the winner share (the correlation is -0.26). The held winner share also declines as losing experience

[13]A more obvious confound may appear to be that as the investor sells winners, there are fewer winners in the portfolio. Under random selling, this reduces the unconditional probability of a winner being sold. It does not, however, reduce the probability of sale conditional on being a winner, which is how the disposition effect is measured.

[14]All our measures are within-investor because our analyses employ investor fixed effects.

accumulates, but less strongly (the correlation is -0.16). Appendix A3 details the data-generating process and plots the simulated distributions. Because selling in the simulation does not depend on experience, any nonzero interaction coefficient in the simulated data is mechanical by construction.

Table 6 reports results under three nulls: random selling, an experience-independent disposition effect, and an experience-independent disposition effect combined with the tendency to sell extreme positions documented by Ben-David and Hirshleifer (2012) ("V-shape selling behavior"). Under random selling, one is not more likely to sell a gain than a loss. Experience correlates with a greater number of positions, but in this simulation there is no baseline disposition effect, and so the propensity to sell a gain over a loss is zero and thus does not fall with experience. This explains the near-zero coefficients on the interaction between gain and both types of experiences.[15] Random selling, combined with experience, cannot mechanically produce the effects we observe.

Our second null is a fixed disposition effect. Indeed, the simulation shows that experience mechanically diminishes this fixed disposition effect: losing and winning experiences lead to a lower propensity to sell winners for the reason we describe at the start of this section. This is potentially an issue for interpreting studies focusing on all experiences, rather than the differential effect of winning and losing experiences. However, it is important to note that under a realistic calibration, the size of this effect is too small to explain the observed diminishment of the disposition effect on the basis of experience. Because winning experiences and losing experiences correlate with large portfolios about equally, the confound that we are concerned about, namely that winning experience indicates more experience overall, and hence a larger portfolio and a smaller disposition effect, cannot explain our results. The reason is that nearly the same dynamics work for losing experiences, resulting in a negative coefficient on the interaction, the opposite sign of what we find. A fixed disposition effect, combined with experience, also cannot mechanically produce the effects we observe.

A close look at the coefficients for this second null reveals an interesting pattern: the

[15] The win-side coefficient under random selling, -0.0022, is small economically, but too large to be simulation noise. Under equal sale probabilities the true winner–loser gap is zero by construction. As mentioned above, the propensity to sell any one position is decreasing in experience, and does so in a non-linear way. The regression, however, estimates a linear effect. The interaction terms between gain and winning and losing experiences pick up the residual from the misspecification.

mechanical effect of experience, while small for both winning and losing experience, is smaller in magnitude (that is, less negative) for winning experience than for losing experience. This is the opposite of what our model predicts and what the data show, but it is worth considering why it occurs. As emphasized above, the main confound arises from portfolio size. A larger portfolio mechanically leads to a reduced disposition effect. A secondary mechanical effect offsets this main confound: as an investor sells winners, there will be fewer as time goes by, raising the probability of a sale under the fixed disposition effect. Winning experiences should thus increase the propensity to sell winners. This portfolio composition effect offsets the portfolio growth effect in our second null, causing the coefficient on the interaction between gain and winning experience to be closer to zero than that between gain and losing experience.

Our third null is a fixed disposition effect, combined with a tendency to sell the largest winners and losers that is not modulated by experience. One might think that a tendency to remove the most extreme positions would work roughly equally on the gain and the loss side. Each period, the extreme positions are removed, but there is an overall tendency to sell winners. As the number of positions grows, the disposition effect diminishes, though less for winners due to the portfolio composition effect described just above. Table 6 shows, however, that the portfolio composition effect is stronger than in the second null, so much so that the sign on the interaction between gain and winning experience is positive. There are two reasons why the portfolio composition effect is stronger. The first is that, due to skewness, extreme returns tend to be winners. The second is that winners tend to be removed earlier in the simulation for a given investor. The effect of removing winners matters more when there are fewer of them. The net effect is the opposite of what we observe in the data. On the loss side, the coefficient is essentially unchanged from the second null.

5.2 Skill and the Portfolio-Level Disposition Effect

An et al. (2024) document that the disposition effect is sharply lower when the investor's portfolio is at a gain. We see a similar pattern in our data (Table 7, Base column). One explanation is skill. Skilled investors are more likely to hold portfolios at a gain, and skill may also reduce the disposition effect. Such an explanation suggests a possible confound.

Table 6: Monte Carlo Assessment of Mechanical Bias ($|X| > 5\%$)

	Gain × LossExp	Gain × WinExp
Paper estimates (Table 3, Column 5)		
Coefficient	+0.0212	−0.0198
t-statistic	3.28	−3.30
Scenario A: Random selling		
MC mean	+0.0006	−0.0022
95% MC CI	[−0.020, +0.020]	[−0.024, +0.018]
p_{MC}	0.025	0.040
Scenario B: Fixed DE		
MC mean	−0.0064	−0.0032
95% MC CI	[−0.025, +0.014]	[−0.023, +0.015]
p_{MC}	0.005	0.045
Scenario C: Fixed DE + V-shaped selling		
MC mean	−0.0068	+0.0019
95% MC CI	[−0.025, +0.012]	[−0.013, +0.018]
p_{MC}	0.005	0.005

Notes. This table reports the Gain × Experience interaction coefficients from 200 Monte Carlo replications of the specification in Table 3, Column (5), at the $|X| > 5\%$ threshold; Figure 4 in Appendix A3 plots the corresponding distributions, where the 30%-threshold pattern is also discussed. Under all three scenarios, selling probabilities are independent of accumulated experience; any nonzero interaction coefficient is therefore a mechanical artifact of portfolio-composition changes. "MC mean" is the average coefficient across the 200 simulations; "95% MC CI" is the 2.5th–97.5th percentile interval across simulations; p_{MC} is the fraction of simulations producing a coefficient at least as extreme as the paper's estimate in the direction of the paper's finding. The simulation is calibrated to both within-investor co-movements that govern this channel: the correlation between experience and portfolio size (+0.61 winning, +0.58 losing) and the correlation between experience and the held winner share (−0.26 and −0.16).

More realized gains would indicate more skill, and hence a smaller disposition effect, while more realized losses would indicate less skill, and hence a larger one. Because our experience measures count realized gains and losses, skill could in principle produce both interactions we document. Investor fixed effects do not rule this out, because skill, like experience, can accumulate.

We therefore condition on the state of the portfolio. Following An et al. (2024), we add an indicator for whether the portfolio is at a gain to the position-level specification, interacted with the position-level gain and with experience.[16] The regression specification is

$$\begin{aligned}
\text{Sale}_{i,j,t} = \alpha &+ \beta_1 \times \text{Gain}_{i,j,t} + \beta_2 \times \text{Portfolio Gain}_{i,t} + \beta_3 \times (\text{Portfolio Gain}_{i,t} \times \text{Gain}_{i,j,t}) \\
&+ \beta_4 \times \text{Experience}_{i,t} + \beta_5 \times (\text{Gain}_{i,j,t} \times \text{Experience}_{i,t}) \\
&+ \beta_6 \times (\text{Portfolio Gain}_{i,t} \times \text{Experience}_{i,t}) \\
&+ \beta_7 \times (\text{Portfolio Gain}_{i,t} \times \text{Gain}_{i,j,t} \times \text{Experience}_{i,t}) + \epsilon_{i,j,t} \qquad (15)
\end{aligned}$$

where Portfolio Gain $= 1$ if the total unrealized return on the account is positive at the time of sale. Table 7 reports the estimates. With the *Portfolio Gain* $\times$ *Gain* interaction in the regression, the *Gain* coefficient (≈ 5.2) measures the disposition effect when the portfolio is at a loss. When the portfolio is at a gain, the disposition effect is the sum of β_1 and β_3, which is negative in our sample. We thus reproduce the effect of An et al. (2024).[17]

We now turn to our main effect. The **Loss Exp.** and **Win Exp.** columns split experience into losing and winning trades. For portfolios at a gain, an additional loss experience adds $\beta_5 + \beta_7 = 0.2466$ percentage points to the disposition effect, whereas an additional win experience adds only 0.0393 (Panel B). We do not report standard errors for these sums; each underlying interaction is individually significant at the 1% level. Having a portfolio at a gain could be correlated with experience, but if so, even experienced investors are influenced by whether these experiences are losses versus gains.

For losing portfolios, we do not find our effect in the baseline numbers: an additional experience of either type compresses the disposition effect by a similar amount ($\beta_5 = -0.0591$ for losing experience, -0.0663 for winning). Note, however, that the previous subsection shows that our effect has to push against the portfolio growth and portfolio composition effects. The portfolio composition effect is especially strong when there are few winners, as is the case when the portfolio is at a loss. We therefore caution against interpreting these results. Regardless of the loss-side interpretation, the survival of our effect on the gain side indicates that time-varying investor skill is not the driver for our results.

[16] We report results for the 30% threshold. Results are robust to the threshold definition.

[17] An et al. (2024) interpret the pattern as mental accounting. Portfolio composition offers a further candidate: a portfolio at a gain typically holds many positions at small gains, so a tendency to sell winners is spread over many positions, and the probability that any one of them is sold is small. We do not take a stance on the source of this effect, which is beyond the scope of our paper.

Table 7: Experience-Driven and Portfolio-Driven Disposition Effect

Panel A: Experience-Driven and Portfolio-Based Disposition Effect

Dependent Variable:	**Sale × 100**			
Experience variable in cols. (2)–(4):	**Base**	**Total Exp.**	**Loss Exp.**	**Win Exp.**
Gain	5.1834***	5.7315***	5.3401***	5.2954***
	(10.01)	(9.83)	(9.89)	(9.96)
Portfolio Gain	2.7031***	3.1320***	2.8601***	2.8106***
	(10.38)	(10.54)	(10.51)	(10.67)
Portfolio Gain × Gain	-7.9336***	-8.6287***	-8.3488***	-8.1018***
	(-13.26)	(-13.52)	(-13.48)	(-13.30)
Experience		0.0035**	0.1001***	0.0712
		(2.42)	(4.36)	(0.82)
Gain × Experience		-0.0041***	-0.0591***	-0.0663***
		(-3.42)	(-3.50)	(-3.08)
Portfolio Gain × Experience		-0.0036***	-0.1215***	-0.0773***
		(-3.48)	(-3.08)	(-6.03)
Portfolio Gain × Experience × Gain		0.0053***	0.3057***	0.1056***
		(4.33)	(3.15)	(3.87)
Constant	8.5824***	8.1990***	8.4401***	8.5357***
	(46.99)	(32.09)	(43.79)	(38.79)
Observations	43,649,867	43,649,867	43,649,867	43,649,867
R-squared	0.141	0.142	0.142	0.142

Panel B: Experience and the Disposition Effect in the Gain State

	Loss Experience	**Win Experience**
Effect on the winner–loser gap ($\beta_5 + \beta_7$)	0.2466	0.0393

Notes. This table presents estimates from the linear probability model described in equation 15, where the dependent variable is a dummy equal to 1 if a stock is sold, multiplied by 100 for ease of interpretation. Panel A examines how individual trading experience and portfolio-level performance jointly shape the disposition effect. The four columns differ in which experience variable enters the regression as "Experience": Column (1) includes baseline controls only; Column (2) uses *Total Experience* – the cumulative number of prior trades; Column (3) uses *Loss Experience* – the cumulative count of prior trades with a realized return below −30%; Column (4) uses *Win Experience* – the cumulative count of prior trades with a realized return above +30%. Row labels "Experience", "Gain × Experience", etc. therefore refer to whichever variable is indicated in the column header.
Gain is an indicator for whether the stock has an unrealized capital gain at the time of decision. *Portfolio Gain* is a dummy indicating whether the investor's portfolio is, on net, in a gain position. The interaction terms *Portfolio Gain* × *Gain* and its higher-order interactions capture how broader portfolio context moderates the disposition effect.
Panel B reports the effect of one additional experience on the winner–loser gap in sale propensity when the portfolio is at a gain, computed as $\beta_5 + \beta_7$ from the corresponding column of Panel A; statistical significance is not reported for these combined effects.

All regressions include investor fixed effects, stock fixed effects, and time fixed effects. Standard errors are clustered by investor, time, and stock. The sample period is July 2013 to February 2016. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

5.3 Persistence Across the Trading History

Our experience counts could be proxying for a mean-reverting process rather than a genuine response to outcomes. Suppose individual disposition effects revert toward a common population level. An investor who realized many winners then had a high past DE that regresses down, while one who realized many losers had a low past DE that regresses up. The experience counts would track these past realizations, and the asymmetric gradients we estimate would emerge mechanically, with no genuine response to accumulated outcomes.

This account predicts that the experience–DE relationship weakens among investors with long trading histories, where the past DE has had more time to revert. We re-estimate the baseline specification on the subsample of investors with account tenure above the sample median. Table 8 reports the results against the full-sample headline from Table 3, Column (5).

The data move the opposite way. Both interactions strengthen in the long-history subsample: *Gain* $\times$ *Loss Experience* rises from +0.0212 ($t = 3.28$) to +0.0302 ($t = 4.37$), and *Gain* $\times$ *Win Experience* moves from −0.0198 ($t = -3.30$) to −0.0219 ($t = -5.31$), both significant at the 1% level. Mean reversion predicts attenuation where reversion has had the most time to operate; the response to accumulated outcomes instead proves durable. Account tenure in this sample varies over months, not decades (interquartile range 266–362 days, Table 1), so this split gauges the stability of the effect within short histories; it does not measure the size of an investor's lifetime stock of competing memories, the margin the age test in Section 6 addresses.

5.4 Extrapolation and Recent Returns

Investors extrapolate recent returns, and recent returns correlate with realized outcomes. An investor who has booked many large gains tends to hold stocks that have risen, and extrapolative beliefs about those risers, rather than the memory of past realized outcomes, could be what shifts selling. Were that the case, the win-side attenuation would be return extrapolation under another name.

To separate the two channels, we re-estimate the full 5% specification adding the held

Table 8: **Persistence in the Long-History Subsample (5% Threshold)**

Dependent Variable:	**Sale × 100**	
	(1) Full sample	**(2) Long-history sub-sample**
Gain	2.4781***	1.5401***
	(4.35)	(2.64)
Loss Experience	0.0109	−0.0101**
	(1.13)	(-2.33)
Win Experience	0.0034	0.0291***
	(0.36)	(5.00)
Gain × Loss Experience	+0.0212***	**+0.0302*** **
	(3.28)	(4.37)
Gain × Win Experience	−0.0198***	**−0.0219*** **
	(-3.30)	(-5.31)
Constant	8.4974***	7.7038***
	(36.63)	(35.75)
Observations	43,649,867	28,734,001
R-squared	0.138	0.137

Notes. Column (1) reproduces Column (5) of Table 3 (the full-sample headline at the $|X| > 5\%$ threshold). Column (2) re-estimates the same specification on the subsample of investors with account tenure above the sample median. The dependent variable is an indicator for sale, multiplied by 100. *Loss Experience* and *Win Experience* count prior trades with realized returns below −5% or above +5%, respectively, relative to the investor's average purchase price. *Gain* equals 1 if the position has an unrealized gain at the decision time. Both columns include individual, stock, and time fixed effects. Standard errors are clustered by investor, time, and stock; t-statistics are in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively. If individual disposition effects mean-reverted to a common population level, the experience interactions should attenuate in the long-history subsample (where past realizations have had more time to revert); instead, both interactions strengthen.

stock's own recent return as a control, measured over the prior week and, separately, the prior month. Table 9 shows the control leaves the experience interactions intact. With the prior-week return, *Gain* × *Loss Experience* is +0.0218 ($t = 3.37$) and *Gain* × *Win Experience* is −0.0199 ($t = -3.35$); with the prior-month return they are +0.0216 ($t = 3.36$) and −0.0200 ($t = -3.36$), against the headline +0.0212 and −0.0198. The recent-return control enters positively (weekly 5.50, $t = 1.75$; monthly 4.04, $t = 3.04$), so investors sell recent risers more often—return-sensitive selling in the spirit of the V-shaped schedule of Ben-David and Hirshleifer (2012), and difficult to square with extrapolative beliefs, which would favor holding a recent riser. It does not touch the asymmetric experience pattern.

Table 9: The Experience-Driven Disposition Effect, Controlling for Recent Returns (5% Threshold)

Dependent Variable:		**Sale × 100**	
	(1) Headline	**(2) + Past-week**	**(3) + Past-month**
Gain	2.4781***	2.2190***	2.1403***
	(4.35)	(3.80)	(3.70)
Loss Experience	0.0109	0.0112	0.0116
	(1.13)	(1.17)	(1.21)
Win Experience	0.0034	0.0033	0.0030
	(0.36)	(0.34)	(0.32)
Gain × Loss Experience	0.0212***	0.0218***	0.0216***
	(3.28)	(3.37)	(3.36)
Gain × Win Experience	-0.0198***	-0.0199***	-0.0200***
	(-3.30)	(-3.35)	(-3.36)
Past-week return		5.4957*	
		(1.75)	
Past-month return			4.0379***
			(3.04)
Constant	8.4974***	8.3418***	8.3255***
	(36.63)	(36.05)	(36.23)
Observations	43,649,867	42,526,853	42,526,853
R-squared	0.138	0.138	0.138

Notes. Column (1) reproduces the full specification from Table 3, Column (5). Columns (2) and (3) add the held stock's cumulative return over the prior week and the prior month, respectively, as a control. *Loss Experience* and *Win Experience* count prior trades with realized returns below -5% or above $+5\%$, respectively, relative to the average purchase price. *Gain* equals 1 if the position has an unrealized gain at the decision time. All columns include investor, stock, and time fixed effects, and standard errors are clustered by investor, time, and stock. Columns (2)–(3) have fewer observations because positions without a full recent-return history are dropped. The sample period is July 2013 to February 2016. *t*-statistics in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

5.5 Portfolio Rebalancing

Selling winners need not reflect any bias. When a holding appreciates, its weight in the portfolio rises above the investor's target, and a risk-managing investor trims it; a holding that falls becomes underweight and is held. A disposition effect—and the experience interactions we document—could be the footprint of rational rebalancing.

The variable driving the decision separates the two accounts. Rebalancing keys off a position's current weight relative to a target and is invariant to the purchase price. Our *Gain* indicator keys off the cost basis, a backward-looking reference point with no role in mean–variance rebalancing. Two investors holding the same stock at the same current weight, one above and one below their purchase prices, sell at the same rate under any rebalancing rule, yet they realize at different rates in the data. The residual predictive power of the cost basis, conditional on the position characteristics we observe, identifies the bias. Frydman and Wang (2020) sharpen the point: when a Chinese brokerage made the purchase price more prominent on its trading interface, realizations rose by roughly 17%. That shock moved a displayed reference point, not weights or risk, and no rebalancing motive explains it.

Rebalancing also cannot generate our central finding. A rebalancing rule maps current weights into trades and is history-independent: an investor with a given portfolio today trades the same way whether her past trades won or lost. Yet the propensity to sell a winner relative to a loser strengthens with accumulated losing experience ($\textit{Gain} \times \textit{Loss Experience} > 0$) and weakens with winning experience ($\textit{Gain} \times \textit{Win Experience} < 0$), the two moving in opposite directions. The pattern is difficult to reconcile with any history-independent rebalancing rule: the rate at which an investor trims overweight winners would have to depend on the valence of her trading history, yet target weights do not shift with realized outcomes. Investor, stock, and time fixed effects identify this variation within investor and within stock, where any stable rebalancing tendency is differenced out. Even if part of the level of winner-selling is rational rebalancing, the experience asymmetry that is this paper's contribution is orthogonal to it.

Three further features resist a rebalancing reading. Rebalancing trims a position toward its target rather than forcing a full exit, yet our results hold when partial sales are dropped and only complete liquidations are kept (Section 2.3), where the motive does not apply (Odean, 1998). The direction of the bias also runs against tax-motivated selling, which favors realizing losses over gains; throughout our sample, individuals' realized gains on A-shares acquired on the secondary market were exempt from individual income tax, so neither loss-harvesting nor gain-deferral can produce the pattern.[18] And the same investors show a reversed disposition effect in mutual funds (Chang, Solomon and Westerfield, 2016*a*), even though diversification and delegation strengthen the rebalancing rationale for funds; a rational-rebalancing account cannot explain why the sign flips across asset classes held by the same person.

Our baseline omits a position's current portfolio weight as an explicit control, so *Gain* might in principle proxy for an overweight position. The experience interactions are shielded by the history-independence argument above, and controlling for the held stock's recent return—the variable most closely tied to a position becoming overweight—leaves them essentially unchanged (Table 9).

5.6 Trading Intensity and Overconfidence

A run of realized gains could breed overconfidence. Overconfident investors trade more (Barber and Odean, 2000*b*; Statman, Thorley and Vorkink, 2006), and heavier trading mixes in marginal sell decisions that mechanically dilute the measured disposition effect. The win-side attenuation would then reflect overconfidence-driven turnover, not memory, and winning experience should raise trading.

We regress each investor's monthly turnover—value bought and sold during the month, scaled by twice average portfolio value—on the cumulative count of winning experiences, with investor fixed effects and standard errors clustered by investor. The sign goes the wrong way for overconfidence. Winning experience lowers turnover (Table 10: -0.29, $t = -17.8$ with investor fixed effects): investors who have realized more gains trade less, not more. The

[18] The exemption is the temporary one in force since Caishuizi [1998] No. 61 (Ministry of Finance and State Administration of Taxation); only transfers of restricted shares are taxed (Caishui [2009] No. 167), which does not apply to the secondary-market trades in our sample. Sell-side trading costs remained: a 0.1% stamp duty levied on the seller only, and dividend taxes differentiated by holding period under Caishui [2012] No. 85, with dividends on shares held over one year exempt from September 2015 (Caishui [2015] No. 101).

Table 10: **Winning Experience and Monthly Turnover**

Dependent Variable:	**Monthly Turnover**	
	(1)	**(2)**
Win Experience	-0.0825***	-0.2911***
	(-10.14)	(-17.83)
Constant	7.3127***	8.4999***
	(116.51)	(82.97)
Observations	881,576	838,399
R-squared	0.001	0.313
Investor FE	No	Yes

Notes. The dependent variable is investor-level monthly turnover, defined as total purchase plus sale value during the month divided by twice the investor's average portfolio value, $(\text{buy} + \text{sell})/(2 \times \text{average portfolio value})$. *Win Experience* is the cumulative count of an investor's prior realized gains exceeding 5%. Column (1) excludes and Column (2) includes investor fixed effects. Standard errors are clustered by investor; t-statistics in parentheses. The unit of observation is the investor–month; the sample period is July 2013 to February 2016. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

negative coefficient rejects the specific overconfidence channel in which realized gains dilute the measured disposition effect through heavier trading; the recall model itself makes no prediction for total turnover, which bundles purchases, loser sales, and portfolio growth.

6 Tests of Additional Predictions

Beyond ruling out alternatives, the memory mechanism carries further implications that we can take to independent variation and to other asset classes. We examine two: how the experience effect varies with investor age, and how it should differ for mutual funds.

6.1 Age and the Stock of Competing Memories

Retrieval in the model is competitive, and that yields a further prediction about who responds most. By Theorem 2, the pull of any one outcome falls as the stock of episodes competing for recall grows. A given realized gain or loss should therefore weigh most on investors who have accumulated the fewest competing memories, so the experience effect should be strongest where trading histories are shortest. Age is our proxy for the size of that stock. This follows the experience-weighting logic of Malmendier and Nagel (2011), who show that individuals overweight recent personal experiences when forming beliefs, and that the young, for whom a recent outcome is a larger share of lifetime experience, respond to it more strongly. Korniotis and Kumar (2011) document a complementary life-cycle channel, in which experience-based investment knowledge is increasingly offset by cognitive aging. We test the prediction by re-estimating the full specification separately above and below the median investor age of 36.

Table 11 reports the estimates. The age gradient is clear on both sides. Among younger investors, an additional large win attenuates the disposition effect more than twice as strongly as among older investors (*Gain* $\times$ *Win Experience* of -0.0354 versus -0.0154), and an additional large loss amplifies it more than twice as strongly (*Gain* $\times$ *Loss Experience* of $+0.0402$ versus $+0.0160$). All four interactions are statistically significant and carry the expected sign. The pattern reinforces the view that younger investors—for whom a given outcome looms larger relative to a shorter lifetime of competing experience—are more responsive to accumulated trading outcomes.

Table 11: Experience-Driven Disposition Effect by Age Group (5% Threshold)

Dependent Variable:	**Sale × 100**			
	Age ≤ 36		**Age > 36**	
$\|X\| > 5\%$	**(1) Main**	**(2) Full**	**(3) Main**	**(4) Full**
Gain	2.3935***	2.2976***	2.5789***	2.5874***
	(4.20)	(3.28)	(5.20)	(4.91)
Loss Experience	0.0172***	0.0037	0.0180*	0.0124
	(3.92)	(0.72)	(1.85)	(1.11)
Gain × Loss Experience		0.0402***		0.0160***
		(3.62)		(2.66)
Win Experience	0.0329***	0.0425***	-0.0070	-0.0026
	(6.69)	(7.44)	(-0.78)	(-0.25)
Gain × Win Experience		-0.0354***		-0.0154**
		(-5.91)		(-2.55)
Constant	8.3169***	8.4253***	8.1570***	8.1998***
	(40.31)	(34.95)	(37.11)	(36.68)
Observations	19,396,148	19,396,148	24,253,718	24,253,718
R-squared	0.145	0.145	0.133	0.133

Notes. This table re-estimates the disposition-effect regression separately for younger (age ≤ 36) and older (age > 36) investors, split at the median investor age of 36. Columns (1) and (3) include the *Loss Experience* and *Win Experience* main effects only; Columns (2) and (4) add both *Gain* × experience interactions, matching the full specification of Table 3, Column (5). *Loss Experience* and *Win Experience* count prior trades with realized returns below −5% or above +5%, respectively, relative to the average purchase price. *Gain* equals 1 if the position has an unrealized gain at the decision time. All specifications include investor, stock, and time fixed effects, and standard errors are clustered by investor, time, and stock. The sample period is July 2013 to February 2016. *t*-statistics are in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

6.2 Cross-Asset Evidence: Mutual Funds

The mechanism also carries a cross-asset implication. Within the model, the disposition effect feeds on the right tail of the return distribution: large gains dominate retrieval under the "stocks I own" cue precisely because individual-stock returns are heavily right-skewed (paper-return skewness ≈ 23 in our sample). Diversification compresses that tail, so the model predicts a weaker disposition effect for mutual funds.

Delegation plausibly pushes the same way, muting the self-attribution that ordinarily reinforces holding loss positions, though this channel lies outside the formal model. Together the two point to a weaker, and potentially reversed, disposition effect for funds.

The evidence lines up. Chang, Solomon and Westerfield (2016*a*) document, within investor, a positive disposition effect for individual stocks alongside a negative one for mutual funds held by the same investors, and interpret the reversal through a delegation-based blame-attribution channel that complements the cued-recall account.[19] The retrieval chan-

[19]Corroborating evidence on the weaker disposition effect for mutual funds and other delegated or diversi-

nel thus predicts attenuation where the right tail is compressed; the sign reversal itself is evidence for the complementary delegation channel rather than an implication of our baseline model.

7 Conclusion

Using transaction records for 189,530 Chinese retail investors over the 2013–2016 boom and bust, we show that the content of past trading experience—not merely its accumulation—governs how the disposition effect evolves. The bias does not simply fade with activity: realized gains weaken it, while realized losses sharpen it. A two-standard-deviation increase in accumulated large losses raises the propensity to sell winners over losers by about 23% of the baseline disposition effect, an equivalent increase in gains lowers it by about 21%, and the loss-side amplification carries the opposite sign from the net mechanical force that portfolio turnover could produce. What moves the bias is the valence of past realized outcomes, not merely their number—the signature a memory account predicts.

To organize these facts, we develop a memory-based cued-recall model that nests the classic disposition effect. Each trade is stored as an episode—its return, the action taken, and the surrounding context—and the cue "stocks I own" retrieves past episodes in proportion to their similarity and frequency, each competing with the others for recall. Because large gains occur more often than large losses, gain memories surface more readily and generate the baseline bias even for novices; realizing a gain then re-tags it as disposed rather than owned, muting its later pull on recall, whereas realizing a loss reinforces it. A single retrieval channel reproduces the baseline effect, its decay after gains, and its amplification after losses, with no shift in preferences and no reference-dependent utility parameters.

The findings also caution against reading experience as uniform learning. Studies conducted mostly in rising markets may have captured the gain channel while obscuring the reinforcing role of losses. Heterogeneity in experience—not only in preferences or risk tolerance—

fied vehicles appears in Bailey, Kumar and Ng (2011), Ivković and Weisbenner (2009), and Calvet, Campbell and Sodini (2009). The cued-recall model unifies the two channels emphasized in this literature—return-distribution shape and decision delegation—as joint determinants of how strongly large gains dominate the recall set under an "investments I own" cue.

generates systematic variation in trading behavior, and that variation is largest where a memory account implies it should be: among younger investors, for whom each outcome is a larger share of a short history.

Two implications follow. First, the win-side effect is about holding, not overconfident churning. One might expect a run of gains to breed overconfidence and heavier trading; instead, winning experience lowers turnover, the opposite of the classic over-trading signature (Barber and Odean, 2000*b*; Statman, Thorley and Vorkink, 2006). Investors who have won realize their gains less readily and hold winning positions longer—the retrieval channel our model formalizes, in which a realized gain is re-tagged as disposed and fades from what comes to mind. Second, Gödker, Jiao and Smeets (2025) show that selectively recalled gains shape beliefs; pairing transaction records with elicited beliefs would test directly how remembered outcomes move return expectations—the natural next step, and a bridge between behavioral finance and the psychology of memory.

References

An, Li, Joseph Engelberg, Matthew Henriksson, Baolian Wang, and Jared Williams. 2024. "The Portfolio-Driven Disposition Effect." *The Journal of Finance*, 79(5): 3459–3495.

Bailey, Warren, Alok Kumar, and David Ng. 2011. "Behavioral biases of mutual fund investors." *Journal of Financial Economics*, 102(1): 1–27.

Barber, Brad M, and Terrance Odean. 2000*a*. "Trading is hazardous to your wealth: The common stock investment performance of individual investors." *The Journal of Finance*, 55(2): 773–806.

Barber, Brad M., and Terrance Odean. 2000*b*. "Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors." *The Journal of Finance*, 55(2): 773–806.

Barberis, Nicholas. 2012. "A Model of Casino Gambling." *Management Science*, 58(1): 35–51.

Barberis, Nicholas, and Wei Xiong. 2012. "Realization utility." *Journal of Financial Economics*, 104(2): 251–271.

Barrot, Jean-Noel, Ron Kaniel, and David Sraer. 2016. "Are retail traders compensated for providing liquidity?" *Journal of Financial Economics*, 120(1): 146–168.

Ben-David, Itzhak, and David Hirshleifer. 2012. "Are investors really reluctant to realize their losses? Trading responses to past returns and the disposition effect." *The Review of Financial Studies*, 25(8): 2485–2532.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini. 2009. "Fight or Flight? Portfolio Rebalancing by Individual Investors." *The Quarterly Journal of Economics*, 124(1): 301–348.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2011. "Robust inference with multiway clustering." *Journal of Business & Economic Statistics*, 29(2): 238–249.

Chang, Tom Y., David H. Solomon, and Mark M. Westerfield. 2016*a*. "Looking for Someone to Blame: Delegation, Cognitive Dissonance, and the Disposition Effect." *The Journal of Finance*, 71(1): 267–302.

Chang, Tom Y, David H Solomon, and Mark M Westerfield. 2016*b*. "Looking for someone to blame: Delegation, cognitive dissonance, and the disposition effect." *The Journal of Finance*, 71(1): 267–302.

Dhar, Ravi, and Ning Zhu. 2006. "Up close and personal: Investor sophistication and the disposition effect." *Management science*, 52(5): 726–740.

Feng, Lei, and Mark S Seasholes. 2005. "Do investor sophistication and trading experience eliminate behavioral biases in financial markets?" *Review of finance*, 9(3): 305–351.

Frydman, Cary, and Baolian Wang. 2020. "The Impact of Salience on Investor Behavior: Evidence from a Natural Experiment." *The Journal of Finance*, 75(1): 229–276.

Frydman, Cary, Samuel M Hartzmark, and David H Solomon. 2018. "Rolling mental accounts." *The Review of Financial Studies*, 31(1): 362–397.

Gennaioli, Nicola, and Andrei Shleifer. 2010. "What Comes to Mind." *The Quarterly Journal of Economics*, 125(4): 1399–1433.

Gödker, Katrin, Peiran Jiao, and Paul Smeets. 2025. "Investor Memory." *The Review* *of Financial Studies*, 38(6): 1595–1640.

Hartzmark, Samuel M. 2015. "The worst, the best, ignoring all the rest: The rank effect and trading behavior." *The Review of Financial Studies*, 28(4): 1024–1059.

Howard, M. W., and M. J. Kahana. 2002. "A distributed representation of temporal context." *Journal of Mathematical Psychology*, 46(3): 269–299.

Ivković, Zoran, and Scott Weisbenner. 2009. "Individual investor mutual fund flows." *Journal of Financial Economics*, 92(2): 223–237.

Jiang, Zhengyang, Hongqi Liu, Cameron Peng, and Hongjun Yan. 2024. "Investor Memory and Biased Beliefs: Evidence from the Field." National Bureau of Economic Research Working Paper 33226.

Kahana, Michael Jacob, James D. Paron, and Jessica A. Wachter. 2025. "Associative learning and representativeness." Working paper, University of Pennsylvania.

Korniotis, George M, and Alok Kumar. 2011. "Do older investors make better investment decisions?" *The review of economics and statistics*, 93(1): 244–265.

List, John A. 2003. "Does market experience eliminate market anomalies?" *The Quarterly Journal of Economics*, 118(1): 41–71.

Lohnas, L. J., S. M. Polyn, and M. J. Kahana. 2015. "Expanding the scope of memory search: Modeling intralist and interlist effects in free recall." *Psychological Review*, 122(2): 337–363.

Luce, Duncan. 2000. *Utility of Gains and Losses.* New York:Psychology Press, Taylor and Francis Group.

Malmendier, Ulrike, and Jessica A Wachter. 2021. "Memory of past experiences and economic decisions." *Available at SSRN 4013583.*

Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression babies: Do macroeconomic experiences affect risk taking?" *The quarterly journal of economics*, 126(1): 373–416.

Odean, Terrance. 1998. "Are investors reluctant to realize their losses?" *The Journal of finance*, 53(5): 1775–1798.

Oh, Sangmin S, and Jessica A Wachter. 2022. "Cross-Sectional Skewness." *The Review of Asset Pricing Studies*, 12(1): 155–198.

Seru, Amit, Tyler Shumway, and Noah Stoffman. 2010. "Learning by trading." *The Review of Financial Studies*, 23(2): 705–739.

Shefrin, Hersh, and Meir Statman. 1985. "The disposition to sell winners too early and ride losers too long: Theory and evidence." *The Journal of finance*, 40(3): 777–790.

Statman, Meir, Steven Thorley, and Keith Vorkink. 2006. "Investor Overconfidence and Trading Volume." *The Review of Financial Studies*, 19(4): 1531–1565.

Appendix A Additional Information

A1 Calculating Trading Returns

Transaction-Level Returns We follow the return methodology of Hartzmark (2015) and Frydman, Hartzmark and Solomon (2018): the realized return on a sale is the difference between the sale price and the average purchase price of the position. When additional shares are acquired after the initial purchase, we measure the return against the volume-weighted average of all purchase prices for the position.

For each buy order b and sell order s of each stock j at time t, the realized return on a sale is

$$return(s,j)_t = \frac{vol(s,j)_t \times (price(s,j)_t - \bar{price}(b,j)_t)}{cost(b,j)_t} \tag{16}$$

where $vol(s,j)_t$ is the volume of stock j sold at time t, $price(s,j)_t$ is the selling price, $\bar{price}(b,j)_t$ is the average purchase price of stock j up to time t, and $cost(b,j)_t$ is the cumulative cost of purchasing stock j up to time t:

$$\bar{price}(b,j)_t = \frac{holding_{j,t-1} \times \bar{price}(b,j)_{t-1} + price(b,j)_t \times vol(b,j)_t}{holding_{j,t-1} + vol(b,j)_t} \tag{17}$$

$$holding_{j,t} = \sum_{t=0}^{t} [vol(b,j)_t - vol(s,j)_t] \tag{18}$$

$$cost(b,j)_t = \sum_{t=0}^{t} [vol(b,j)_t \times price(b,j)_t] \tag{19}$$

For example, suppose five transactions comprise three buy orders and two sell orders. The sell orders leave $cost(b,j)_t$ unchanged, since $vol(b,j)_t = 0$ on those dates; only the buy orders move it. The five transactions yield two realized returns, one per sell order. A stock is treated as new if an investor, having previously held shares in it, sells the entire position before buying again.

This procedure yields a return for every sale. The same average-purchase-price reference also defines a paper return for every position an investor continues to hold. Paper returns, though more often negative than positive, are strongly right-skewed (skewness $\approx$ 23)—the input the recall model of Section 4 acts on. Realized returns are far more symmetric, since investors close positions from across the distribution.

Stock-Level Returns We also compute a total return at the stock level, from the first purchase to the final sale that liquidates the holding:

$$return_j = \frac{\sum_{t=0}^{t}[vol(s,j)_t \times price(s,j)_t]}{\sum_{t=0}^{t}[vol(b,j)_t \times price(b,j)_t]} - 1 \tag{20}$$

This is the net gain from trading stock j relative to its total purchase cost.

Other Approaches Barrot, Kaniel and Sraer (2016) consolidate trades on a per-individual, per-stock, and per-day basis, sourcing daily stock returns from EUROFIDAI. This approach simplifies the computation of stock returns, enabling the subsequent analysis to regress daily abnormal returns for each stock.

A2 Empirical Strategies for Quantifying Investor Experience

We construct trading episodes following Seru, Shumway and Stoffman (2010). Each investor's history is segmented into episodes as follows:

1. **Initiation of a Trading Episode:**
   - An episode begins when the investor purchases shares of a stock.

2. **Termination of a Trading Episode:**
   - An episode ends at the first sale, partial or complete, of those shares.
   - The realized return at that sale classifies the episode as a winning or losing experience.

3. **Implications of Partial Sales:**
   - A partial sale ends the current episode.
   - The remaining shares begin a new episode from that point.

Segmenting each trading history into episodes gives a clean measure of investor experience and lets us relate past outcomes to the current sell decision.

A3 Mechanical Bias from Portfolio Composition

Our experience measures count an investor's own past realized gains and losses, so they are built from the same portfolio process that generates the current selling decision. Two features of that process could reach the estimated interactions mechanically. The first is scale: as experience accumulates the book grows, and the probability of selling any one position falls with the number of positions, compressing the winner–loser sale gap on both sides. The second is composition: realized trades thin the book, and realizing gains leaves fewer winners. Under random selling, fewer winners lower the unconditional probability that the next sale is a winner, but not the probability of sale conditional on being a winner, which is how the disposition effect is measured. With a baseline disposition effect, however, there is less competition among the remaining winners, and the probability of sale conditional on being a winner rises, pushing the Gain $\times$ WinExp interaction upward, against our negative estimate. The same arithmetic applied to realized losses thins the losers and pushes the Gain $\times$ LossExp interaction upward, toward our positive estimate. This appendix measures the net effect of these channels.

Three features of the design and the data bound this channel.

Portfolios evolve between sales. The composition channel requires that the thinning of winners or losers at t persists to the sell decision at $t+1$. In practice, both prices and holdings evolve substantially between consecutive trades. Price movements reclassify positions from gains to losses (and vice versa), while new purchases add stocks whose gain/loss status is unrelated to the stock sold at t. In our data, the median interval between consecutive sales by the same investor is approximately 12 trading days, during which the typical stock experiences price changes large enough to alter its gain/loss classification. Consequently, a sale at time t need not systematically shift the winner–loser composition at $t+1$.

Identification is within-investor. Our main specification (equation 3) includes investor fixed effects, which absorb any time-invariant tendency of an investor's portfolio to be tilted toward gains or losses. The identifying variation comes from within-investor changes in experience interacted with within-investor changes in the gain indicator. For the mechanical channel to bias the experience–gain interaction coefficients, it would need to generate a systematic within-investor correlation between the cumulative count of past large gains (losses) and the conditional probability of the next sale being a winner, net of investor-level means. In our data this within-investor co-movement is substantial: the within-investor correlation between portfolio size and experience is +0.61 for winning experience and +0.58 for losing, and the correlation between the held winner share and experience is -0.26 and -0.16, respectively. Rather than assume the mechanical channel away, we calibrate the simulation below to reproduce these correlations and measure the resulting bias directly.

A calibrated simulation. To quantify the mechanical channel directly, we conduct a Monte Carlo simulation described in detail below. We simulate 5,000 investors whose number of trades is drawn from a right-skewed distribution (mean $\approx$ 27), and we calibrate portfolio sizes, experience counts, and—critically—both within-investor co-movements that drive the mechanical channel to the data: the correlation between experience and portfolio size (+0.61 winning, +0.58 losing) and the correlation between experience and the held winner share (−0.26 and −0.16; untabulated calibration moments). Under three null scenarios—random selling, a fixed (experience-independent) disposition effect, and a fixed disposition effect combined with a V-shaped selling schedule—we estimate the interaction specification of Table 3, Column (5), adapted to the simulated panel; any nonzero interaction is, by construction, mechanical.

Monte Carlo Assessment of Mechanical Bias

This subsection details the simulation. Selling is generated under null rules independent of past experience, so that any nonzero Gain $\times$ Experience coefficient is, by construction, a mechanical artifact of portfolio composition.

Simulation design. We simulate 5,000 investors and repeat the exercise 200 times. The number of sell decisions per investor is drawn from a log-normal distribution (median $\approx$ 9, mean $\approx$ 27, with a heavy right tail capped at 600), reproducing the dispersion and skew of the experience counts in Table 1; a fixed count per investor cannot match these moments. Each investor starts with about five stocks and, after each sale, buys back toward a size target that drifts upward over the trading sequence. This growth is the feature that lets the simulated panel reproduce both the mean portfolio size (7.6) and the substantial within-investor correlation between experience and portfolio size in the data—+0.61 for winning experience and +0.58 for losing; a constant target produces a within-investor correlation near zero. The held winner share, in turn, falls as experience accumulates (−0.26 and −0.16 in the data): we reproduce this with a bounded market level tied to each investor's position in her own trading sequence—a gentle boom-to-bust path that lowers the winner share late in the sequence. Tying the decline to sequence progress rather than calendar time caps its total size, so it shifts composition without manufacturing a fat left tail in returns. Between sell decisions, prices evolve over approximately five trading days as a log-normal process with two independent jump components—frequent moderate jumps and rare large jumps (crashes and spikes)—calibrated to the mean counts of winning and losing experiences at the 5% and 30% thresholds in Table 1. Algorithm 1 summarizes the data-generating process for a single investor.

Algorithm 1. Data-generating process for one investor

1. Draw the number of sales $n = \text{round}\big(\text{LogNormal}(2.20,\ 1.48)\big)$, capped to $[1, 600]$.
2. Draw the starting size $k_0 = \max\big(3,\ \text{round}(5 + \mathcal{N}(0, 0.6^2))\big)$; open $K = k_0$ positions at the purchase price. Draw the investor's market-decline amplitude a_i (mean 0.08).
3. For each sell event $\tau = 1, \ldots, n$:
   (a) If $K < 3$, replenish to k_0 positions.
   (b) Evolve prices for ≈ 5 days: $p \leftarrow p \cdot \exp\big(\mathcal{N}(\mu, \sigma^2) + J^+_{\text{mod}} + J^+_{\text{big}} + J^-_{\text{mod}} + J^-_{\text{big}}\big)$, with independent moderate and rare-large jumps.
   (c) Scale prices to the bounded market level $M(\tau) = \exp\big(-a_i(\tau-1)/(n-1)\big)$, applied incrementally relative to $\tau - 1$. Tying the level to sequence progress caps the total move at a_i, so the held winner share declines with experience.
   (d) Classify each holding as a gain or loss from its unrealized return.
   (e) Assign sell weights—**(A)** equal; **(B)** 1.5$\times$ for gains; **(C)** rule (B)$\times(1 + 0.8\,|r|/0.10)$—and sell one position with probability proportional to its weight.
   (f) Record all current holdings: the sale indicator, the gain dummy, and the running counts WinExp and LossExp.
   (g) If the realized return exceeds $\pm 5\%$ (resp. $\pm 30\%$), increment the matching experience counter.
   (h) Remove the sold position; set the target $= \min\big(k_0 + \min(0.22(\tau-1),\ 2.5),\ 40\big)$ and stochastically buy back toward it.

We consider three null scenarios:

(A) **Random selling.** Each stock in the portfolio is equally likely to be sold, regardless of whether it is at a gain or loss. This isolates the pure portfolio-composition channel.

(B) **Fixed disposition effect.** Winners are sold with 1.5$\times$ the probability of losers, but this ratio is constant across all experience levels. This tests whether a standard (experience-independent) disposition effect, combined with the mechanical channel, can generate the estimated interaction terms.

(C) **Fixed DE with V-shaped selling.** Following Ben-David and Hirshleifer (2012), we additionally weight stocks with extreme absolute returns more heavily in the sell decision, while keeping the selling rule independent of accumulated experience. This is the most conservative null.

For each simulated dataset, we estimate the Column (5) specification of Table 3 adapted to the simulated panel:

$$\begin{aligned}\text{Sale}_{i,j,t} \times 100 = \alpha_i &+ \beta_1 \,\text{Gain}_{i,j,t} + \beta_2 \,\text{LossExp}_{i,t} + \beta_3 \,\text{WinExp}_{i,t} \\ &+ \beta_4 \,\text{Gain}_{i,j,t} \times \text{LossExp}_{i,t} + \beta_5 \,\text{Gain}_{i,j,t} \times \text{WinExp}_{i,t} + \varepsilon_{i,j,t}\end{aligned}$$

with investor fixed effects and standard errors clustered by investor. The empirical specification adds stock and time fixed effects with three-way clustering; the simulated panel has no meaningful stock or calendar dimension, so those terms would absorb nothing by construction. The coefficients of interest are β_4 and β_5. Under all three null scenarios, any nonzero estimate is purely mechanical.

Results. The main-text Table 6 reports the mean simulated coefficient, its 95% Monte Carlo confidence interval, and the Monte Carlo p-value across 200 replications; Figure 4 plots the corresponding distributions at the 5% threshold. The one-sided p_{MC} gives the fraction of simulations producing a coefficient at least as extreme as the empirical estimate in the direction of the paper's finding.

Interpretation. None of the calibrated nulls produces the loss-side gradient. Under both disposition-effect nulls the mean Gain × LossExp coefficient is negative (−0.0064 fixed-DE, −0.0068 V-shape)—the opposite sign from our estimate of +0.0212 ($t = 3.28$)—and it is negligible under random selling (+0.0006). A channel that pushes the loss-side interaction negative or to zero cannot manufacture a strongly positive empirical coefficient; our loss-side finding is obtained despite this drag, so the true behavioral effect is, if anything, larger than the raw estimate. In no scenario does more than one simulation in forty ($p_{\text{MC}} \leq 0.025$) reach the empirical value.

The win-side gradient is also largely behavioral. The data exhibit two offsetting within-investor co-movements: portfolios grow with experience (correlation +0.61), which pushes the mechanical Gain × WinExp coefficient negative, while the held winner share falls with experience (correlation −0.26), which pushes it positive. With both calibrated to the data they nearly cancel: the mean mechanical coefficient is −0.0032 (Scenario B) and +0.0019 (Scenario C), far short of our estimate of −0.0198, and fewer than one simulation in twenty ($p_{\text{MC}} \leq 0.045$) reaches the empirical value. The win-side attenuation we estimate is thus statistically distinguishable from the mechanical null and is left overwhelmingly to behavior.

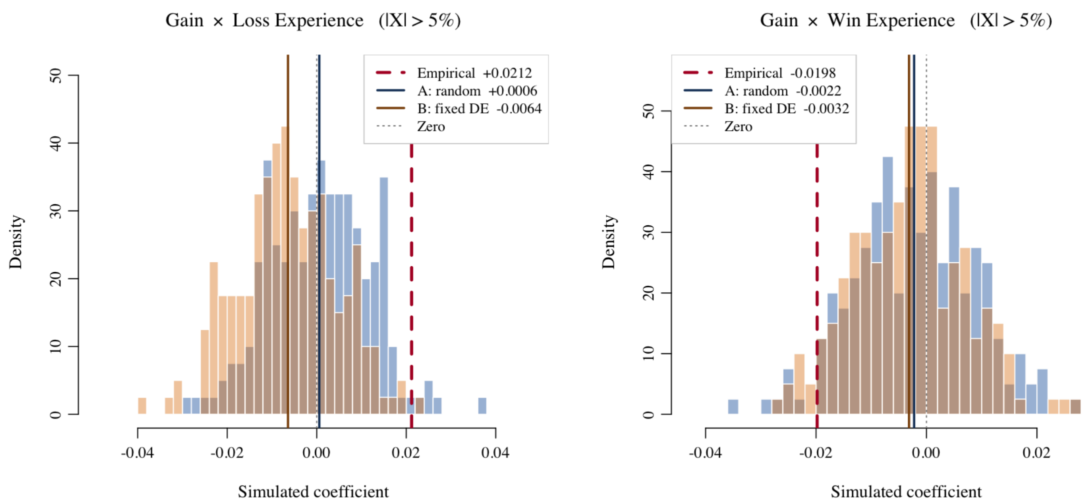

Figure 4: Distribution of Simulated Gain $\times$ Experience Coefficients ($|X| > 5\%$)

Notes. The two panels show the distribution of the Gain × Loss Experience (left) and Gain × Win Experience (right) coefficients across 200 Monte Carlo replications at the $|X| > 5\%$ threshold, under the data-matched calibration (within-investor experience–portfolio-size correlation $+0.61$ and experience–winner-share correlation -0.26). Blue histograms correspond to Scenario A (random selling); orange histograms to Scenario B (fixed disposition effect); Scenario C (fixed DE + V-shape) is reported in Table 6 but not plotted here. Dashed red lines mark the paper's point estimates from Table 3. Both mechanical distributions are centered near zero—the loss-side slightly negative, opposite the empirical $+0.0212$, and the win-side near zero, far short of the empirical -0.0198—so each empirical estimate falls well into the tail of what portfolio arithmetic alone produces.

A calibration that matched only the size correlation, omitting the offsetting composition correlation, would overstate the mechanical bias—roughly half the estimate—and understate this conclusion. Under random selling the coefficient is -0.0022, so most of the already small mechanical bias appears with no disposition effect at all. It is a specification artifact. With equal sale probabilities the true winner–loser gap is zero at every point; the per-position sale probability, however, declines with experience along a curve, steep while the book is small and flat once it is large, while the regression fits straight lines to gain and loss observations separately. Because the winner share falls as experience accumulates, gain observations concentrate on the steep part of the curve, and their fitted slope is slightly steeper. The artifact is embedded in all three null distributions, so the comparisons with the empirical estimates already account for it.

No null reproduces our central pattern, a strongly positive loss-side coefficient. Across all three rules the mechanical Gain $\times$ LossExp estimate stays negative or negligible (-0.0068 to $+0.0006$), never approaching $+0.0212$: on the loss side the portfolio-growth channel dominates the offsetting composition channel and pushes the interaction negative. No experience-independent selling rule generates a positive loss-side interaction of the magnitude we estimate. This opposite-signed response, the disposition effect strengthening after losses while weakening after gains, is the clearest evidence on this point: it is a behavioral signature, not an artifact.

The same qualitative pattern holds at the 30% threshold: the mechanical Gain $\times$ LossExp remains negative against a strongly positive empirical estimate ($+0.0907$), while the weaker empirical win-side estimate (-0.0634, $t = -1.79$) and the wide sampling spread there lead us to draw no conclusion from the 30% win-side comparison.

A portfolio-composition channel, even under a DGP calibrated to both the growth in portfolio size and the decline in the held winner share that experience brings in our data, and even with the V-shaped selling propensity of Ben-David and Hirshleifer (2012), does not generate the loss-side amplification under any null we consider (the bias has the wrong sign) and accounts for at most a sixth of the win-side attenuation. Both effects are statistically distinguishable from the mechanical null ($p_{\mathrm{MC}} \leq 0.045$); the loss-side response is the sharpest, since there the mechanical force runs actively opposite to what we estimate.