The Watchlist Trap*

Minxia Chen†

This draft: July 31, 2026

Updated Frequently: Most Recent Version Here ↗

Abstract

A watchlist records a form of attention prior work could not see. Where the literature measures attention in its *transient* form—a search spike, a volume jump, a login, caught near the trade and gone within the week—the watchlist reveals its *persistent* form: a self-authored, dated commitment to follow a stock, returned to over time, before any purchase. Using 28,990 Chinese investors whose every purchase is linked to a timestamped watchlist entry, I ask whether routing a purchase through this consideration stage helps. It need not. The stocks an investor watches but does not buy outperform the stocks she buys; yet among the stocks she buys, the ones she watched first (*warm* purchases) underperform the ones she buys on sight (*cold* purchases) by about one percentage point of characteristic-adjusted return within investor–day pairs, and by more against never-watched purchases. The list's signal is good; the damage is in the conversion from watching to buying. I take the gap apart in two stages. First, *composition*: warm purchases carry more of the attention-type speculative tendencies this literature prices as costly, which accounts for roughly 45% of the gap. Second, the mechanism that produces the exposure—*reference-dependent conversion*: the daily first-purchase hazard changes slope sharply as the price crosses the level at which the stock was added to the list and the highs and lows witnessed while it waited, identified within (stock, day) cells, so it reflects the investor's own reference point; what investors read does not carry the gap. Extending the trading-tendencies program of Han et al. (2025) from *which* stock retail investors buy to *when* they buy it, I locate a cost that registry and execution data cannot see: not in watching, but in the reference-dependent conversion the watchlist sets up.

*Chen gratefully acknowledges support from the Ian Potter '93D PhD Award at INSEAD.

†INSEAD, 1 Ayer Rajah Avenue, Singapore 138676, Singapore; Email: minxia.chen@insead.edu.

1 Introduction

A retail investor cannot evaluate the thousands of stocks she might buy, so attention does the first work of choosing: it forms the small *consideration set* she deliberates over, and preferences act only within it (Barber and Odean, 2008). How that set is formed and carried over time has stayed largely unobserved, because attention has been measured only in its *transient* form—a search spike, a jump in volume, a login (Da et al., 2011; Gargano and Rossi, 2018), caught near the moment of the trade and gone within the week. A watchlist reveals its *persistent* form: a self-curated roster of stocks, each entry stamped with its date, displayed beside the trade panel, that an investor *returns to*—a stock added to it stays until she removes it, and she rarely does. It is attention of a different *type*: persistent where prior measures are transient, self-authored, and pre-purchase—not merely more of it. The list therefore records the *consideration stage* a purchase passes through, or skips. Using 28,990 Chinese investors whose every purchase is linked to a timestamped watchlist entry—2.86 million dated entries against 745,040 first purchases—I observe this stage directly and price its effect on the trades that come out of it.

Whether passing through the stage should help or hurt is not obvious. On one view it helps: watching directs attention, invites comparison, and lets a first impression cool before money moves—paying attention has been shown to pay (Gargano and Rossi, 2018), and, observed directly, a consideration set is one within which retail investors choose well (Imas et al., 2026). On another, the stage is not neutral: writing down a stock fixes a price the investor saw, so the purchase, when it finally fires, may chase that self-set mark rather than the stock's prospects; and the attention the list sustains need not improve the decision it feeds—Hartzmark et al. (2021) show that attention to an asset one *owns* distorts how its owner learns from identical information, so channeling attention to a stock need not mean a better-timed trade. Which force dominates is an empirical question the linked watchlist can settle.

Two facts frame the puzzle. First, the list's *signal* is good: within an investor on a given day, the stocks she watches but does not buy outperform the stocks she buys, which in turn outperform the stocks she watched before buying—a staircase of about 2.2 percentage points from top to bottom. Whatever draws stocks into consideration works. Second, and against the natural prior, among the stocks she buys those she *watched* first—I call these *warm* purchases—underperform those she buys *cold*, with no prior watching, by about one percentage point of characteristic-adjusted return within investor–day pairs, and by more against never-watched purchases. The damage is not in watching; it is in the *conversion* from watching to buying: what an investor watches beats what she buys from it, and the loss sits in that conversion, not in the watching.

Why does converting a watched stock into a purchase go wrong? My answer is that conversion is *reference-dependent*. Estimating the daily first-purchase hazard on the panel of watched stock-days, I find it changes slope sharply as the price crosses the level at which the stock was added to the list, and as it crosses the highs and lows the investor witnessed while it waited. Because the hazard is identified within (stock, day) cells—so that every quantity common to a stock on a date, including its speculative-tendency profile and its recent momentum, is absorbed—the kink

reflects the individual investor's own reference point, formed from when *she* started watching, not which stocks are watched or how they were trending. How much investors read about a stock, by contrast, does not carry the gap. The mechanism is a property of *when* the purchase fires, not of information; I read it as consistent with reference dependence in the timing of conversion, of which anchoring on the add price is one case, and take no stand on whether it reflects loss aversion.

Reference-dependent conversion has a cost because of *where* in a stock's path it places the purchase. Following the trading-tendencies program of Han et al. (2025), I flag each purchase for the speculative situations retail investors are documented to buy into—each tendency implemented from its original source—and find that warm purchases carry more of the attention-type tendencies (recent limit-ups, proximity to highs, lottery-like payoffs) that this literature shows are costly. This tendency composition accounts for roughly 45% of the warm–cold gap. The reference-dependent conversion is what produces the exposure: it releases the purchase the list has been holding into local heat.

Contributions. First, measurement: where clickstream evidence observes a within-session consideration set (Imas et al., 2026), the watchlist reveals consideration as a *persistent, dated commitment*—and, uniquely, the purchases that skip it (the cold buys)—linking dated entries, news reading, and platform activity to every first purchase, so that the *effect* of routing a purchase through consideration, not only selection within it, becomes observable at scale. Second, a mechanism for a puzzle the executions literature could not pose: the watchlist's signal is good, yet buying from it underperforms, because conversion from watching to buying is reference-dependent and places the purchase at attention-heavy moments. Third, an extension of the trading-tendencies program (Han et al., 2025): where that work prices *which* stock retail investors buy, holding timing fixed, the linked watchlist lets me price *when*, and locate the cost in the reference-dependent conversion the list sets up—a stage registry and execution data cannot see.

Related literature. The paper speaks to four literatures. *Retail attention and consideration sets*: Odean (1999) and Barber and Odean (2008) established attention-driven buying from executions, where the attention is *transient*—imputed from volume or extreme returns, or measured as a search or a login (Da et al., 2011; Gargano and Rossi, 2018), contemporaneous with the trade. The watchlist instead records a *persistent*, self-curated stage that precedes it. Two papers observe consideration directly and sit closest: Imas et al. (2026) shows that within a *passively viewed* fund set retail picks well, and Levy (2025) reads the watchlist as an information set; neither observes purchases made *outside* the set, nor prices the watched-versus-unwatched purchase. Relatedly, Hartzmark et al. (2021) show that *owning* an asset channels attention to it and distorts how its owner learns from its signals; the watchlist locates this attention pull one step earlier—at the pre-ownership consideration stage—and prices it through the timing of the eventual purchase. I observe an *active* commitment with an outside, and show that consideration helps in selection but its conversion

into a trade hurts. *Reference dependence, memory, and anchoring in trading*: the disposition effect and the salience of the 52-week high and of purchase prices (George and Hwang, 2004; Huddart et al., 2009; Hartzmark, 2015) shape *selling*, and a memory-based account casts such reference points as encoded by past experience and cued at the moment of decision (Bordalo et al., 2020; Wachter and Kahana, 2024; Enke, 2026); I document a reference-dependent kink in *buying*—in the conversion hazard—with the watchlist add as the encoding event. *Trading tendencies*: the tendency battery consolidates the buy-side literature and adopts the joint design of Han et al. (2025), whose ordering (attention costliest, familiarity benign) the Chinese estimates reproduce. *Price limits*: exchanges impose daily limits as regulatory pauses (Kim and Rhee, 1997), whose attention spillovers Seasholes and Wu (2007) and Chen et al. (2019) document from executions; the watchlist reveals what the pause produces on the demand side—a persistent, salient reference point rather than a cooled impulse.

Roadmap. Section 2 describes the setting, data, and institutions. Section 3 establishes the watchlist as a persistent, self-authored form of attention and the good signal it carries. Section 4 documents the warm–cold gap and its robustness to price-state composition. Section 5 measures what warm purchases are made of, and Section 6 prices that composition—roughly 45% of the gap. Section 7 gives the reference-dependent-conversion mechanism that produces the exposure. Section 8 converts the estimates to money and concludes.

2 Setting, Data, and Institutions

The data come from a large Chinese retail trading platform. The application's home screen is a watchlist: a self-maintained roster of stocks the investor has chosen to follow, displayed alongside the trade panel from which orders are placed. Additions to the list are accepted around the clock, on trading days and non-trading days alike, and are recorded with their dates; the same accounts' trades, news-reading events, and platform messages are observed over the same window. All results concern A-share common stocks on the Shanghai and Shenzhen main boards and the ChiNext and STAR growth boards; B-shares, the Beijing exchange, funds, and convertibles are excluded. A-share trading operates under daily price limits (Kim and Rhee, 1997; Chen et al., 2019): a stock may move at most $\pm 10\%$ from the prior close on the main boards, $\pm 20\%$ on ChiNext and STAR, and $\pm 5\%$ for main-board ST stocks. When buying pressure carries the price to the upper bound and holds it there through the close, the stock ends the day *locked at the limit*: a would-be buyer cannot execute, and an investor who wants the stock can add it to the watchlist but must wait until the constraint stops binding.

Table 1 summarizes the sample. The window runs from October 2024 through December 2025 and covers 28,990 investors, who place 2,855,109 watchlist additions and make 745,040 first purchases of stocks they had never previously held. Of these first purchases, 316,612 were preceded by a watchlist addition of the same stock on a strictly earlier date and 428,428 were not. Return-

based estimation applies canonical filters: at least 252 days of price history, returns winsorised at the 1st and 99th percentiles, ST stocks excluded. Designs that condition on the watching state rather than on the purchase run on an at-risk panel of (investor, stock, day) rows — 157,347,124 rows in total — estimated after 1:10 within-stratum risk-set sampling, which keeps 250,907 first-buy events and 1,485,644 rows across 124,323 (stock × day) strata.

Table 1: Sample. Chinese A-share retail brokerage panel, 2024-10-01 to 2025-12-24

Watchlist additions	2,855,109
arriving with ≥1 same-day companion add	83.5%
placed on a limit-up day	823,420
First purchases of never-held stocks	745,040
previously watched ("warm")	316,612
not previously watched ("cold")	428,428
At-risk (investor, stock, day) panel rows	157,347,124
Investors	28,990
First-buy events in the estimation panel	250,907
Rows after 1:10 within-stratum risk-set sampling	1,485,644
(stock × day) case strata	124,323

Note: Sample. Chinese A-share retail brokerage panel, 2024-10-01 to 2025-12-24. The estimation panel applies the canonical filters (minimum 252 days of price history, winsorised returns at the 1st/99th percentiles, ST stocks excluded).

Two definitions carry through every table. A first purchase preceded by a strictly earlier watchlist addition of the same stock is *warm* ("warm"); a first purchase with no prior addition is *cold* ("cold"). Two conventions accompany the split. First, the platform automatically adds a purchased stock to the watchlist, so a same-day addition does not count as prior watching: same-day add-and-buy events are classified cold. Second, purchases whose only addition comes *after* the buy date — post-hoc tracking of an already-held position, roughly two percent of first purchases — are excluded from the sample. Each addition is further classified by the state of the stock on its addition day: (i) closed locked at the daily limit, so that execution was blocked; (ii) an ordinary open day, on which the investor could have bought and did not; (iii) a non-trading day, split by the last close before the addition into weekend and holiday additions that follow a limit-up close and those that follow a quiet eve; and (iv) made during a trading suspension — a small cell that is quarantined from every claim and shown only as an audit row.

Outcomes are cumulative abnormal returns over 20, 60, and 80 trading days from the purchase. Eight benchmark adjustments are computed — raw returns; market-adjusted returns; CAPM, Fama–French three- and five-factor residuals; the Liu et al. (2019) Chinese three- and four-factor residuals; and DGTW characteristic matching (Daniel et al., 1997) — and the tables report the characteristic-matched (DGTW) and Chinese four-factor (CH-4) adjustments as the headline pair. Unless a table states otherwise, events are weighted equally per investor and standard errors are clustered on investor, following the design conventions of Akepanidtaworn et al. (2023). Effects on binary outcomes are reported in percentage points throughout.

Three descriptive facts about the list itself frame what follows. Additions arrive in bursts:

83.5% of them come with at least one same-day companion addition. They cluster on days the stock could not be bought: 823,420 of the 2,855,109 additions — roughly twenty-nine percent — were placed on a day the stock closed locked at its upper limit. And most additions never become trades: 11.09% ever convert into a first purchase of the added stock, 9.22% within sixty trading days, and the conversions that do happen come fast, with a median lag of six trading days.

3 A Persistent, Pre-Purchase Type of Attention

Before pricing what the watchlist does, I fix what it *is*: not attention in its transient, flow form—a search spike, a volume jump, a login—but a stock the investor *wrote down* to follow, dated, and returned to. This persistent, self-authored attention carries a good signal, which frames the paper's puzzle—the signal is good, yet buying on it is where the money is lost.

Table 2 places the watchlist's own content next to the investor's executions. The design (script 391) uses investor-days on which the same investor did three things at once, to three different stocks: added a buyable stock to the watchlist (an open-day or quiet-eve addition; additions blocked at the limit or suspended are dropped from this leg), bought a stock on sight, and bought a previously watched stock. Panel A gives the sixty-day levels of the three baskets: the stocks added and not bought that day return -4.67 (DGTW), the cold purchases -4.79, and the warm purchases -6.40. Panel B takes the pairwise differences within each investor-day cell, and Panel C provides the inference by forming daily calendar-time portfolios of the same differences over 239 trading days with Newey–West standard errors: watched-not-bought beats cold by $+1.59$ (DGTW) and beats warm by $+2.17$, both significant at the one-percent level, while cold beats warm by $+0.59$ (positive but not significant in calendar time, though significant within investor-day). Cold $\equiv$ never-watched (strict); MINHIST_252 + EXCLUDEST on all three legs. Panel D repeats the first comparison on the full sample of investor-days with both a buyable addition and a cold purchase (45,431 cells; script 390): -0.10 paired, $+1.48$ in calendar time.

The ordering is the section's result. Within the same investor on the same day, the stocks written onto the list and left unbought outperform the stock bought on sight, which in turn outperforms the stock bought after a spell on the list — the leg that passed through the list before execution is the worst of the three. The sections that follow take this warm–cold difference apart: Section 4 documents the return gap and its robustness to price-state composition, Section 5 measures what warm purchases are made of, Section 6 prices that composition, Section 7 gives the reference-dependent-conversion mechanism that produces the exposure, and Section 8 converts the estimates to money.

4 The Warm–Cold Gap

Table 3 establishes the gap and its robustness to price-state composition (scripts 301 and 324). Paired within (investor, day) — the same investor buying warm and cold on the same day — warm

Table 2: The list's signal against the executions (scripts 391, 390; strict-cold, canonical refit)

	DGTW	CH-4
Panel A: sixty-day levels, three-leg cells		
Watched that day, not bought (29,612)	−4.67	−2.94
Cold purchases (12,027)	−4.79	−3.43
Warm purchases (12,350)	−6.40	−4.73
Panel B: within-cell differences (point estimates)		
Watched − cold (9,256 cells)	+0.12	+0.52
Cold − warm (9,258 cells)	+1.49	+1.29
Watched − warm (9,296 cells)	+1.59	+1.74
Panel C: calendar-time inference		
Watched − cold	+1.59***	+1.76***
	(0.22)	(0.36)
Cold − warm	+0.59	+0.20
	(0.50)	(0.48)
Watched − warm	+2.17***	+1.95***
	(0.42)	(0.35)
Panel D: two-leg comparison, full sample (script 390)		
Watched − cold, paired (45,431 cells)	−0.10	−0.04
calendar-time version	+1.48***	+1.30***
	(0.21)	(0.27)

Note: The list's signal against the executions (scripts 391, 390; strict-cold, canonical refit). Panels A–C use investor-days on which the same investor did all three things at once, to three different stocks: added a buyable stock to the watchlist (open or quiet non-trading day; additions blocked at the limit or suspended are dropped), bought a stock on sight (no prior watching), and bought a previously watched stock. Cold ≡ never-watched (strict); MINHIST_252 + EXCLUDEST at the flag anchor on all three legs. Sixty-day cumulative abnormal returns in percentage points, per-cell basket means; characteristic-adjusted (DGTW) and four-factor (CH-4) columns. Panel A: the three baskets' levels. Panel B: the pairwise differences within investor-day; by within-cell arithmetic the first two rows sum to the third. Panels A–B are point estimates from the bundle; inference is in Panel C, which forms daily calendar-time portfolios of the same differences (239 trading days; Newey–West standard errors in parentheses). Panel D repeats the two-leg comparison on the full sample of investor-days with both a buyable addition and a cold purchase (script 390; 45,431 cells; calendar-time over 239 days). The stocks the investor wrote down and did not buy beat the ones bought on sight, which beat the watched ones bought later. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

purchases underperform by -0.31, -0.59, and -1.49 percentage points (DGTW) at five, twenty, and sixty days. Because warm purchases could simply sit at different points of their stocks' price cycles, Panel A adds fixed effects for the bought stock's position in its trailing price range at the buy date, in deciles: with the sixty-day range the gap is -0.31, -0.54, -1.40; with the 252-day range, -0.27, -0.47, -1.24. The state controls move the estimates by almost nothing. Panel B re-estimates the origin split with the state control in place — limit-up-origin warm purchases at -3.97 (sixty days) against $+0.13$ for ordinary-day origins — and Panel C gives the deliberation-length split.

Table 3: The deliberation penalty is not price-state composition (script 324; strict-cold, 15-member canonical refit)

	5 days	20 days	60 days
Panel A: the gap, with and without state controls			
Warm − cold, no state control	-0.31^{**}	-0.59^{**}	-1.49^{***}
	(0.15)	(0.23)	(0.34)
+ buy-day state-decile FE (60d range)	-0.31^{**}	-0.54^{**}	-1.40^{***}
	(0.15)	(0.23)	(0.34)
+ state-decile FE (252d range)	-0.27^{*}	-0.47^{**}	-1.24^{***}
	(0.15)	(0.23)	(0.34)
Panel B: by origin of the addition, state-controlled			
Limit-up-origin warm (state-controlled)	-0.92^{***}	-1.67^{***}	-3.97^{***}
	(0.29)	(0.41)	(0.58)
Ordinary-day warm (state-controlled)	-0.11	-0.11	$+0.13$
	(0.18)	(0.29)	(0.45)
Other-origin warm (state-controlled)	$+0.27$	$+0.30$	-0.89
	(0.35)	(0.51)	(0.73)
Panel C: by deliberation length, state-controlled			
Deliberation 0–5 days	-0.40^{*}	-0.39	-0.46
	(0.22)	(0.32)	(0.46)
5–20 days	-0.19	-0.84^{*}	-2.59^{***}
	(0.31)	(0.45)	(0.69)
20–60 days	-0.26	-0.49	-2.90^{***}
	(0.32)	(0.59)	(0.84)
Over 60 days	-0.21	-0.76	-2.05^{***}
	(0.31)	(0.52)	(0.79)

Note: The deliberation penalty is not price-state composition (script 324; strict-cold, 15-member canonical refit). Paired within-(investor, day) difference between warm (warm) and cold (≡ never-watched) first purchases, DGTW-adjusted, in percentage points. State controls: deciles of the purchased stock's position in its trailing price range at the buy date (60- and 252-day windows); the DGTW benchmark additionally matches on momentum by construction. Panel B and C re-estimate the origin and deliberation-length decompositions with the 60-day state control in place (cold purchases are the base category). $N = 41,889$ at 60 days. SE clustered on investor. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The gap is not explained by warm purchases sitting at systematically different points of their price cycles: the state-decile controls leave it essentially unchanged. Section 6 asks what it *is* made of, and Section 7 why conversion places the purchase where it does.

5 What the List Holds

This section measures the composition of what gets bought after watching, using the speculative-tendency battery of Han et al. (2025): fifteen tendencies of the purchased stock's state, measurable in these data on the day before each purchase, grouped into attention-type tendencies (recent price-path landmarks such as a recent limit-up, a thirty-day run-up, or proximity to the 52-week high), contrarian-type, familiarity-type, and an income-type tendency (dividend payer). The count of active tendencies, zero to fifteen, is the outcome throughout.

Table 4, Panel A compares across investors (script 650). Cold purchases carry 4.65 tendencies on average; warm purchases carry $+0.952$ more. The difference is not uniform across addition origins: purchases whose addition was made on a limit-up close average 6.88 tendencies, a difference of $+2.259$ relative to cold purchases, while purchases whose addition was made on an ordinary open day — a voluntary deferral — carry -0.078 *fewer*. The two warm origins sit on opposite sides of the cold base.

Table 4: What the list holds: the tendency composition of warm purchases (scripts 650, 655; strict-cold, 15-member canonical refit)

	Tendency count
Panel A: across investors (script 650)	
Mean tendency count, cold purchases	4.65
Mean tendency count, blocked-at-the-limit origin	6.88
Warm − cold difference	$+0.952^{***}$
	(0.020)
blocked-at-the-limit origin vs. cold	$+2.259^{***}$
	(0.024)
voluntary-deferral origin vs. cold	-0.078^{***}
	(0.028)
Blocked − voluntary-deferral contrast	$+0.813^{***}$
	(0.148)
Panel B: within investor, confirmatory (script 655)	
Warm − cold difference, investor fixed effects	$+0.525^{***}$
	(0.017)
with month × size-tercile fixed effects	$+0.515^{***}$
	(0.017)

Note: What the list holds: the tendency composition of warm purchases (scripts 650, 655; strict-cold, 15-member canonical refit). The outcome is the number of speculative-tendency tendencies (of fifteen measurable) the purchased stock carries on the day before the purchase. Panel A compares across investors: all first purchases, cold (≡ never-watched) purchases as the base ($N = 488,863$: 209,719 cold, 111,891 blocked-at-the-limit-origin, 68,015 voluntary-deferral warm purchases; SE clustered on investor). Panel B is the pre-registered within-investor confirmation (510,956 purchases): investor and month fixed effects, so the difference is read off the same investor buying with and without prior watching. Frozen verdict: gate *C-B-composition* PASS (revalidated on the canonical sample) — threshold: within-investor difference $\geq +0.30$ tendencies with $|t| \geq 4$; recorded $+0.525$ ($t = +31.3$). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Panel B is the pre-registered within-investor confirmation (script 655; 510,956 purchases). With investor and month fixed effects, so that the contrast is read off the same investor buying with and

without prior watching, the warm–cold difference is +0.525 tendencies ($t = +31.3$), and +0.515 with month × size-tercile fixed effects. The frozen gate *C-B-composition* required a within-investor difference of at least +0.30 tendencies with $|t| \geq 4$; it passed, and passed again on revalidation.

6 The Tendency Decomposition

This is the accounting stage of the decomposition—it prices the composition and asks how much of the gap it spans, a footprint the mechanism of Section 7 then explains. The composition tilt of Section 5 does not merely describe what warm purchases contain; the tendencies carry a return price, so pricing them decomposes the warm–cold gap. I estimate the prices on the fifteen-member battery at the paper's headline convention—characteristic-adjusted (DGTW) returns at sixty trading days, with never-watched purchases as the baseline. The Chinese four-factor companions to the three exhibits below are Appendix Tables A2, A3, and A4.

Table 5 prices the battery on *all* first purchases: each additional tendency a purchase carries costs a significant characteristic-adjusted return, the population-level analog of the catalog evidence in Han et al. (2025) (the per-member mapping to the source literature is Appendix Table A1).

Table 6 cuts the sample into its exact three-way partition—warm, same-day (the platform's automatic-addition mixture), and never watched—showing that every group pays for tendencies while warm purchases both carry the most and pay the steepest per-flag price.

Table 7 carries the headline decomposition. Against the never-watched baseline the warm–cold gap is −234 basis points at sixty days ($N = 401{,}814$; two-way investor-and-day $t = -11.3$), and the tendency basket explains 45% of it. The fifteen members absorb the buy-day price state: adding deciles of the stock's trailing sixty- and 252-day price position at the purchase date as fixed effects moves the tendency-adjusted gap only from −88 to −90 basis points, so the difference is not warm purchases sitting at a different point of the price cycle but the tendencies they carry.

The gap, then, is largely a composition effect: warm purchases load on the attention-type tendencies this literature prices as costly. Composition is the footprint, not the cause; Section 7 asks what *produces* the exposure.

7 Mechanisms

The composition is a footprint; this section asks what *produces* it. Candidate mechanisms fall into four groups: the deferral motive that sets up the wait, the attention that shapes beliefs during it (as when owning an asset channels attention to it and distorts learning from its signals, Hartzmark et al. 2021), the reference dependence that fires the conversion, and the selection of who converts. The full paper adjudicates all four; the mechanism that survives and carries the exposure is reference-dependent conversion.

Why does converting a watched stock into a purchase place it at attention-heavy moments? Because conversion is *reference-dependent*. Table 8 estimates the daily first-purchase hazard on the

Table 5: The fifteen-member tendency battery priced on all first purchases (characteristic-adjusted)

	Characteristic-adjusted (DGTW), bps		
Member (share flagged)	5 days	20 days	60 days
Attention-type (chasing what's hot)			
Lottery-like (limit-up run) (21.5%)	-115.9^{***}	-212.8^{***}	-629.4^{***}
	(7.9)	(12.7)	(16.1)
Recent limit-up close (45.5%)	-113.3^{***}	-149.4^{***}	-227.6^{***}
	(6.2)	(10.4)	(14.7)
Thirty-day winner (59.2%)	-1.4	-16.4	-57.3^{***}
	(5.9)	(10.6)	(15.8)
Abnormal volume (55.7%)	-37.6^{***}	-95.9^{***}	-109.0^{***}
	(5.4)	(9.0)	(12.7)
High turnover (40.1%)	$+9.5$	-66.4^{***}	-245.0^{***}
	(6.0)	(10.3)	(14.1)
High volatility (40.6%)	-49.9^{***}	-138.0^{***}	-295.6^{***}
	(6.3)	(10.3)	(14.6)
Near the 52-week high (54.5%)	-34.8^{***}	-90.3^{***}	-149.5^{***}
	(5.6)	(9.9)	(13.8)
Salient returns (Cosemans–Frehen) (54.2%)	-44.7^{***}	-28.7^{***}	$+59.3^{***}$
	(5.7)	(9.6)	(13.4)
Contrarian-type (fishing for what's cheap)			
Thirty-day loser (8.9%)	-30.2^{***}	-72.5^{***}	-98.8^{***}
	(6.0)	(11.1)	(17.3)
Low price (9.5%)	$+80.4^{***}$	$+289.5^{***}$	$+336.6^{***}$
	(8.2)	(14.0)	(21.0)
Small cap (5.2%)	-44.1^{***}	-20.6	$+137.6^{***}$
	(9.0)	(16.9)	(24.6)
Value (top book-to-market quintile) (19.4%)	-36.9^{***}	-131.9^{***}	-120.4^{***}
	(5.8)	(9.3)	(13.6)
Familiarity-type			
Local company (8.1%)	$+12.2$	$+49.2^{***}$	$+89.6^{***}$
	(7.9)	(13.4)	(19.7)
Familiar industry (84.1%)	-13.7^{**}	-44.9^{***}	$+50.1^{***}$
	(5.7)	(9.4)	(13.0)
Income-type			
Dividend payer (63.3%)	-22.5^{***}	$+2.4$	$+59.6^{***}$
	(5.2)	(8.3)	(11.5)
Price of one additional tendency (separate specification)	-43.5^{***}	-85.3^{***}	-161.0^{***}
	(1.1)	(1.9)	(2.6)

Note: Every first purchase pays for the tendencies it carries (script 672, canonical refit): the fifteen-member battery priced on all first purchases, no warm buying split. *Specification.* Sample: all first purchases of never-previously-held stocks (745,040), keeping stocks with at least 252 trading days of history at the anchor day (716,595) and normal trading status (714,556 events; the anchor day is the last trading day strictly before the purchase). Battery: the fifteen canonical members — thirteen recorded members plus the value and salient-returns members admitted after the availability re-screen; the salient-returns construction follows Cosemans and Frehen (salience-weighted minus equal-weighted trailing daily returns, twenty-one days), not Han, He, and Weagley's exact implementation. Outcome: characteristic-adjusted (DGTW) cumulative abnormal returns over five, twenty, and sixty trading days, in basis points (levels in bps per the locked reporting-units rule); windows are complete per adjustment (an event enters a column only when every day in its window carries an adjusted return), which excludes purchases whose sixty-day window crosses the characteristic-assignment boundary (19% of full-window events; sixty-day coverage 99.5% of the feasible sample, $N = 568,039$). Estimation: weighted least squares with equal-per-fund weights. The member panel is the single joint specification $\mathrm{CAR}^{h}_{i,s,t} = \alpha_i + \gamma_{m(t)} + \sum_{j=1}^{15} \beta_j F_{j,s,t^-} + \varepsilon_{i,s,t}$, where $\mathrm{CAR}^{h}_{i,s,t}$ is investor i's adjusted return on stock s over the h trading days after the purchase at t, α_i and $\gamma_{m(t)}$ are investor and purchase-month fixed effects, and $F_{j,s,t^-} \in \{0,1\}$ is member j measured on the anchor day; cells report the $\hat{\beta}_j$ in basis points, so each member's price is conditional on the other fourteen. The bottom row reports $\hat{\theta}$ from the *separate* specification $\mathrm{CAR}^{h}_{i,s,t} = \alpha_i + \gamma_{m(t)} + \theta\, C_{s,t^-} + \varepsilon_{i,s,t}$ with the tendency count $C_{s,t^-} = \sum_{j=1}^{15} F_{j,s,t^-} \in \{0,\ldots,15\}$ — the average cost of carrying one additional tendency, not the sum of the member coefficients. Members are grouped into the paper's four pre-registered families; Han, He, and Weagley group differently (their attention group has six members and their taxonomy has no contrarian family), and the two admitted members are assigned salient returns to attention and value to contrarian. SE clustered on investor in parentheses; two-way clustering by investor and day gives $t = -16.4$ on the sixty-day count. Member prevalence on the canonical sample in parentheses after each label. Two member results read against their home literatures: value purchases lose -120 bps at sixty days ($t = -8.8$), and the salient-returns member prices *positively* conditional on the other fourteen members ($+59$ bps, $t = +4.4$) — the raw salience tilt's damage is carried by the correlated attention members. The four-factor (CH-4) companion is Appendix Table A2; the Han–He–Weagley mapping is Appendix Table A1. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6: The fifteen-member anatomy of the purchase poles: warm versus never watched (characteristic-adjusted, sixty days)

	DGTW-adjusted, 60 days (bps)	
Member (share flagged)	Warm	Never watched
First purchases (canonical sample)	303,059	209,719
Mean tendency count (of 15, weighted)	5.32	4.60
Attention-type (chasing what's hot)		
Lottery-like (limit-up run) (21.5%)	-679.4^{***}	-526.1^{***}
	(21.4)	(30.1)
Recent limit-up close (45.5%)	-261.6^{***}	-242.0^{***}
	(20.2)	(27.0)
Thirty-day winner (59.2%)	-44.5^{**}	-69.8^{**}
	(21.9)	(28.0)
Abnormal volume (55.7%)	-85.2^{***}	-64.3^{***}
	(17.7)	(21.5)
High turnover (40.1%)	-259.3^{***}	-296.4^{***}
	(19.1)	(25.8)
High volatility (40.6%)	-293.8^{***}	-210.0^{***}
	(20.2)	(26.4)
Near the 52-week high (54.5%)	-165.6^{***}	-141.6^{***}
	(19.7)	(24.0)
Salient returns (Cosemans–Frehen) (54.2%)	$+90.8^{***}$	$+64.0^{***}$
	(18.3)	(23.9)
Contrarian-type (fishing for what's cheap)		
Thirty-day loser (8.9%)	-119.9^{***}	-100.0^{***}
	(25.2)	(28.5)
Low price (9.5%)	$+334.6^{***}$	$+362.5^{***}$
	(29.3)	(39.1)
Small cap (5.2%)	$+205.9^{***}$	$+130.1^{***}$
	(40.6)	(49.2)
Value (top book-to-market quintile) (19.4%)	-126.7^{***}	-93.4^{***}
	(18.8)	(24.7)
Familiarity-type		
Local company (8.1%)	$+101.4^{***}$	$+112.6^{***}$
	(27.2)	(33.9)
Familiar industry (84.1%)	$+37.6^{*}$	+38.0
	(20.0)	(23.7)
Income-type		
Dividend payer (63.3%)	$+82.0^{***}$	$+46.1^{**}$
	(16.2)	(20.7)
Price of one additional tendency (separate specification)	-166.2^{***}	-141.6^{***}
	(3.6)	(4.6)

Note: Both poles pay for the tendencies they carry, and warm purchases pay the steepest per-tendency price while carrying the most (script 673, canonical refit): the fifteen-member battery estimated separately on the two poles of the purchase sample — warm (added to the watchlist on a strictly earlier day) versus never watched (no addition at any date). *Specification.* Sample, battery, outcome convention, weights, fixed effects, and clustering are identical to Table 5: of 745,040 first purchases of never-previously-held stocks, 716,595 keep at least 252 trading days of history at the anchor day and 714,556 keep normal trading status (the canonical sample); the fifteen canonical members carry their pooled-sample prevalence in parentheses; outcomes are DGTW-adjusted sixty-trading-day returns with complete windows, in basis points, under equal-per-fund weights. The poles are two groups of the exact three-way partition of the canonical sample (303,059 warm + 201,778 same-day + 209,719 never watched); the same-day mixture (added and bought the same day; the platform's automatic-addition mixture) is omitted from the table: its tendency loading sits between the poles (recorded cell: intermediate) and it is excluded from the gap tables; its cells remain in the results bundle. Per-group feasible coverage 99.6% / 99.5%; estimation samples $N = 238,497/165,251$. Within each group g the member panel is the joint specification $\mathrm{CAR}^{60}_{i,s,t} = \alpha_i + \gamma_{m(t)} + \sum_{j=1}^{15} \beta^g_j F_{j,s,t^-} + \varepsilon_{i,s,t}$ and the bottom row reports $\hat{\theta}^g$ from the *separate* count specification, both in Table 5's notation — each member's price is conditional on the other fourteen, and $\hat{\theta}^g$ is the average cost of carrying one additional tendency in that group, not the sum of the member coefficients. The mean-tendency-count row is the equal-per-fund-weighted mean of the count on the pre-outcome canonical panel. SE clustered on investor in parentheses; two-way clustering by investor and day gives $t = -16.4$ on the warm count price. Descriptive levels only — no difference column; the formal warm-minus-never-watched contrast is Table 7. Value purchases lose in both groups (-127 / -93 bps — the value trap is not warm buying-specific); the salient-returns member prices *positively* conditional on the other fourteen in the warm group (+91 bps, $t = +5.0$). The four-factor (CH-4) companion is Appendix Table A3. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 7: Which tendencies carry the warm–cold gap, Panel A (characteristic-adjusted, with and without state controls)

Panel A: the pooled gap, by horizon

	DGTW-adjusted warm–cold gap, bps					
	Without state controls			With state controls		
Member	5 days	20 days	60 days	5 days	20 days	60 days
Warm buying gap (no members)	−30.9***	−107.3***	−234.0***	−22.7***	−83.7***	−184.4***
	(6.7)	(11.2)	(16.2)	(6.6)	(11.1)	(16.1)
Warm buying gap, tendency-adjusted (members held)	+30.4*	+42.6	−88.3**	+26.1	+37.5	−89.5**
	(17.1)	(29.4)	(41.0)	(17.0)	(29.4)	(41.1)
Attention-type (chasing what's hot)						
Lottery-like (limit-up run)	−100.0***	−147.6***	−166.7***	−100.0***	−146.9***	−166.2***
	(18.7)	(30.1)	(37.2)	(18.7)	(30.0)	(37.2)
Recent limit-up close	−19.2	−93.9***	−83.0**	−16.4	−89.8***	−80.2**
	(14.2)	(24.4)	(34.1)	(14.2)	(24.5)	(34.1)
Thirty-day winner	+25.2*	−1.9	+26.5	+26.1*	−0.8	+27.1
	(13.8)	(24.9)	(36.7)	(13.8)	(24.9)	(36.7)
Abnormal volume	−12.4	+9.9	+40.8	−11.8	+10.3	+41.3
	(12.3)	(21.0)	(28.8)	(12.3)	(20.9)	(28.8)
High turnover	−24.5*	−86.9***	+15.9	−25.0*	−88.4***	+14.1
	(14.3)	(25.1)	(33.4)	(14.3)	(25.1)	(33.4)
High volatility	+16.6	+37.6	−48.8	+16.8	+38.1	−47.7
	(15.2)	(26.0)	(34.7)	(15.2)	(26.0)	(34.7)
Near the 52-week high	−14.8	−14.4	−71.7**	−14.5	−14.1	−75.9**
	(12.8)	(22.3)	(31.7)	(12.8)	(22.3)	(31.7)
Salient returns (Cosemans–Frehen)	+3.0	+48.1**	+48.5	+3.9	+49.6**	+51.0
	(13.5)	(22.8)	(31.1)	(13.4)	(22.9)	(31.1)
Contrarian-type (fishing for what's cheap)						
Thirty-day loser	−12.9	+1.0	−3.4	−11.5	+4.7	−3.0
	(13.2)	(24.6)	(38.3)	(13.2)	(24.7)	(38.4)
Low price	+62.7***	+30.3	+53.0	+62.3***	+28.7	+52.8
	(19.3)	(33.0)	(48.9)	(19.2)	(32.9)	(48.9)
Small cap	+17.0	+48.0	+117.1*	+15.2	+45.3	+113.5*
	(21.0)	(39.5)	(60.1)	(21.0)	(39.5)	(60.1)
Value (top book-to-market quintile)	−22.2	−62.4***	−49.0	−21.3	−61.5***	−49.0
	(13.8)	(22.7)	(32.4)	(13.8)	(22.6)	(32.5)
Familiarity-type						
Local company	+4.5	−8.4	−49.6	+4.0	−8.6	−49.5
	(17.7)	(31.3)	(43.9)	(17.6)	(31.3)	(43.9)
Familiar industry	−26.1**	−19.1	−0.7	−25.5**	−18.0	+0.2
	(12.8)	(21.1)	(29.0)	(12.8)	(21.1)	(29.0)
Income-type						
Dividend payer	+10.2	−9.2	+61.6**	+10.2	−9.4	+61.7**
	(11.7)	(19.2)	(26.8)	(11.7)	(19.2)	(26.8)
Additional gap per additional tendency	−13.0***	−28.0***	−33.2***	−12.5***	−28.3***	−36.0***
	(2.3)	(3.9)	(5.4)	(2.3)	(3.9)	(5.4)

Note: Which tendencies carry the warm–cold gap (script 674, canonical refit): the fifteen-member decomposition of the warm-minus-never-watched gap. Panel A (this page): the pooled gap at five, twenty, and sixty trading days, without and with buy-day price-state controls; Panels B and C (continued on the next page, script 679): the same decomposition on subsamples split by addition origin. *Specification.* Sample, battery, outcome convention, weights, fixed effects, and clustering are identical to Tables 5 and 6: the gap pool is the 303,059 warm and 209,719 never-watched purchases of the canonical sample (same-day excluded), DGTW-adjusted returns with complete windows, in basis points, equal-per-fund weights, investor and purchase-month fixed effects, SE clustered on investor. Member cells are the $\hat{\theta}^h_j$ from $\text{CAR}^h_{i,s,t} = \alpha_i + \gamma_{m(t)} + \delta^h W_{i,s,t} + \sum_{j=1}^{15} \beta^h_j F_{j,s,t^-} + \sum_{j=1}^{15} \theta^h_j W_{i,s,t} F_{j,s,t^-} + \varepsilon_{i,s,t}$, where W indicates a warm purchase — member j's *extra* gap when warm, conditional on the other members; member prevalences are in Table 6. The top row is the raw gap from the specification with no members; the tendency-adjusted row is $\hat{\delta}^h$ from the interacted specification (the gap for a purchase carrying no tendency); the bottom row is the warm buying–count interaction from the separate count specification. The with-state-controls columns add deciles of the stock's trailing sixty-day and 252-day price position at the purchase date as fixed effects: the sixty-day tendency-adjusted gap moves from −88 to −90 bps, so the fifteen members absorb the buy-day price state (recorded cell: members-sufficient). The raw sixty-day gap is −234 bps ($N = 401,814$; two-way investor-and-day $t = -11.3$), and the tendency basket explains 45% of it (count difference times count price). The four-factor companion is Appendix Table A4. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

panel of watched stock-days as a stratified conditional logit with (stock, day) strata, so that every quantity common to a stock on a date—its tendency profile, its momentum—is absorbed and the estimates come from within-stock-day variation in the price relative to the individual investor's own reference points. The hazard changes slope sharply as the price crosses the level at which the stock was added to the list—a kink of -2.56 ($z = -17.3$), a steep drop in the propensity to convert once the price rises above the add price (Table 8). A companion specification (Table 9) lets the add price compete with the highs and lows witnessed during the wait; distance above the witnessed trough carries its own positive slope. Because the strata absorb everything stock-date, this is the investor's own reference point driving *when* she converts, not which stocks are watched or how they trend—a conversion-timing mechanism consistent with reference dependence, of which anchoring on the add price is one case; I take no stand on loss aversion.

8 Magnitude, Welfare, and Design

Table 10 converts the per-trade estimates into money (script 225). The per-trade cost is the sixty-day within-investor-day action cost of buying a watched stock sitting below its addition price, relative to the same investor's watched losers left unbought that day: -246.4 basis points market-adjusted per purchase (-261.2 under the Chinese four-factor adjustment, -217.3 DGTW; 89,188 such purchases by 11,996 investors; Fama–MacBeth standard errors). Investors make 6.2 such purchases per investor-year, an annualized drag of $-1{,}527$ basis points market-adjusted. Applied to the mean stock-deployed capital of 177,111 RMB per investor (median 40,000; monthly account snapshots of 17,205 investors), the drag is $-27{,}040$ RMB per investor-year market-adjusted ($-28{,}658$ CH-4, $-23{,}843$ DGTW). This is a back-of-envelope whose every input — the per-trade cost, the purchase frequency, and the capital base — is recorded in the table, so a reader who prefers the median capital base, or a different benchmark, can rescale it directly.

The paper's question was whether the watchlist—the consideration stage of household investing—helps investors buy better. Measured against the only counterfactual the data make visible, the answer is two-sided: the stocks an investor watches carry a good signal, but converting them into purchases is where the money is lost, because conversion is reference-dependent and fires the trade at attention-heavy moments. Watching helps; the trap is in the buying. It is separable in data platforms already record, and it is a design margin they could act on.

Table 8: Reference-dependent conversion: the kink at the add price

Panel A: The kink at the add price (script 490)					
	(M1)	(M2)	(M3)	(M4)	(M5)
Return since add, below the add price (loss-side slope)	$+6.789^{***}$	$+2.975^{***}$	$+3.037^{***}$	$+3.083^{***}$	$+3.033^{***}$
	(0.201)	(0.113)	(0.113)	(0.115)	(0.113)
Level jump at the add price	-0.382^{***}	-0.115^{***}	-0.117^{***}	-0.119^{***}	-0.117^{***}
	(0.022)	(0.012)	(0.012)	(0.012)	(0.012)
Slope change above the add price (the kink)	-9.017^{***}	-2.492^{***}	-2.560^{***}	-2.609^{***}	-2.558^{***}
	(0.276)	(0.146)	(0.147)	(0.150)	(0.148)
Log days on watchlist	–	-0.894^{***}	-0.871^{***}	-0.871^{***}	-0.867^{***}
		(0.016)	(0.015)	(0.015)	(0.015)
Log watchlist size	–	–	-0.052^{***}	-0.052^{***}	-0.061^{***}
			(0.014)	(0.014)	(0.015)
Log prior buys	–	–	-0.171^{***}	-0.196^{***}	-0.188^{***}
			(0.017)	(0.016)	(0.017)
Log cash remaining	–	–	–	$+0.016^{***}$	$+0.017^{***}$
				(0.003)	(0.003)
Log news-reading minutes	–	–	–	–	$+0.057^{***}$
					(0.007)
Log chatbot queries	–	–	–	–	$+0.087^{***}$
					(0.007)
Buy events (cases)	250,907	250,907	250,907	250,907	250,907
(Stock × day) strata	123,599	123,599	123,599	123,599	123,599
Pseudo-R^2	0.353	0.411	0.413	0.413	0.413

Note: **The watched stock is bought at a reference point (Panel A).** Conditional-logit (stratified Cox) estimates of the daily first-buy hazard on the warm at-risk panel (investor × stock × day), script 490. Strata are (stock, day) with at least one buy, so every stock-date quantity is absorbed; coefficients are identified *within* stock-day from investor-specific variation in the price relative to that investor's own add price. The running variable is the return since the add price; the kink is the change in hazard slope as the price crosses the add price from below. Canonical conditions (252-day history floor, ST excluded, running variable winsorized at 1st/99th percentiles); cases matched 1:10 to risk-set controls (seed 20260609); covariates within-investor demeaned; day-clustered SE in parentheses. Read as *consistent with reference dependence*, not as evidence on loss aversion, and not as a channel for the warm–cold return gap. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 9: Reference-dependent conversion: the add price against the witnessed peak and trough

Panel B: Add price vs. witnessed peak and trough (script 495)				
	(U0)	(U1 peak)	(U2 trough)	(U3 both)
Return since add, below the add price (loss-side slope)	$+3.047^{***}$	$+3.059^{***}$	$+2.973^{***}$	$+2.968^{***}$
	(0.110)	(0.114)	(0.111)	(0.115)
Level jump at the add price	-0.113^{***}	-0.114^{***}	-0.045^{***}	-0.044^{***}
	(0.012)	(0.012)	(0.012)	(0.011)
Slope change above the add price (the kink)	-2.561^{***}	-2.568^{***}	-3.126^{***}	-3.108^{***}
	(0.144)	(0.146)	(0.154)	(0.154)
Distance below the witnessed peak	–	$+0.181$	–	$+0.200^{*}$
		(0.116)		(0.118)
New high since add today	–	$+0.035$	–	$+0.055^{**}$
		(0.024)		(0.024)
Reclaim near the witnessed peak	–	-0.109^{***}	–	-0.105^{***}
		(0.022)		(0.022)
Distance above the witnessed trough	–	–	$+0.623^{***}$	$+0.612^{***}$
			(0.058)	(0.056)
New low since add today	–	–	$+0.163^{***}$	$+0.178^{***}$
			(0.022)	(0.023)
Recovery near the witnessed trough	–	–	-0.074^{***}	-0.063^{***}
			(0.021)	(0.021)
Buy events (cases)	250,907	250,907	250,907	250,907
Pseudo-R^2	0.413	0.413	0.414	0.414

Note: **The add-price reference against the witnessed extremes (Panel B).** The add-price kink of Table 8 (script 495) competes with distance from the running peak and trough the investor witnessed over the watching window; distance above the witnessed trough carries its own positive slope. Same panel, strata, matching, and conditions as Table 8; day-clustered SE in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: What forced deferral costs in money (script 225)

	Market-adjusted	CH-4	DGTW
Cost per such purchase (basis points)	-246.4^{***}	-261.2^{***}	-217.3^{***}
	(33.2)	(38.0)	(38.2)
Such purchases per investor-year: 6.2.			
Annualized drag (basis points per year)	$-1{,}527^{***}$	$-1{,}618^{***}$	$-1{,}346^{***}$
	(206)	(235)	(237)
Mean stock-deployed capital: 177,111 RMB per investor.			
Annualized drag (RMB per investor-year)	$-27{,}040^{***}$	$-28{,}658^{***}$	$-23{,}843^{***}$
	(3,641)	(4,171)	(4,195)

Note: What forced deferral costs in money (script 225). The per-trade cost is the sixty-day within-investor-day action cost of buying a watched stock sitting below its addition price, relative to the same investor's watched losers left unbought that day (89,188 such purchases by 11,996 investors; Fama–MacBeth standard errors in parentheses). The annualization multiplies by the inferred 6.2 such purchases per investor-year and applies the resulting drag to the mean stock-deployed capital of 177,111 RMB per investor (median 40,000; monthly account snapshots of 17,205 investors) — a back-of-envelope whose every input is recorded so the reader can rescale it. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A The Tendency Battery: Mapping to the Source Literature

The battery adapts the twenty-two trading tendencies of Han et al. (2025) into fifteen estimable members. Table A1 gives the full correspondence: twelve map one-to-one, three are carried by an adapted neighbor, seven are excluded (for the first-purchase event grain, for data availability, or because the Chinese price-limit institution removes the object), and six are added beyond their roster; three of the eighteen frozen definitions are not estimated, leaving fifteen.

B Chinese Four-Factor Companions

The decomposition of Section 6 is reported on characteristic-adjusted (DGTW) returns in the main text; the Chinese four-factor (CH-4) companions are collected here.

Table A1: Mapping the twenty-two Han–He–Weagley tendencies into the eighteen frozen battery definitions

Han et al. (2025) tendency	This paper	Relation
Mapped one-to-one (twelve)		
Abnormal Trading Volume	Abnormal volume	Adapted: above-median cut → top quintile; volume baseline 252 → 30 trading days.
Momentum	Thirty-day winner	Adapted: 20-day window → 30 trading days; top decile → top quintile.
Contrarian	Thirty-day loser	Adapted: 20-day window → 30 trading days; bottom decile → bottom quintile.
Volatility	High volatility	Adapted substantively: portfolio-matched risk habitat → unconditional top volatility quintile.
Lottery-Like Stocks	Lottery-like	Adapted to the limit institution: volatility–skewness–price sort → consecutive-limit-up run. Dormant in Finland; live here.
Low Price	Low price	Matched: bottom decile of the prior day's nominal price. Dormant in Finland; live here.
Small Cap	Small cap	Adapted: bottom decile → bottom quintile; free-float capitalization. Dormant in Finland; live here.
High Dividend	Dividend payer	Adapted: top-decile dividend level → extensive-margin payer flag.
Local Company	Local company	Adapted: shared municipality of residence → shared native province.
Top Industry Experience	Familiar industry	Adapted: personal-success filter and one-year window dropped.
Value	Value	Adapted: top decile of book-to-market → top quintile, from the annual June characteristic formations of the benchmark pipeline. Dormant in Finland; live here. Admitted at the design revision.
Salient Returns	Salient returns	Matched with documented adaptations: their flag is the top decile of the Cosemans and Frehen 2021 measure by calendar month; here the same measure on a rolling twenty-one-day window at the anchor day, cut at the top quintile. Admitted at the design revision.
Carried by an adapted neighbor (three)		
Return Sorting Top	Thirty-day winner	Yesterday's-return top decile, carried by the past-winner member.
Return Sorting Bottom	Thirty-day loser	Symmetric: carried by the past-loser member.
Chase Dividend	Dividend payer	Announcement-to-record-date timing not implemented; payer status carries the dividend margin.
Excluded (seven)		
Holding Period	—	By grain: buying a stock already held six months cannot occur at a first purchase.
Easy Name	—	No analog among short Chinese security-name abbreviations.
Integer Price	—	Requires intraday transaction prices; the battery is daily.
Naive Buying Diversification	—	Requires comparing same-day purchase values across stocks; not implemented at this grain.
Peer Influence Industry	—	Requires fine residence geography and full-population purchase shares; the platform records native province only.
Peer Influence Stock	—	Same requirement.
Max Return	—	The daily limit censors the maximum-return sort toward limit-up days; terrain carried by the limit-based members.
Added beyond the twenty-two (six)		
—	Recent limit-up close	The Chinese price-limit institution (Seasholes and Wu 2007; Chen et al. 2019).
—	Near the 52-week high	The 52-week-high anchor (George and Hwang 2004; Huddart et al. 2009).
—	Near the 52-week low	Mirror of the above (George and Hwang 2004; Grinblatt and Keloharju 2001b); fails the availability screen on both the full and the filtered sample.
—	New listing	Retail demand for recent listings (Ritter 1991; Dorn 2009); excluded from the estimated battery at the design revision (the history filter removes young listings).
—	High turnover	Turnover level as an attention proxy (Barber and Odean 2008; Da et al. 2011).
—	Bought before (repurchase)	From the broader literature (Strahilevitz et al. 2011); not constructible at this grain.

Note: Mapping the twenty-two trading tendencies of Han et al. (2025) into the eighteen frozen battery definitions: twelve map one-to-one onto members (two of them, Value and Salient Returns, admitted at the design revision after an availability re-screen on the filtered estimation sample), three are carried by an adapted neighbor, and seven are excluded—for the first-purchase event grain, for data availability, or because the Chinese institution removes the object—while six members are added beyond the twenty-two. Original references for the excluded tendencies are as consolidated in Han et al. 2025. "Dormant in Finland" marks the four tendencies Han, He, and Weagley find no more prevalent in real than in simulated trades and exclude from their main analysis. Three definitions are not estimated — near the 52-week low (availability floor, both screens), bought before (impossible at this grain), and new listing (removed by the 252-day history filter) — leaving the fifteen canonical members of Table 5.

Table A2: The fifteen-member panel under the four-factor adjustment (appendix companion)

Member (share flagged)	Four-factor (CH-4), bps		
	5 days	20 days	60 days
Attention-type (chasing what's hot)			
Lottery-like (limit-up run) (21.5%)	−139.6***	−266.5***	−523.4***
	(8.1)	(12.7)	(16.6)
Recent limit-up close (45.5%)	−109.4***	−204.2***	−341.7***
	(6.4)	(10.5)	(14.9)
Thirty-day winner (59.2%)	−16.1***	−62.2***	−134.1***
	(6.1)	(10.7)	(16.0)
Abnormal volume (55.7%)	−53.0***	−78.3***	−125.0***
	(5.6)	(9.1)	(12.7)
High turnover (40.1%)	+4.4	−43.7***	−164.1***
	(6.2)	(10.2)	(14.0)
High volatility (40.6%)	−37.5***	−116.9***	−174.0***
	(6.6)	(10.4)	(15.1)
Near the 52-week high (54.5%)	−21.9***	−36.0***	−4.6
	(5.7)	(10.1)	(14.1)
Salient returns (Cosemans–Frehen) (54.2%)	−33.5***	+36.9***	+182.8***
	(5.9)	(9.7)	(13.8)
Contrarian-type (fishing for what's cheap)			
Thirty-day loser (8.9%)	−15.7**	−15.5	−76.5***
	(6.1)	(11.4)	(18.0)
Low price (9.5%)	+92.2***	+269.3***	+185.3***
	(8.5)	(14.4)	(20.2)
Small cap (5.2%)	−41.9***	−18.2	+153.1***
	(9.1)	(15.9)	(24.1)
Value (top book-to-market quintile) (19.4%)	−47.3***	−168.1***	−328.9***
	(6.0)	(9.5)	(13.5)
Familiarity-type			
Local company (8.1%)	+11.0	+41.6***	+71.7***
	(7.9)	(13.8)	(19.4)
Familiar industry (84.1%)	−13.6**	−15.8*	+103.2***
	(5.9)	(9.5)	(13.4)
Income-type			
Dividend payer (63.3%)	−17.7***	−12.8	+53.2***
	(5.3)	(8.4)	(11.5)
Price of one additional tendency (separate specification)	−45.2***	−84.0***	−131.0***
	(1.1)	(1.9)	(2.6)

Note: The fifteen-member panel under the four-factor adjustment (script 672, canonical refit; appendix companion to Table 5): identical sample, battery, weights, fixed effects, and clustering, with the outcome replaced by Liu–Stambaugh–Yuan four-factor (CH-4) cumulative abnormal returns over five, twenty, and sixty trading days, in basis points. Windows are complete per adjustment; the four-factor residuals extend past the panel boundary, so no calendar tail is excluded (sixty-day coverage 100.0%, $N = 705,972$). Cells: coefficient with SE clustered on investor in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A3: The warm-versus-never-watched anatomy under the four-factor adjustment (appendix companion)

	Four-factor (CH-4), 60 days (bps)	
Member (share flagged)	Warm	Never watched
Attention-type (chasing what's hot)		
Lottery-like (limit-up run) (21.5%)	-529.2^{***}	-509.3^{***}
	(22.3)	(32.7)
Recent limit-up close (45.5%)	-353.4^{***}	-388.6^{***}
	(20.3)	(28.6)
Thirty-day winner (59.2%)	-132.7^{***}	-149.3^{***}
	(22.3)	(28.7)
Abnormal volume (55.7%)	-98.2^{***}	-98.5^{***}
	(17.8)	(22.4)
High turnover (40.1%)	-181.4^{***}	-165.4^{***}
	(19.1)	(26.8)
High volatility (40.6%)	-166.8^{***}	-125.7^{***}
	(20.5)	(27.6)
Near the 52-week high (54.5%)	-9.5	$+85.5^{***}$
	(20.2)	(24.6)
Salient returns (Cosemans–Frehen) (54.2%)	$+222.2^{***}$	$+217.3^{***}$
	(18.9)	(25.2)
Contrarian-type (fishing for what's cheap)		
Thirty-day loser (8.9%)	-71.7^{***}	-56.1^{*}
	(25.7)	(30.1)
Low price (9.5%)	$+175.3^{***}$	$+141.8^{***}$
	(27.8)	(39.5)
Small cap (5.2%)	$+219.3^{***}$	$+140.4^{***}$
	(39.2)	(44.8)
Value (top book-to-market quintile) (19.4%)	-357.7^{***}	-240.3^{***}
	(18.8)	(24.9)
Familiarity-type		
Local company (8.1%)	$+85.1^{***}$	$+61.6^{*}$
	(27.4)	(35.2)
Familiar industry (84.1%)	$+89.4^{***}$	$+153.5^{***}$
	(20.2)	(25.1)
Income-type		
Dividend payer (63.3%)	$+54.0^{***}$	$+118.5^{***}$
	(16.6)	(21.8)
Price of one additional tendency (separate specification)	-129.3^{***}	-107.6^{***}
	(3.6)	(4.8)

Note: The warm-versus-never-watched anatomy under the four-factor adjustment (script 673, canonical refit; appendix companion to Table 6): identical sample, poles, battery, weights, fixed effects, and clustering, with the outcome replaced by Liu–Stambaugh–Yuan four-factor (CH-4) cumulative abnormal returns over sixty trading days, in basis points. Windows are complete per adjustment; the four-factor residuals extend past the panel boundary, so no calendar tail is excluded (per-group sixty-day coverage 100.0% / 100.0%; estimation samples $N = 300,097/207,459$). Cells: coefficient with SE clustered on investor in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A4: The gap decomposition under the four-factor adjustment (appendix companion)

	Four-factor (CH-4) warm–cold gap, bps					
	Without state controls			With state controls		
Member	5 days	20 days	60 days	5 days	20 days	60 days
Warm buying gap (no members)	−31.6***	−96.4***	−186.9***	−24.5***	−82.4***	−170.3***
	(6.8)	(11.5)	(16.4)	(6.7)	(11.5)	(16.4)
Warm buying gap, tendency-adjusted (members held)	+52.5***	+125.0***	−6.2	+47.5***	+117.4***	−25.3
	(17.6)	(30.0)	(42.4)	(17.6)	(30.0)	(42.4)
Attention-type (chasing what's hot)						
Lottery-like (limit-up run)	−122.9***	−110.3***	−96.7**	−123.2***	−111.7***	−100.6***
	(18.9)	(30.2)	(38.6)	(18.9)	(30.2)	(38.6)
Recent limit-up close	−20.5	−31.0	−44.3	−17.8	−28.4	−43.4
	(14.7)	(24.9)	(34.4)	(14.7)	(24.9)	(34.4)
Thirty-day winner	+33.6**	+15.6	+30.7	+33.8**	+15.1	+31.1
	(14.2)	(25.6)	(37.3)	(14.2)	(25.6)	(37.3)
Abnormal volume	+1.8	+12.3	+21.7	+2.5	+13.2	+24.3
	(12.6)	(21.1)	(29.1)	(12.6)	(21.1)	(29.1)
High turnover	−39.8***	−82.3***	−24.1	−39.8***	−81.7***	−20.0
	(14.5)	(24.8)	(33.5)	(14.5)	(24.7)	(33.5)
High volatility	+21.0	+54.8**	−32.6	+20.9	+55.4**	−31.1
	(15.5)	(25.8)	(34.7)	(15.5)	(25.8)	(34.7)
Near the 52-week high	−22.4*	−86.8***	−105.7***	−21.3	−81.7***	−90.8***
	(13.2)	(22.9)	(31.8)	(13.2)	(23.0)	(31.9)
Salient returns (Cosemans–Frehen)	−13.7	−3.8	+58.0*	−12.7	−3.6	+58.4*
	(13.8)	(23.2)	(32.1)	(13.7)	(23.2)	(32.1)
Contrarian-type (fishing for what's cheap)						
Thirty-day loser	−13.1	−35.4	+8.6	−11.4	−31.9	+11.7
	(13.6)	(25.4)	(40.4)	(13.6)	(25.4)	(40.4)
Low price	+73.2***	+101.3***	+83.4*	+73.0***	+101.7***	+86.4*
	(20.1)	(33.3)	(47.0)	(20.1)	(33.2)	(47.0)
Small cap	+18.2	+57.9	+66.9	+16.5	+58.0	+68.5
	(21.6)	(38.2)	(58.0)	(21.5)	(38.1)	(58.0)
Value (top book-to-market quintile)	−32.0**	−102.4***	−129.1***	−31.3**	−102.1***	−128.1***
	(14.1)	(23.1)	(32.1)	(14.1)	(23.1)	(32.1)
Familiarity-type						
Local company	−2.0	+0.5	−47.2	−2.0	−0.2	−46.2
	(18.2)	(31.6)	(43.7)	(18.2)	(31.6)	(43.6)
Familiar industry	−23.9*	−57.1***	−20.7	−23.2*	−56.3***	−19.8
	(13.1)	(21.6)	(30.3)	(13.1)	(21.6)	(30.2)
Income-type						
Dividend payer	−11.6	−77.6***	−24.4	−11.3	−76.9***	−22.0
	(12.0)	(20.1)	(27.5)	(12.0)	(20.1)	(27.5)
Additional gap per additional tendency	−16.7***	−30.5***	−31.7***	−16.1***	−29.8***	−29.7***
	(2.4)	(4.0)	(5.5)	(2.4)	(4.1)	(5.5)

Note: The fifteen-member gap decomposition under the four-factor adjustment (script 674, canonical refit; appendix companion to Table 7): identical pool, battery, specifications, weights, fixed effects, and clustering, with the outcome replaced by Liu–Stambaugh–Yuan four-factor (CH-4) cumulative abnormal returns over five, twenty, and sixty trading days, in basis points. Windows are complete per adjustment; the four-factor residuals extend past the panel boundary, so no calendar tail is excluded (sixty-day estimation sample $N = 505{,}726$; 507,556 with state controls). Cells: coefficient with SE clustered on investor in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

References

Akepanidtaworn, K., R. Di Mascio, A. Imas, and L. Schmidt, 2023, Selling fast and buying slow, *Journal of Finance* 78, 3055–3098.

Barber, B., and T. Odean, 2008, All that glitters, *Review of Financial Studies* 21, 785–818.

Bordalo, P., N. Gennaioli, and A. Shleifer, 2020, Memory, attention, and choice, *Quarterly Journal of Economics* 135, 1399–1442.

Chen, T., Z. Gao, J. He, W. Jiang, and W. Xiong, 2019, Daily price limits and destructive market behavior, *Journal of Econometrics* 208, 249–264.

Cosemans, M., and R. Frehen, 2021, Salience theory and stock prices: Empirical evidence, *Journal of Financial Economics* 140, 460–483.

Da, Z., J. Engelberg, and P. Gao, 2011, In search of attention, *Journal of Finance* 66, 1461–1499.

Daniel, K., M. Grinblatt, S. Titman, and R. Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* 52, 1035–1058.

Dorn, D., 2009, Does sentiment drive the retail demand for IPOs?, *Journal of Financial and Quantitative Analysis* 44, 85–108.

Enke, B., 2026, The cognitive turn in behavioral economics, working paper.

Gargano, A., and A. Rossi, 2018, Does it pay to pay attention?, *Review of Financial Studies* 31, 4595–4649.

George, T., and C.-Y. Hwang, 2004, The 52-week high and momentum investing, *Journal of Finance* 59, 2145–2176.

Grinblatt, M., and M. Keloharju, 2001b, What makes investors trade?, *Journal of Finance* 56, 589–616.

Han, H.-S., X. He, and D. Weagley, 2025, Dissecting retail investor trading tendencies, working paper.

Hartzmark, S., 2015, The worst, the best, ignoring all the rest: The rank effect and trading behavior, *Review of Financial Studies* 28, 1024–1059.

Hartzmark, S., S. Hirshman, and A. Imas, 2021, Ownership, learning, and beliefs, *Quarterly Journal of Economics* 136, 1665–1717.

Huddart, S., M. Lang, and M. Yetman, 2009, Volume and price patterns around a stock's 52-week highs and lows: Theory and evidence, *Management Science* 55, 16–31.

Imas, A., L. Jin, and X. Lü, 2026, Limited consideration sets in financial decision-making, working paper.

Kim, K. A., and S. G. Rhee, 1997, Price limit performance: Evidence from the Tokyo Stock Exchange, *Journal of Finance* 52, 885–901.

Levy, 2025, What's in investors' information sets?, working paper.

Liu, J., R. Stambaugh, and Y. Yuan, 2019, Size and value in China, *Journal of Financial Economics* 134, 48–69.

Odean, T., 1999, Do investors trade too much?, *American Economic Review* 89, 1279–1298.

Ritter, J., 1991, The long-run performance of initial public offerings, *Journal of Finance* 46, 3–27.

Seasholes, M., and G. Wu, 2007, Predictable behavior, profits, and attention, *Journal of Empirical Finance* 14, 590–610.

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

Wachter, J., and M. Kahana, 2024, A retrieved-context theory of financial decisions, *Quarterly Journal of Economics* 139, 1095–1147.