Chapter 2

Stock Ownership and Inflation Expectations:

Evidence from the 2021–2025 Inflation Surge

Giovanni Borraccia (Geneva Graduate Institute)

Abstract

This paper investigates whether stock ownership shapes how households form inflation expectations, and why the gap between stockholders and non-stockholders, stable near 0.76 percentage points for two decades, widened sharply during the 2021–2025 inflation surge. Using the University of Michigan Survey of Consumers and a difference-in-differences design, I find that stockholders' one-year-ahead expectations rose 1.70 percentage points less than non-stockholders' across the surge, and their absolute forecast errors fell 1.39 points more. The divergence is concentrated in the upper tail of the belief distribution, where the effect reaches 1.63 points at the top of the distribution and fades toward the median, and is absent during shorter inflationary episodes. To explain it, I develop a two-signal rational-inattention model in which stock ownership lowers the cost of processing an unbiased financial signal. Structural estimation reproduces the group means with a far smaller attention cost for stockholders, and traces the widening to an amplification of experienced-price bias among non-stockholders anchored to salient food and energy prices, which rises from 1.23 to 2.13 percentage points as the surge persists, while stockholders reallocate attention toward aggregate financial signals.

Keywords: inflation expectations, stock ownership, rational inattention, perceptual bias, difference-in-differences, household heterogeneity.

JEL: E31; D84; G51; D83.

2.1 Introduction

Between 2021 and 2025, the global economy experienced its most sustained wave of inflation in four decades (Bernanke and Blanchard, 2023; Ball et al., 2022). Household inflation expectations grew substantially in parallel, but did not do so uniformly. The dispersion of one-year ahead inflation expectations displayed a sharp widening: a household's position in that distribution turned out to depend on something apparently unrelated to what is usually conceived as underpinning forecasting dynamics, stock ownership. Based on American data from the University of Michigan Survey of Consumers, prior to 2021 the gap between stockholders' and non-stockholders' expectations was consistently around 1 percentage point for almost two decades. As the surge in inflation materialized and persisted, this gap more than tripled and non-stockholders accounted for nearly all of the extreme beliefs at the top of the distribution.

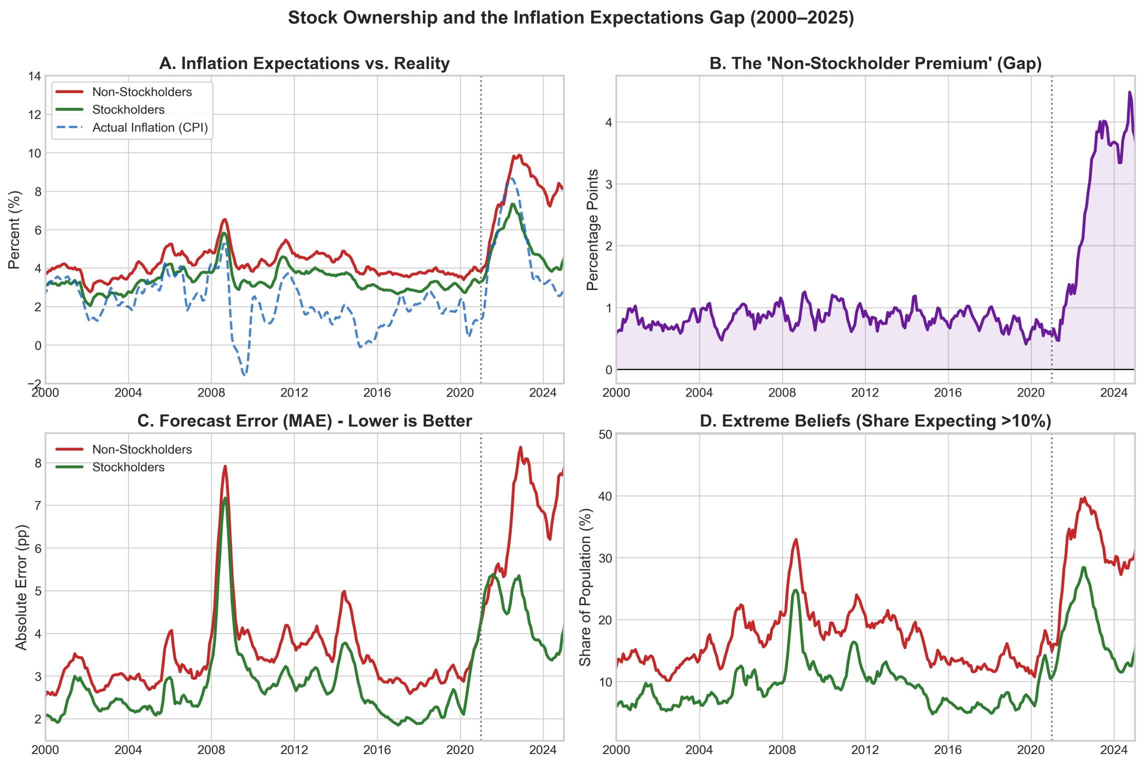

Figure 1: Stock ownership and the inflation-expectations gap, 2000–2025. Panel A: mean one-year-ahead expectations of stockholders and non-stockholders against realized CPI inflation. Panel B: the stockholder–non-stockholder expectations gap. Panel C: mean absolute forecast errors by group. Panel D: share of each group holding extreme beliefs (> 10%). Data: Michigan Survey of Consumers; the vertical line marks January 2021.

The obvious explanation to this would be selection: stockholders are richer and better educated. In this sample 55 percent are college graduates against 22 percent of non-stockholders and 54 percent are high income against 14 percent. These households may simply forecast better for reasons that are unrelated to the stock market. However, selection would have a sharp implication which the data seem to reject. If the gap purely reflected household traits, it should be stable over time, instead it was small and flat for two decades and then tripled after 2021. Importantly, this widening did not materialize in 2008, an earlier episode of comparable inflation volatility. The timing dimension is the puzzle the paper seeks to unpack: is the gap causal or compositional, what mechanism generates it, where in the distribution does it operate, and why did it appear in 2021 but not in 2008?

I examine these questions using the University of Michigan Survey of Consumers, a monthly repeated cross-section which has been running since the 1970s and reporting households inflation expectations and whether respondents own stocks. The estimation sample contains 166,125 responses covering both the pre- and post-2021 periods. The empirical design I adopt is a difference-in-differences specification which compares stockholders and non-stockholders before and after the onset of the surge. Common movements in expectations are absorbed by year fixed effects, so identification is grounded in the differential change between the two groups. I add demographic and socioeconomic controls in layers since the survey composition shifts slightly over time. In the preferred specification, I let each demographic group have its own post-2021 shift, so the headline estimates are net of any tendency for women, younger, less-educated or lower-income households to respond differently to the shock. Stockholders' one-year-ahead expectations rose 1.7 percentage points less than those of similar non-stockholders across the surge. The estimate is stable as controls are added: in its rawest form the specification yields a coefficient of 2.48 falling monotonically to 1.94 with socioeconomic controls and 1.7 in the most conservative one. This divergence appears to be grounded in accuracy rather than optimism, stockholders' absolute forecast error fell 1.39 percentage points more than non-stockholders' and their upward forecast bias 1.47 points more both compared to the pre-2021 baseline period. The analysis identifies information access as mechanism rather than sophistication mainly through three specific patterns. First, the effect of stock ownership turns out to have the largest size precisely where sophistication would predict it to be smallest: it is 1.84 times larger for low-income households and 1.79 times larger for women. Interactions with college education and age are insignificant and I interpret this as households having relatively fewer alternative information channels gain the most from owning stocks. The nature of this advantage is also fully specific to consumer prices since stockholders' accuracy edge does not extend to their home price expectations (difference-in-differences interaction on forecast error is small and insignificant). This is consistent with a signal that provides valuable information about CPI inflation rather than general forecasting skill. Finally, the effect is concentrated in the upper tail of the distribution with the gap being six times larger at the 75th percentile of beliefs than at the 25th (-2.77 against -0.45) and winsorizing the extreme responses removes most of it. Stock ownership does not shift beliefs down uniformly; it prevents the extreme high expectations that the surge produced among those inattentive to the financial signal. The same pattern recurs within stockholders, where those holding larger portfolios revised both expectations and errors down more steeply as the inflation surge persisted. This dynamic displays a gradient which is absent in a pre-2021 placebo; the divergence scales with stakes. What appears to separate 2021 from 2008 is duration rather than magnitude: only the more recent episode was sustained enough for the perceptual bias anchored to salient food and energy prices to accumulate, widening the gap as attentive households (stockholders) raised their weight on less biased financial signals. On the other hand, non-stockholders, anchored to their grocery basket, absorbed a steadily rising bias. I develop a two-signal rational inattention model to organize these findings. Households forecast inflation by combining a costless but positively biased signal from their own experienced prices together with a costly, unbiased signal from macro-financial news and they choose how much attention to pay to the latter. Building on Sims (2003) and Mankiw and Reis (2002), I assume an information-bundling discount that lowers stockholders' marginal cost of attention, given that monitoring a portfolio provides exposure to macroeconomic information. As inflation persists, the error in the experienced price signal grows raising the optimal attention of those already paying it and leaving the gap to widen. A moment-matching estimator disciplined by the model's first order condition recovers a stockholder attention weight rising from 0.39 to 0.50 across the episode and a perceptual bias rising from 1.23 to 2.13 percentage points, while non-stockholders' attention weight stays at its corner of zero. The two margins of the reduced form (rising stockholder attention and rising non-stockholder bias) appear directly in the structural estimates. This paper bridges two branches of the literature: one on inflation expectations formation and its connection with individual experience and salient prices (Malmendier and Nagel, 2016; D'Acunto et al., 2021a), and one applying rational inattention to expectations formation (Sims, 2003; Coibion and Gorodnichenko, 2015). The remainder of the paper is structured as follows. Section 2.2 reviews the related work; Section 2.3 develops the model; Sections 2.4 and 2.5 describe the data and present the main estimates; Sections 2.6 through 2.8 turn to mechanism, robustness, and structural estimation; Section 2.9 concludes.

2.2 Related Literature

Households disagree about future inflation to an extent that full-information rational expectations cannot accommodate, and that disagreement is neither random noise nor a fixed feature of the cross-section. Instead, it is something deeply rooted into their socio-economic and demographic characteristics (Coibion and Gorodnichenko, 2015, 2012; Mankiw et al., 2003). The dispersion of household forecasts is large, persistent and systematically related to observable characteristics, and it widens along with higher volatility in realized inflation (Mankiw et al., 2003; Patton and Timmermann, 2010). A large body of survey-based work has progressively emerged treating these subjective expectations as core economic objects: they can be measured (Manski, 2004; Engelberg et al., 2009), they differ sharply across households and they shape economic decisions, chiefly consumption and saving behavior (Coibion et al., 2018; Weber et al., 2022; Bachmann et al., 2015). Households appear to act on these beliefs: experimental and survey evidence shows that the inflation people expect feeds into their spending and portfolio choices (Armantier et al., 2015; Roth and Wohlfart, 2020). The divergence identified by several papers is deep enough that even when households and professionals are handed an identical information set, their beliefs about the direction and size of macroeconomic shocks remains widely dispersed (Andre et al., 2022). Observing the degree to which disagreement is systematic and responsive to conditions calls for a structural investigation of where it originates, what places a given household high or low in the distribution to begin with, and what moves it as inflation rises.

A first explanation for the positioning in this distribution identifies the source of disagreement in what households personally experience, rather than in what official statistics report (Malmendier and Nagel, 2016; D'Acunto et al., 2021a; Cavallo et al., 2017). Individuals overweigh the inflation they have lived through, so that lifetime experience predicts expectations and younger households, having shorter memories, update more aggressively (Malmendier and Nagel, 2016). At higher frequency, they anchor on the prices they encounter most often, the price changes in their grocery basket, weighted by how frequently goods are purchased rather than by expenditure share, and with increases having greater impact on memories than decreases (D'Acunto et al., 2021a). Reliance on personally experienced prices appears to persist even when respondents are shown official aggregate figures, pointing to the presence of genuine information frictions (Cavallo et al., 2017). Because the anchor is what a household sees rather than what it earns, the resulting bias is largest where price exposure is most salient: women, who historically perform most household grocery shopping, hold systematically higher expectations (D'Acunto et al., 2021b), households extrapolate from their own recent experience when forming aggregate forecasts (Kuchler and Zafar, 2019; Bordalo et al., 2020), and the belief gap across socioeconomic groups itself widens and narrows over the cycle (Das et al., 2020). Experienced prices, however, are only one of the inputs a household can draw on; another key one is the aggregate signal carried in macroeconomic news. Which of the two a household relies on more is itself a choice about how to allocate scarce attention, rather than a fixed trait.

Why households differ in how heavily they lean on experienced prices rather than aggregate signals is, in turn, within the analytical scope of the rational-inattention tradition. Here, attention is treated as a costly and endogenous choice that responds to the information environment (Sims, 2003; Mankiw and Reis, 2002). Agents with finite processing capacity optimally direct attention toward whatever is most variable or most payoff-relevant, so the same household attends more to inflation exactly when inflation matters more (Sims, 2003; Maćkowiak and Wiederholt, 2009). Sticky-information and inattentive-consumer models make the point dynamically: households update infrequently, and those facing lower information costs update more often and forecast more accurately (Mankiw and Reis, 2002; Reis, 2006). Households moreover absorb macroeconomic news indirectly, through media reports of professional forecasters' views, so the diffusion of information is itself a margin of attention (Carroll, 2003; Lamla and Lein, 2014); professional forecasters themselves display inattention to public information, updating in ways consistent with sticky-information dynamics (Andrade and Le Bihan, 2013); and when attention is scarce, agents reason with sparse, simplified models that retain only first-order variables (Gabaix, 2014; Maćkowiak et al., 2023). If attention is chosen in this way, it tracks the information environment a household happens to inhabit and that environment is not common to all. Some households move through daily life being in closer contact with macroeconomic information than others, which raises the question of what observable traits marks them out.

A smaller and more recent literature focuses precisely on identifying such traits, showing that equity-market participation in particular places households in a richer macroeconomic information environment (Ahn and Xie, 2024; Peress, 2014; Jeong et al., 2025). Most directly, Ahn and Xie (2024) document on the Michigan survey that stockholders forecast inflation more accurately, disagree less, and revise more responsively to news than non-stockholders, evidence that participation lowers the effective cost of attending to macroeconomic conditions. On the supply side, Peress (2014) show that the financial media causally diffuse macro-relevant information to market participants: when newspapers strike, trading volume and the dispersion of returns both fall. Jeong et al. (2025) formalize the link in a Gabaix-style attention model in which belief revisions to policy news concentrate among attentive, high-financial-stakes households, a first-order condition structurally isomorphic to the one developed below. This reading sits within a longer household-finance literature on why so few households hold equities and what separates participants from non-participants (Haliassos and Bertaut, 1995; Mankiw and Zeldes, 1991; Vissing-Jorgensen, 2003; Guiso and Sodini, 2013), with financial literacy playing a first-order role in determining who holds stocks (van Rooij et al., 2011; Lusardi and Mitchell, 2014) and the resulting investment quality differing substantially across the population (Calvet et al., 2007). Alongside theories in which attention is allocated toward higher-stakes positions (Kacperczyk et al., 2016), direct measures confirm that investors attend to their portfolios selectively (Da et al., 2011; Sicherman et al., 2016). Read together, these strands describe a household whose place in the distribution of inflation beliefs is set at once by the prices it encounters, the attention it can spare, and the macroeconomic information its financial life brings within reach.

What these strands of literature have in common is that they read belief differences as standing feature. A prolonged inflation episode unsettles that reading, and poses a question none of them is built to answer: not how far apart two groups of households stand on average, but whether and why that distance changes when a shock endures. Placing an observable financial trait at the center, the paper takes the duration of an inflation episode as its source of variation and asks how the structure of disagreement responds as the episode lengthens. The contribution here is to show that a belief cleavage long treated as a fixed feature of the cross-section is in fact state-dependent, responding to how long inflation lasts rather than only how high it rises, and to trace that response to a single mechanism in which attention and experienced-price bias move together. In doing so I link the household-finance question of who holds stocks to the macroeconomic question of how expectations are formed, areas of the literature that have largely developed apart. The remainder of the paper develops the two-signal attention model in Section 2.3, describes the Michigan survey data in Section 2.4, and then turns to the reduced-form, distributional, and structural evidence.

2.3 A Model of Attention, Bundling, and Bias Amplification

2.3.1 Setup: signals and two household groups

A household forms a point forecast of next-period headline inflation π from two noisy signals. A *financial* signal summarises macro-financial news; it is unbiased, but attending to it is costly as it requires attention. An *experienced-price* signal summarises the prices the household experiences in daily life; it is free, since a household is exposed to it naturally, but it runs systematically above headline inflation because salient, frequently purchased items (food and energy above all) are over-weighted relative to their expenditure share (D'Acunto et al., 2021a; Malmendier and Nagel, 2016; Cavallo et al., 2017):

$$s_i^F = \pi + \varepsilon_i^F, \quad \varepsilon_i^F \sim \mathcal{N}(0, \sigma_F^2); \qquad s_i^E = \pi + b_\tau + \varepsilon_i^E, \quad \varepsilon_i^E \sim \mathcal{N}(0, \sigma_E^2), \tag{2.3.1}$$

where $b_\tau > 0$ is the experienced-price bias at horizon τ. The household combines the two signals with an attention weight $m_i \in [0, 1]$ on the financial signal,

$$E_i \;=\; m_i\, s_i^F + (1 - m_i)\, s_i^E \;=\; \pi + (1 - m_i)\, b_\tau + \left[m_i \varepsilon_i^F + (1 - m_i)\varepsilon_i^E\right], \tag{2.3.2}$$

so that a household attending more to financial news carries less of the experienced-price bias into its forecast.

Households belong to one of two groups, stockholders and non-stockholders, $g \in \{S, N\}$. Because the model refers to group averages, I define the group attention cost as the within-group mean $\kappa_g \equiv \mathbb{E}[\kappa_i \mid g]$; any within-group distribution of individual κ_i is consistent with what follows. Stock ownership enters through a single channel: monitoring a portfolio bundles exposure to macro-financial information and so lowers the marginal cost of attending to it. Formally, stockholders receive a *bundling discount* $\delta > 0$ on their attention cost.

$$\kappa_S \;=\; \kappa_N - \delta, \qquad \delta > 0, \tag{2.3.3}$$

This discount is grounded in the rational-inattention principle that attention costs are endogenous to the information environment (Sims, 2003; Mankiw and Reis, 2002) and in direct evidence that stockholders attend more closely to macroeconomic conditions (Ahn and Xie, 2024; Peress, 2014). Stockholder status is taken as given throughout: the model describes how each group forms expectations, not who comes to hold stocks.

2.3.2 Optimal attention and the bundling discount δ

Each household chooses its attention weight to minimise expected squared forecast error plus the cost of attention,

$$\min_m \,(1 - m)^2 \Sigma_{E,\tau}^2 + m^2 \sigma_F^2 + \tfrac{\kappa_g}{2} m^2, \tag{2.3.4}$$

where the effective noise of the experienced-price signal combines its dispersion and its bias, $\Sigma_{E,\tau}^2 = \sigma_E^2 + b_\tau^2$: for forecasting purposes a more biased signal is a noisier one. The first-order condition gives the optimal weight

$$m^*_{g,\tau} \;=\; \frac{\Sigma^2_{E,\tau}}{\sigma^2_F + \Sigma^2_{E,\tau} + \kappa_g/2}, \qquad \Sigma^2_{E,\tau} = \sigma^2_E + b^2_\tau. \tag{2.3.5}$$

Two comparative statics organise everything that follows. The weight is increasing in $\Sigma^2_{E,\tau}$: when experienced prices become a worse guide to headline inflation, attending to financial news is more valuable, and decreasing in κ_g. The bundling discount (2.3.3) therefore implies $m^*_{S,\tau} > m^*_{N,\tau}$ at every date. For non-stockholders the attention cost is high enough that the optimal weight sits at its corner, $m^*_{N,\tau} \approx 0$: they place essentially no weight on the financial signal and stay anchored to experienced prices. The structural estimates bear this out, recovering $\kappa_S \approx 17$ for stockholders against an effectively unbounded κ_N for non-stockholders (Section 2.8). Taking group means in (2.3.2), the two groups' average forecasts and the cross-sectional gap are

$$\mu_{N,\tau} = \pi + b_\tau, \qquad \mu_{S,\tau} = \pi + (1 - m^*_{S,\tau})\, b_\tau, \qquad \Delta_\tau \equiv \mu_{N,\tau} - \mu_{S,\tau} = m^*_{S,\tau}\, b_\tau. \tag{2.3.6}$$

The gap is the product of two objects: how much attention stockholders pay, $m^*_{S,\tau}$, and how biased the experienced-price signal has become, b_τ. Appendix 2.A reports the full derivation steps.

2.3.3 Duration and bias amplification

The experienced-price bias is not a constant but grows with the cumulative salience of the inflation episode. Let S_τ denote accumulated salience: the longer headline inflation stays elevated, and the longer food and energy prices keep climbing, the more those increases dominate a household's impression of prices. The bias rises linearly in this stock,

$$b_{i,\tau} \;=\; b_0 \;+\; \gamma_i\, S_\tau, \tag{2.3.7}$$

with $\gamma_i \geq 0$ the household's sensitivity of bias to salience and b_0 a baseline; the group average is $b_{g,\tau} = b_0 + \gamma_g S_\tau$. This equation, dated by how long the episode has run, is where duration enters the model, and it drives the gap (2.3.6) through two margins at once. As S_τ accumulates, b_τ rises;

this raises $\Sigma^2_{E,\tau} = \sigma^2_E + b^2_\tau$, which through (2.3.5) raises stockholder attention $m^*_{S,\tau}$, the *intensive margin*, and at the same time feeds directly into the non-stockholder anchor $\mu_{N,\tau} = \pi + b_\tau$, *bias amplification*. Because $\Delta_\tau = m^*_{S,\tau}\, b_\tau$ is the product of the two, the gap widens on both counts as the episode lengthens.

Two features of this mechanism deserve emphasis. First, there is no participation margin: non-stockholders never leave their corner, $m^*_{N,\tau} \approx 0$ throughout, so the widening reflects rising attention among stockholders and rising bias among non-stockholders, not any compositional movement of households into financial-signal use. Second, what moves the gap is the *duration* of the episode through S_τ, not its peak magnitude. Two episodes with the same peak in Σ^2_E but different lengths accumulate different salience: a brief spike leaves S_τ, and hence b_τ and $m^*_{S,\tau}$, close to baseline and the gap near zero, whereas a sustained episode accumulates enough salience to move both margins. This is the model's reading of why 2008, an episode of comparable volatility but far shorter duration, produced no gap, while the prolonged post-2021 surge did.

2.3.4 Testable predictions

The model delivers three predictions that map onto the empirical sections.

Prediction 1 (Duration widening). The gap $\Delta_\tau = m^*_{S,\tau} b_\tau$ is increasing in cumulative salience S_τ. It stays near baseline in tranquil periods and widens only as an episode persists. This is tested by the difference-in-differences widening after 2021 set against the placebo null in the shorter 2008 episode (Section 2.5).

Prediction 2 (Tail compression). Because stockholder attention pulls forecasts toward the unbiased signal while non-stockholders carry the full amplified bias, the divergence concentrates where beliefs are most extreme. The gap is widest in the upper tail of the belief distribution rather than a uniform downward shift, a distributional signature tested by the quantile gradient.

Prediction 3 (Heterogeneity in γ_i). Households whose price exposure makes salient items loom larger (e.g. lower-income households and women, who shop more often and spend a larger share on food and energy) have higher sensitivity γ_i and accumulate bias faster, producing larger individual gaps. Sensitivity should track exposure rather than analytical capacity, so education and age should not predict it. This is tested by the triple-interaction heterogeneity results.

2.4 Data

2.4.1 The Michigan Survey of Consumers

The data come from the University of Michigan Survey of Consumers, a nationally representative survey of U.S. households which has been conducted monthly since 1978. The survey is of rotational nature with a repeated cross-section: each month the survey interviews around five hundred households, a portion of whom are re-interviewed six months later. The central variable is the household's quantitative one-year-ahead inflation expectation: the expected percent change in prices over the next twelve months. Following standard practice I discard non-substantive responses and restrict the variable to the range $[-5, 50]$ percent to remove implausible outliers. All means, regressions, and distributional statistics use the survey's sampling weights, except where an estimator does not admit them, as noted in place. Because the stock-ownership question is asked on a continuous monthly basis only from 1997 onward, the estimation sample runs from 1997 through 2025, spanning the tranquil pre-pandemic decades, the 2008 financial crisis, and the 2021–2025 inflation surge.

2.4.2 Variable construction and stockholder status

A household's classification as a stockholder is based on the survey's investment item, which asks whether the household holds any investments in the stock market: I set $Stock_i = 1$ for respondents who answer yes and 0 for those who answer no, dropping the small number who do not respond. The treatment timing follows the surge, with $Post_t = 1$ from January 2021 onward and 0 before. The coefficient of interest throughout is on the interaction $Stock_i \times Post_t$, which measures how stockholders' expectations moved relative to non-stockholders' as the surge took hold.

The remaining covariates are built from the survey's standard demographic and socioeconomic items. Two demographic indicators identify respondents under 45 and women. Two socioeconomic indicators identify college graduates, those holding at least a bachelor's degree, and high-income households, defined as the top two quintiles of the survey's income distribution. A homeownership indicator is constructed in the same binary fashion for use in the housing robustness check.

Two further items are constructed for later use. A specific question records the dollar value a stockholder reports holding in the market. I discard the survey's sentinel non-response codes and, where the reported amount is positive, take natural logs. A separate item, available from 2007, records the household's one-year-ahead expectation for home prices and is cleaned in the same way as the inflation expectation.

2.4.3 Summary statistics

Table 1 reports survey-weighted means for the pre-2021 sample, split by stock ownership. Before the surge the two groups already differed in expectations, but only modestly: stockholders expected inflation of 3.30 percent on average against 4.07 percent for non-stockholders, a gap of 0.76 percentage points that had been roughly stable for two decades. Stockholders were also more educated and higher-income, as one would expect (55 against 22 percent college graduates, and 54 against 14 percent in the top income quintiles), while the two groups were similar in age and gender composition. These compositional differences are precisely why the design conditions on them: the question is not whether stockholders differ from non-stockholders, but whether the gap in expectations *widened* as the surge persisted, beyond what these stable traits can explain.

Table 1: Summary statistics, pre-2021 period

	Stockholders	Non-stockholders	Difference
One-year inflation expectation (%)	3.30 (3.44)	4.07 (4.22)	−0.76
Age below 45 (share)	0.39 (0.49)	0.41 (0.49)	−0.02
Female (share)	0.49 (0.50)	0.57 (0.50)	−0.08
College graduate (share)	0.55 (0.50)	0.22 (0.42)	+0.33
High income (share)	0.54 (0.50)	0.14 (0.34)	+0.40
Observations	90,636	54,448	

Notes: Pre-2021 sample (survey months before January 2021) with non-missing stock-ownership status. Entries are survey-weighted means with standard deviations in parentheses; the final column is the stockholder-minus-non-stockholder difference in means. The inflation expectation is the cleaned one-year-ahead `PX1` series (N = 145,084).

2.5 Empirical Strategy and Main Results

This section will introduce the main Difference-in-differences specification, detailing the features of the econometric techniques adopted. The main results will be presented and analysed next, together with some of the core robustness checks which are specific to the Difference-in-differences approach.

2.5.1 Difference-in-differences specification

The estimating equation compares the change in stockholders' inflation expectations across the surge to the contemporaneous change for non-stockholders:

$$E_{i,t} \;=\; \beta\left(Stock_i \times Post_t\right) + \gamma\, Stock_i + X'_{i,t}\theta + \lambda_t + \varepsilon_{i,t}, \tag{2.5.1}$$

where $E_{i,t}$ is the one-year-ahead expectation, λ_t are survey-year fixed effects that absorb the $Post_t$ main effect and any aggregate movement in expectations, $X_{i,t}$ is a vector of controls, and standard errors are clustered by survey month. The parameter of interest is β, the coefficient on $Stock_i \times Post_t$ which measures how much stockholders' expectations rose relative to non-stockholders' once the surge took hold. Because stockholder status is time-invariant and the treatment timing is common to all treated units, the standard two-way fixed effects estimator is appropriate here. For reference, heterogeneous-timing extensions that allow for differential treatment paths (Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfœuille, 2020) confirm the headline estimate in Section 2.7.

I report four specifications that add controls in layers. The *baseline* includes only the stockholder indicator and year fixed effects. *+Demographics* adds the age and gender indicators; *+SES* adds the college-graduate and high-income indicators. The headline *+Demog.×Post* specification interacts every demographic and socioeconomic control with $Post_t$, so that β is identified only off the stockholder-specific change across the surge, net of any change in expectations common to younger, female, college-educated, or higher-income households over the same period. This is the most demanding of the four and the one I treat as the headline. A fifth specification adding a homeownership interaction is reported as a robustness check in

Section 2.7.

2.5.2 Baseline estimates

Table 2 reports the four specifications. In the baseline, stockholders' expectations rose by roughly 2.48 percentage points less than non-stockholders' across the surge. Adding demographic controls leaves this essentially unchanged (−2.49); adding the socioeconomic controls pulls it to −1.94; and the headline specification, which interacts every control with *Post*, settles it at −1.70 (SE 0.190) on 166,125 observations. The attenuation is meaningful: stockholders are older, wealthier, and better educated, and those traits do predict lower expectations after the surge. However, more than two-thirds of the raw gap survives the most demanding specification, and all four estimates are significant at the 1% level. The headline −1.70 is the conservative anchor net of all demographic effects which is carried through the rest of the paper.

Table 2: Stock ownership and inflation expectations: difference-in-differences

	Baseline	+Demog.	+SES	+Demog.×Post
Stock × Post	-2.477^{***}	-2.487^{***}	-1.938^{***}	-1.702^{***}
	(0.220)	(0.221)	(0.205)	(0.190)
Survey-year FE	Yes	Yes	Yes	Yes
Demographic controls	No	Yes	Yes	Yes
Socioeconomic controls	No	No	Yes	Yes
Controls × Post	No	No	No	Yes
Observations	166,125			

Notes: Dependent variable: cleaned one-year-ahead inflation expectation. Each column reports the coefficient on *Stock*×*Post* from a separate survey-weighted regression; standard errors clustered by survey month in parentheses. All columns include survey-year fixed effects. The headline specification is the +Demog.×Post column. $^{***}p < 0.01$.

2.5.3 Event study

To see *when* the gap opened, I replace the single *Post* interaction with a full set of $Stock_i$ × year interactions, normalised to the last pre-surge year (2019). Figure 2 plots the resulting coefficients. Through 2020 they are small and hover near zero. The largest pre-period deviation, in 2010, is only 0.49 percentage points and a joint test that the 2012–2018 leads are zero cannot reject it

(p = 0.279), so the parallel-trends assumption is supported. The gap then opens only after the surge and *widens as it persists*: the 2021 coefficient is still statistically indistinguishable from zero, but it falls to −1.32 in 2022, −2.32 in 2023, and −2.38 in 2024. The effect is thus not an instantaneous response to the onset of inflation but a divergence that builds over the duration of the episode, the reduced-form signature of Prediction 1.

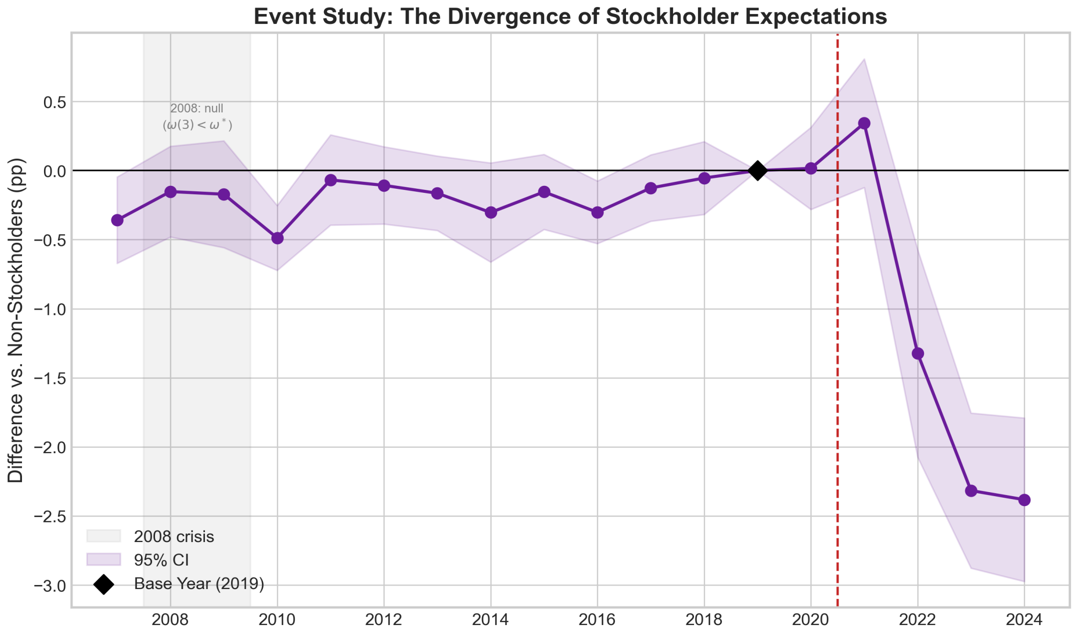

Figure 2: Event study of the stockholder–non-stockholder expectations gap. Coefficients on *Stock* × year interactions from the difference-in-differences specification, normalised to 2019; bands are 95% confidence intervals from standard errors clustered by survey month.

2.5.4 The 2008 placebo: duration, not magnitude

The event study shows the gap opening only after 2021, but the sharpest identification check compares the surge to an earlier episode of comparable inflation volatility and far shorter duration: the 2008 commodity-price spike. If the widening were a generic response to volatile inflation, 2008 should have produced a similar gap. It did not. Re-estimating the difference-in-differences around a 2008 cutoff, with the same controls, yields a stockholder effect of −0.085 (SE 0.097, p = 0.38, N = 40,291), economically and statistically zero.

The contrast is the empirical core of the duration mechanism. Crucially, it is not a contrast in the salient price pressure that drives experienced-price bias: peak 24-month cumulative food-and-energy exposure was almost identical across the two episodes (a ratio of 1.04). What separated them was duration. Only the post-2021 surge sustained high inflation long enough for cumulative salience S_τ to accumulate, the experienced-price bias b_τ to amplify, and the gap to widen. The transitory 2008 spike ended before salience could build, leaving b_τ near baseline and the change in the gap near zero, exactly what (2.3.7) predicts. Two further placebo cutoffs (2015 and 2018), reported in Section 2.7, similarly produce no comparable widening.

2.5.5 Forecast accuracy and Mincer–Zarnowitz

A lower expectation is an improvement only if it is also more accurate, and it is. Stockholders' *absolute* forecast errors fell by 1.39 percentage points more than non-stockholders' across the surge (SE 0.18), and their *signed* errors (the upward bias) by 1.47 points more (SE 0.19), both on 159,865 observations (Table 3). The gain is one of bias, not efficiency. Column 3 of Table 3 reports an augmented Mincer–Zarnowitz regression in which realised inflation is regressed on the individual forecast together with $Stock \times Post$. If stockholders processed any given signal more accurately, the interaction would enter with a negative coefficient. Instead the coefficient on the expectation is 1.002 (SE 0.001), correctly scaled in aggregate, and $Stock \times Post$ is +0.014 (SE 0.011), indistinguishable from zero. Conditional on the forecast level, stockholder status predicts nothing about who came closer to realised inflation. The accuracy advantage operates entirely through the forecast level: attending to the financial signal produces lower expectations, and in a prolonged surge where food and energy prices ultimately subsided faster than headline inflation, those lower expectations were systematically, not randomly, closer to realised inflation.

The Mincer–Zarnowitz decomposition sharpens the picture. Regressing realised inflation on each group's mean forecast over the post-2021 window, stockholders retain a positive slope ($\beta = 0.10$) while non-stockholders' slope turns *negative* ($\beta = -0.49$, SE 0.15): as the surge wore on, non-stockholders' forecasts moved opposite to realised headline inflation, tracking the salient food and energy prices that anchor experienced-price beliefs rather than the headline index. Over the full sample both slopes are positive (0.30 and 0.19). The post-2021 sign flip for non-stockholders is the reduced-form counterpart of the bias-amplification mechanism, and a Wald test rejects forecast rationality for both groups.

Table 3: Stock ownership and forecast accuracy

	Absolute error	Signed bias	Efficiency
Stock × Post	-1.390^{***}	-1.467^{***}	0.014
	(0.178)	(0.189)	(0.011)
Stockholder	-0.470^{***}	-0.493^{***}	-0.014^{**}
	(0.024)	(0.028)	(0.006)
Expectation (E_π)	—	—	1.002^{***}
			(0.001)
Headline controls	Yes	Yes	Yes
Observations		159,865	

Notes: The dependent variable is the absolute forecast error (column 1) and the signed forecast error, i.e. expectation minus realised inflation (columns 2–3). The efficiency column adds the expectation E_π as a regressor: an insignificant *Stock* × *Post* there indicates that stockholders' accuracy gain is a reduction in bias rather than a difference in efficiency. All columns use the headline controls; standard errors clustered by survey month in parentheses. $^{**}p < 0.05$, $^{***}p < 0.01$.

2.6 Mechanism

This section is built on a comprehensive set of post-estimation tests which are aimed at unpacking the mechanisms behind the observed widening gap in inflation expectations between stockholders and non-stockholders.

2.6.1 Quantile regression

The headline coefficient is an average, but the mechanism predicts the effect should be concentrated in the upper tail rather than spread uniformly (Prediction 2). Quantile regressions (Koenker and Bassett, 1978) confirm this. Estimated at the 25th, 50th, and 75th percentiles of the expectations distribution, the stockholder effect steepens sharply across the distribution: −0.45 (SE 0.04) at the 25th percentile, −1.12 (SE 0.04) at the median, and −2.78 (SE 0.06) at the 75th, a gradient of roughly 6.1 to one between the upper and lower quartiles (Figure 3). The gap is thus overwhelmingly a right-tail phenomenon: stockholders are far less likely to hold the extreme high expectations that the surge produced. Two features guard against reading this as a mechanical artifact of one group's dispersion. Each coefficient is the effect of *Stock* × *Post at* that quantile, not the level of non-stockholder beliefs and the pattern is one-sided (negligible at the 25th percentile, large at the 75th), rather than a symmetric widening. It is also dynamic: the tail gap is absent before 2021 and in the 2008 placebo (Section 2.5), so it cannot reflect a fixed tendency for non-stockholders to hold more dispersed beliefs. Winsorising confirms that this is a tail effect rather than a uniform shift (Section 2.7). A mild outlier trim leaves the headline estimate essentially intact, -1.43 at the 1st/99th percentiles against the headline -1.70, so the result is not driven by a handful of extreme respondents. Capping beliefs at the 95th percentile (around 13 percent), which removes precisely the extreme expectations the mechanism concerns, cuts the estimate to -0.39 but leaves it strongly significant ($p < 0.001$). The action lives in the upper tail, yet a significant effect survives even when that tail is compressed.

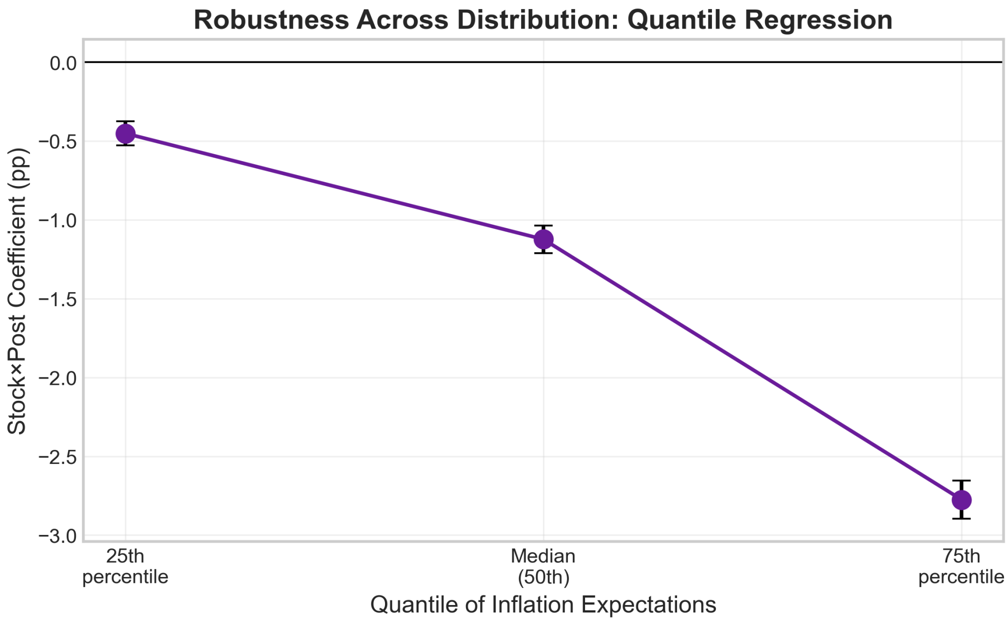

Figure 3: Quantile gradient of the stockholder–non-stockholder gap. Coefficients on *Stock* × *Post* estimated at each quantile of the one-year-ahead expectations distribution (N = 169,517; unweighted quantile regression, asymptotic standard errors).

2.6.2 Heterogeneity

If the bias channel runs through experienced-price salience, the effect should be larger for households more exposed to salient prices and should not depend on analytical sophistication (Prediction 3). Triple-interaction regressions, which add $Stock \times Post \times X_i$ to the headline specification, where X_i is a binary demographic characteristic, bear this out (Table 4). The effect is significantly larger for women (triple interaction -1.419, SE 0.302), about 1.79 times the male effect, and for lower-income households (the high-income interaction is +1.056, SE 0.415, so the effect is roughly 1.84 times larger at the bottom of the income distribution). Both groups shop more frequently and spend a larger share on food and energy, exactly the exposure that raises the bias-amplification sensitivity γ_i. Neither college education (+0.358, SE 0.326) nor youth (+0.184, SE 0.297) moves the effect. The channel appears to be grounded in exposure, not analytical capacity or shorter memory. This is the pattern γ_i predicts and a sophistication-based account does not.

Table 4: Heterogeneity in the stockholder effect (triple interactions)

	Gender (female)	Income (high income)	Education (college)	Age (young)
Stock × Post × Char.	-1.419^{***}	1.056^{**}	0.358	0.184
	(0.302)	(0.415)	(0.326)	(0.297)
Observations		166,125		

Notes: Each column adds the full set of two-way and three-way interactions of *Stock*, *Post*, and the demographic characteristic to the headline specification; the row reports the triple interaction. A negative coefficient means the stockholder effect is larger (more negative) for that group. Survey-weighted estimates; standard errors clustered by survey month in parentheses. $^{**}p < 0.05$, $^{***}p < 0.01$.

2.6.3 Within-stockholder dose-response

The group-average effect also masks a gradient *within* stockholders. Sorting stockholders by the dollar amount they report holding in the market, the effect rises monotonically with the size of the position. Relative to the smallest-holding quartile, expectations fall by 0.45, 0.89, and 1.63 percentage points more across the second, third, and top quartiles after the surge and, more tellingly, absolute forecast errors fall by 0.32, 0.48, and 1.00 points (Table 5); the continuous specification gives a log-holding interaction of −0.32 (SE 0.04). The accuracy gradient in the second column is the key result: larger holders do not merely report lower numbers, they forecast more accurately. A placebo that assigns a pseudo-surge in 2014 to the pre-2021 sample produces no gradient on any quartile, and the pattern survives saturating the regression with observables (about a tenth of the gradient attenuates).

I read this as descriptive enrichment of the mechanism, not a structural estimate. It shows that the group-average treatment effect summarises a distribution of responses that widens with stakes. A plausible reading is stakes-driven attention allocation in the spirit of Sims (2003) (households with more at risk attend more), but because the design conditions on stockholder status and observed holdings and cannot absorb unobserved wealth or preferences, I do not interpret the gradient as a structural attention parameter.

Table 5: Within-stockholder dose-response by holding size

INVAMT quartile × Post	Expectation	Absolute error
Q2	-0.452^{**}	-0.320^{*}
	(0.218)	(0.178)
Q3	-0.887^{***}	-0.483^{***}
	(0.169)	(0.133)
Q4 (largest)	-1.625^{***}	-1.002^{***}
	(0.232)	(0.234)
Observations	92,980	88,687

Notes: Within-stockholder sample. Quartiles of reported INVAMT; the smallest-holding quartile (Q1) is the omitted reference. Each column is a separate survey-weighted regression with the headline controls; standard errors clustered by survey month in parentheses. The continuous analogue is $\log(\texttt{INVAMT}) \times Post = -0.321$ (SE 0.043). A placebo assigning a pseudo-post in 2014 yields no significant gradient. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

2.7 Robustness

The headline survives a battery of checks, collected in Table 6 and summarised here.

Additional placebo cutoffs. Beyond the 2008 placebo (Section 2.5), assigning a pseudo-treatment in 2015 or 2018 should also produce no significant results. The 2015 cutoff is precisely null (+0.099, $p = 0.15$). The 2018 cutoff is small but marginally significant with the *wrong* sign (+0.198, $p = 0.008$), about a tenth of the main effect and opposite in direction, the kind of false positive expected among several placebos rather than a confound that could generate the post-2021 widening.

Sensitivity to parallel-trend violations. The pre-trend is clean (joint F-test on 2012–2018 leads, $p = 0.279$), and the 2008 placebo (−0.085, $p = 0.380$) independently rules out a confound that operated through stockholder composition in a prior episode of elevated inflation volatility. The Rambachan and Roth (2023) smoothness restriction formalises this: allowing the post-period trend to deviate from a linear extrapolation of the pre-period, the confidence interval at the extrapolation itself ($M = 0$) is $[-1.73, -1.09]$, and the result breaks only once the permitted deviation reaches $M^* \approx 0.10$ pp per year. Under the relative-magnitudes restriction, overturning the average 2021–2024 estimate requires post-period violations roughly 0.55 times the largest pre-period deviation, sustained across four consecutive years. This is a cumulative drift that the flat pre-trend and the 2008 placebo jointly make difficult to attribute to confounding.

Alternative treatment timing. Moving the post-period start to March 2021, July 2021, or January 2022 leaves the estimate in the same range (−1.71, −1.63, −2.34); the effect is not an artifact of the exact cutoff.

Sample and composition. Restricting to phone respondents (−1.892) and reweighting cells to hold demographic composition fixed (−2.402) both preserve the effect, as does adding a homeownership interaction (−1.575). The widening is not an artifact of the survey's mode transition or of shifting sample composition.

Compositional change in the stockholder group. The 2020–21 retail-trading boom raises a sharper version of this concern: the post-period stockholders are not the same people as the pre-period stockholders. The share reporting stock-market investments does rise across the cutoff, from 60 to 68 percent, so the treatment group's membership changes materially. Three facts contain the threat. First, the observable part of the shift is fully absorbed by the design as post-2021 stockholders are more educated (55 $\rightarrow$ 67 percent college) and slightly older, but the headline $Demog. \times Post$ specification lets every demographic group take its own post-period shift. Therefore, β is identified within demographic$\times$period cells rather than off the changing mix, and the cell-reweighted estimate which holds the pre-period demographic distribution fixed (−2.402) leaves the effect intact. Second, the within-household panel holds the stockholder label constant by construction, it follows the same households across the cutoff, and returns the same sign (−0.124), insignificant only due to power. Third, new entrants are by revealed preference the most marginal holders, with the smallest positions and the weakest portfolio engagement. The dose-response of Section 2.6 locates the effect among large, established holders and near zero at the smallest quartile, so the divergence is not generated by who newly entered. Because these marginal entrants resemble non-stockholders in their information environment, their inclusion in the post-period treatment group attenuates the gap, making the headline a lower bound with respect to this channel.

Partisanship. Since partisan alignment tracks the party in power, I add a triple interaction with an out-of-power indicator on the full-SES sample; the stockholder effect is unchanged (-2.075), so it is not a relabelled partisan response.

Specificity to inflation. The accuracy advantage is specific to consumer prices. The difference-in-differences interaction on *home-price* forecast errors is small and insignificant (-0.088, SE 0.132), and the home-price expectation series shows a non-parallel pre-trend, marking it a distinct (rate-anticipation) channel rather than general forecasting skill. I make no accuracy claim about housing forecasts.

Within-household panel. Exploiting the survey's short rotating panel, a within-household specification gives the same sign but is insignificant (-0.124, $p = 0.67$). This is expected, since fewer than five percent of households are re-interviewed across the cutoff, too little power to detect the effect rather than evidence against it.

Winsorization. As discussed in Section 2.6, a mild $1/99$ trim leaves the estimate essentially intact (-1.428) while capping at the 95th percentile shrinks it (-0.388) but keeps it significant, confirming the tail-compression reading rather than fragility.

Table 6: Consolidated robustness

	Coef.	SE	N
Headline			
+Demog.×Post	-1.702^{***}	0.190	166,125
Sample and composition			
Phone-only (CATI)	-1.892^{***}	0.228	146,801
Cell-reweighted	-2.402^{***}	0.228	164,766
+ Homeownership interaction	-1.575^{***}	0.183	164,766
Alternative timing			
Post = March 2021	-1.709^{***}	0.261	166,125
Post = July 2021	-1.630^{***}	0.256	166,125
Post = January 2022	-2.343^{***}	0.192	166,125
Partisanship			
Triple-diff (full SES)	-2.075^{***}	0.271	77,574
Placebo cutoffs			
2015	+0.099	0.069	58,285
2018	$+0.198^{**}$	0.074	49,319
Winsorization			
1/99	-1.428^{***}	0.141	164,380
5/95	-0.388^{***}	0.082	164,380
Specificity			
Home-price forecast errors (MAE)	−0.088	0.132	90,821
Within-household panel			
Panel DiD	−0.124	0.292	51,664

Notes: Each row reports the *Stock* × *Post* coefficient from a separate regression; the headline +Demog.×Post row repeats Table 2 for reference. All rows use one-year CPI inflation expectations as the dependent variable except the specificity row, which uses absolute home-price forecast errors. Standard errors clustered by survey month (the within-household panel clusters by household). The partisan triple-diff is estimated on the POLAFF subsample. $^{**}p < 0.05$, $^{***}p < 0.01$.

2.8 Structural Estimation

The reduced form establishes a widening gap, its concentration in the upper tail, and its accuracy. The structural exercise asks whether the model's two margins can reproduce the group means quantitatively, and what attention costs they imply. Through the forecast rule (2.3.2) and the attention FOC (2.3.5), group g's mean expectation in period τ is

$$\mu_{g,\tau} \;=\; m^*_{g,\tau}\,\pi^F_\tau + \left(1 - m^*_{g,\tau}\right)\left(\pi^E_\tau + b_\tau\right). \tag{2.8.1}$$

I proxy the financial signal π^F by the Cleveland Fed expected-inflation series (unbiased and professionally informed) and the experienced-price signal π^E by food-and-energy inflation (the salient, frequently purchased prices that anchor experienced beliefs), and calibrate the signal noises to $\sigma_F^2 = 0.25$ and $\sigma_E^2 = 4.0$, tying each to an observable dispersion: $\sigma_F \approx 0.5$ pp matches the cross-sectional standard deviation of professional (SPF) forecasts, and $\sigma_E \approx 2$ pp the standard deviation of the food-and-energy-minus-headline wedge, a signal-to-noise ratio of about 16 to 1 in variance. This leaves four parameters, the two attention costs κ_S, κ_N and the bias in each period $b_{\text{pre}}, b_{\text{post}}$,[1] to match the four observed group means ($\bar{E}_S, \bar{E}_N$ before and after 2021), with the attention weights following from the FOC (2.3.5). The four moment conditions are

$$\begin{aligned}
m^*(\kappa_S, b_{\text{pre}})\,\pi^F_{\text{pre}} + \big[1 - m^*(\kappa_S, b_{\text{pre}})\big](\pi^E_{\text{pre}} + b_{\text{pre}}) &= \bar{E}_{S,\text{pre}},\\
m^*(\kappa_N, b_{\text{pre}})\,\pi^F_{\text{pre}} + \big[1 - m^*(\kappa_N, b_{\text{pre}})\big](\pi^E_{\text{pre}} + b_{\text{pre}}) &= \bar{E}_{N,\text{pre}},\\
m^*(\kappa_S, b_{\text{post}})\,\pi^F_{\text{post}} + \big[1 - m^*(\kappa_S, b_{\text{post}})\big](\pi^E_{\text{post}} + b_{\text{post}}) &= \bar{E}_{S,\text{post}},\\
m^*(\kappa_N, b_{\text{post}})\,\pi^F_{\text{post}} + \big[1 - m^*(\kappa_N, b_{\text{post}})\big](\pi^E_{\text{post}} + b_{\text{post}}) &= \bar{E}_{N,\text{post}},
\end{aligned} \tag{2.8.2}$$

where $m^*(\kappa, b) = (\sigma_E^2 + b^2)/(\sigma_F^2 + \sigma_E^2 + b^2 + \kappa/2)$ from (2.3.5), and the right-hand sides are the observed survey group means. The system is solved exactly by a standard nonlinear equation solver. Identification works as follows. Because non-stockholders place essentially no weight on the financial signal, their level pins the bias directly, $b_\tau \approx \bar{E}_{N,\tau} - \pi^E_\tau$ (about 1.23 before and 2.13 after); given b_τ, the stockholder level pins the attention weight $m_{S,\tau}$ through the forecast rule, and a single time-invariant κ_S reconciles both periods' weights via the FOC. The cost κ_N is pinned by whether the non-stockholder level requires any weight on the financial signal at all: it does not, so the estimator drives κ_N to its bound, the $m_N = 0$ corner is identified, not imposed.

Table 7 reports the estimates, with three core insights. First, the perceptual bias rises sharply, from 1.23 to 2.13 percentage points, the amplification margin, consistent with evidence that bias scales with salience (D'Acunto et al., 2021a). Second, stockholder attention rises from

[1] The two states for b match the two-regime difference-in-differences design: b_{pre} and b_{post} are the regime-average biases the design identifies, and at the non-stockholder corner each is read directly from the data as $\bar{E}_{N,\tau} - \pi^E_\tau$. A continuous law of motion is not warranted at this frequency: both a salience form $b_\tau = b_0 + \gamma S_\tau$ and a persistence form $b_\tau = \rho b_{\tau-1} + \gamma S_\tau$ are poorly identified, because the experienced-price signal is far more volatile than the sticky survey expectations it must track. The duration dynamics are therefore located in the reduced form (Section 2.5), not in a structural process for b.

0.39 to 0.50. As experienced prices become a worse guide (Σ_E^2 rises with b), the FOC pulls stockholders toward the financial signal. Together the two margins widen the implied gap from 0.82 to 3.03 percentage points, matching the observed means by construction. Third, and most consequentially, the non-stockholder attention cost κ_N is driven to its upper bound and m_N sits at zero in both periods. Non-stockholders place no weight on the financial signal at any point. Their expectations are $\pi^E + b$ throughout, rising only because b rises. The stockholder cost $\kappa_S \approx 16.8$ is finite and far smaller, the structural counterpart of the bundling discount δ.

Both margins are statistically distinguishable from zero. A moving-block bootstrap over months ($L = 12$), which treats the official inflation series as fixed and resamples only the survey group means, puts the rise in bias at $\Delta b = 0.89$ (95% CI $[0.32, 2.54]$) and the rise in stockholder attention at $\Delta m_S = 0.11$ (95% CI $[0.035, 0.271]$), and bounds $\kappa_S = 16.8$ (SE 5.5) well below the non-stockholder corner. The point estimates fit the four group means exactly and the bootstrap shows the two margins are identified, not artifacts of the calibration.

Table 7: Structural estimates (FOC-disciplined moment matching)

	Pre-2021	Post-2021
Calibrated signals		
Financial signal π^F (Cleveland Fed)	2.01	2.59
Experienced signal π^E (food & energy)	2.86	6.57
Target moments (observed group means)		
Stockholders $\bar{E}_S$	3.28	5.67
Non-stockholders $\bar{E}_N$	4.10	8.70
Estimated parameters (bootstrap SE in parentheses)		
Perceptual bias b (pp)	1.23 (0.22)	2.13 (0.62)
Stockholder attention m_S	0.39 (0.07)	0.50 (0.09)
Non-stockholder attention m_N	0.00	0.00
Attention cost κ_S (time-invariant)	16.75 (5.53)	
Attention cost κ_N (time-invariant)	10^6 (corner)	

Notes: Just-identified non-linear moment matching (four moments, four parameters; solved exactly by a standard nonlinear equation solver), exact fit at the point estimate. The attention weights $m^*_{g,\tau}$ are derived from the FOC (2.3.5) given b_τ and κ_g. Calibrated: $\sigma_F^2 = 0.25$, $\sigma_E^2 = 4.0$. Pre-period 1997m7–2020m12 (273 months); post-period 2021m1–2025m11 (59 months). κ_N at its upper bound denotes a corner solution ($m_N = 0$). Standard errors in parentheses are from a moving-block bootstrap over months ($L = 12$, $B = 2000$) that holds the official signal series π^F, π^E fixed and resamples the survey group means; the increases Δb and Δm_S have 95% bootstrap confidence intervals of $[0.32, 2.54]$ and $[0.035, 0.271]$, both excluding zero.

2.8.1 Assessing extensive-margin entry

The same widening admits a potential compositional interpretation: perhaps, as inflation persisted, some non-stockholders began attending to financial signals, an extensive-margin *entry* into the informed group, rather than accumulating bias. The two stories share a reduced form but differ structurally, and the data favour the intensive one decisively. The moment-matching estimator already forecloses entry: the best level-fit pins m_N at zero in both periods and κ_N at its bound, so non-stockholders have no attention to the financial signal that could grow as the surge persists. A direct test of the entry prediction agrees. If duration drove non-stockholders to become informed, the share holding 'informed' expectations (close to the professional signal) should have *risen* post-2021, most where the gain is largest. Instead it *fell* by about 18 to 19 percentage points in every demographic subset, with difference-in-differences interactions that are uniformly small and insignificant (-0.009, $p = 0.65$ by income; +0.010, $p = 0.46$ by gender). In addition, the decline is *larger* in 2023–25 than in 2021–22, the opposite of progressive entry. A structural mixture model and a battery of alternative informed-band definitions confirm the pattern (Appendix 2.F). The widening is intensive-margin attention and bias amplification, not composition.

2.8.2 Robustness of the structural estimates

The estimates rest on two calibrated choices, the signal-noise variances and the experienced-price signal, and survive both (Table 8). Across a 3×3 grid of (σ_F^2, σ_E^2) the stockholder cost stays in $[12.8, 22.2]$, the bias always rises (Δb between +0.9 and $+1.2$), stockholder attention always rises, and non-stockholders sit at, or within 0.06 of, the $m_N = 0$ corner. The qualitative structure is essentially invariant to σ_F^2. Replacing food-and-energy with headline, median, trimmed-mean, or sticky-price inflation leaves the same pattern in every case: rising bias, rising stockholder attention, and the corner. Only the levels move, smoother signals sit further below expectations and so imply a larger bias, which makes food-and-energy the conservative choice, delivering the *smallest* attention shift. A finite stockholder cost, rising bias, rising stockholder attention, and a non-stockholder corner are thus features of the data, not of the specific calibration choice.

Table 8: Robustness of the structural estimates

	b_{pre}	b_{post}	Δm_S	κ_S	m_N
Panel A. Experienced-price signal π^E ($\sigma_F^2 = 0.25, \sigma_E^2 = 4.0$)					
Food & energy (baseline)	1.23	2.13	0.107	16.8	0.00
Headline CPI	1.87	4.17	0.241	41.6	0.00
Median CPI	1.55	4.18	0.272	41.6	0.00
Trimmed-mean CPI	1.86	4.42	0.261	45.7	0.00
Sticky-price CPI	1.64	4.41	0.280	45.4	0.00
Panel B. Signal noise σ_E^2 (food & energy, $\sigma_F^2 = 0.25$)					
$\sigma_E^2 = 2.25$	1.21	2.12	0.144	13.0	0.00
$\sigma_E^2 = 4.00$	1.23	2.13	0.107	16.8	0.00
$\sigma_E^2 = 6.25$	1.31	2.49	0.110	22.0	0.06

Notes: Each row re-solves the just-identified four-moment system of Table 7 under a different calibration. $\Delta m_S = m_{S,\text{post}} - m_{S,\text{pre}}$; m_N is the post-period non-stockholder attention weight. "0.00" denotes the κ_N corner; in the sole off-corner cell ($\sigma_E^2 = 6.25$) the estimator returns $\kappa_N \approx 422$, still an order of magnitude above κ_S, so $m_N \approx 0.06$. Results are essentially invariant to $\sigma_F^2 \in \{0.16, 0.25, 0.36\}$ (not shown). π^F is the Cleveland Fed series throughout.

2.8.3 Modelling tail compression

The moment-matching estimator targets four group means. The same parameters, however, carry an untargeted prediction about the entire cross-sectional distribution, which can be compared to the data as an over-identifying check. The procedure has three steps.

First, because $m_N = 0$, non-stockholders place no weight on the financial signal and their forecast collapses to $E_i^N = \pi_i^E + b$. Their observed expectation distribution is therefore identical to the latent experienced-signal-plus-bias distribution, it is the raw unanchored baseline.

Second, stockholders draw from that same latent distribution of experienced signals but pull each forecast toward π^F with weight m_S, so $E_i^S = m_S\pi^F + (1 - m_S)E_i^N$ in distribution. Because this map is affine and monotone, every quantile transforms the same way:

$$Q_\tau\big(E^S\big) \;=\; m_S\,\pi^F + \big(1 - m_S\big)\,Q_\tau\big(E^N\big). \tag{2.8.3}$$

Third, the predicted stockholder quantiles (2.8.3) are computed using only m_S and π^F from the moment-matching step (no new parameters) and compared to the actual stockholder quantiles in the data. The predicted quantile difference-in-differences (*Stock* × *Post* effect at each percentile) is then set against the raw data and the reduced-form quantile-regression estimates from Section 2.6.

Since $(1-m_S) < 1$, the map pulls every quantile toward π^F, but the pull is proportional to how far $Q_\tau(E^N)$ sits above the anchor, which is largest in the upper tail. The gap $Q_\tau(E^N) - Q_\tau(E^S) = m_S[Q_\tau(E^N) - \pi^F]$ is therefore near zero at low percentiles, where non-stockholder beliefs already sit close to π^F, and large at the top. Table 9 and Figure 4 confirm that the model reproduces the gradient independently: a *Stock* × *Post* effect near zero at the 25th percentile, −0.81 at the median, and −2.51 at the 75th, against the reduced-form quantile-regression estimate of −2.78, steepening further into the upper tail. The tail-compression signature of Section 2.6 is therefore not merely a descriptive regularity, but rather a structural prediction of the attention model.

Two caveats need to be reported. The center of the distribution is partly mechanical: m_S was calibrated to the group means, so the median is reproduced almost by construction. The genuine, untargeted content is the upper tail, where the same m_S that fits the means also governs the dispersion, which it was under no constraint to do. And at the extreme top the model slightly *over*-compresses (a 90th-percentile effect of −8.0 against −7.0 in the data).

Table 9: Tail compression as a structural prediction

Quantile τ	Model (2.8.3)	Raw quantile DiD	QR coefficient
25th	−0.10	0.00	−0.45
50th	−0.81	−1.00	−1.12
75th	−2.51	−3.00	−2.78
90th	−8.02	−7.00	-5.73^{***}

Notes: Stock × *Post* effect on the τ-quantile of one-year expectations. *Model*: the unconditional quantile difference-in-differences implied by (2.8.3), using the estimated m_S ($0.39 \to 0.50$) and π^F ($2.01 \to 2.59$) and the observed non-stockholder quantiles—no additional free parameters. *Raw quantile DiD*: the corresponding unconditional 2×2 statistic in the microdata. *QR coefficient*: the conditional quantile-regression estimate of Section 2.6 (controls and year fixed effects), reproduced for reference; being conditional it differs from the unconditional columns, most at the 25th percentile.

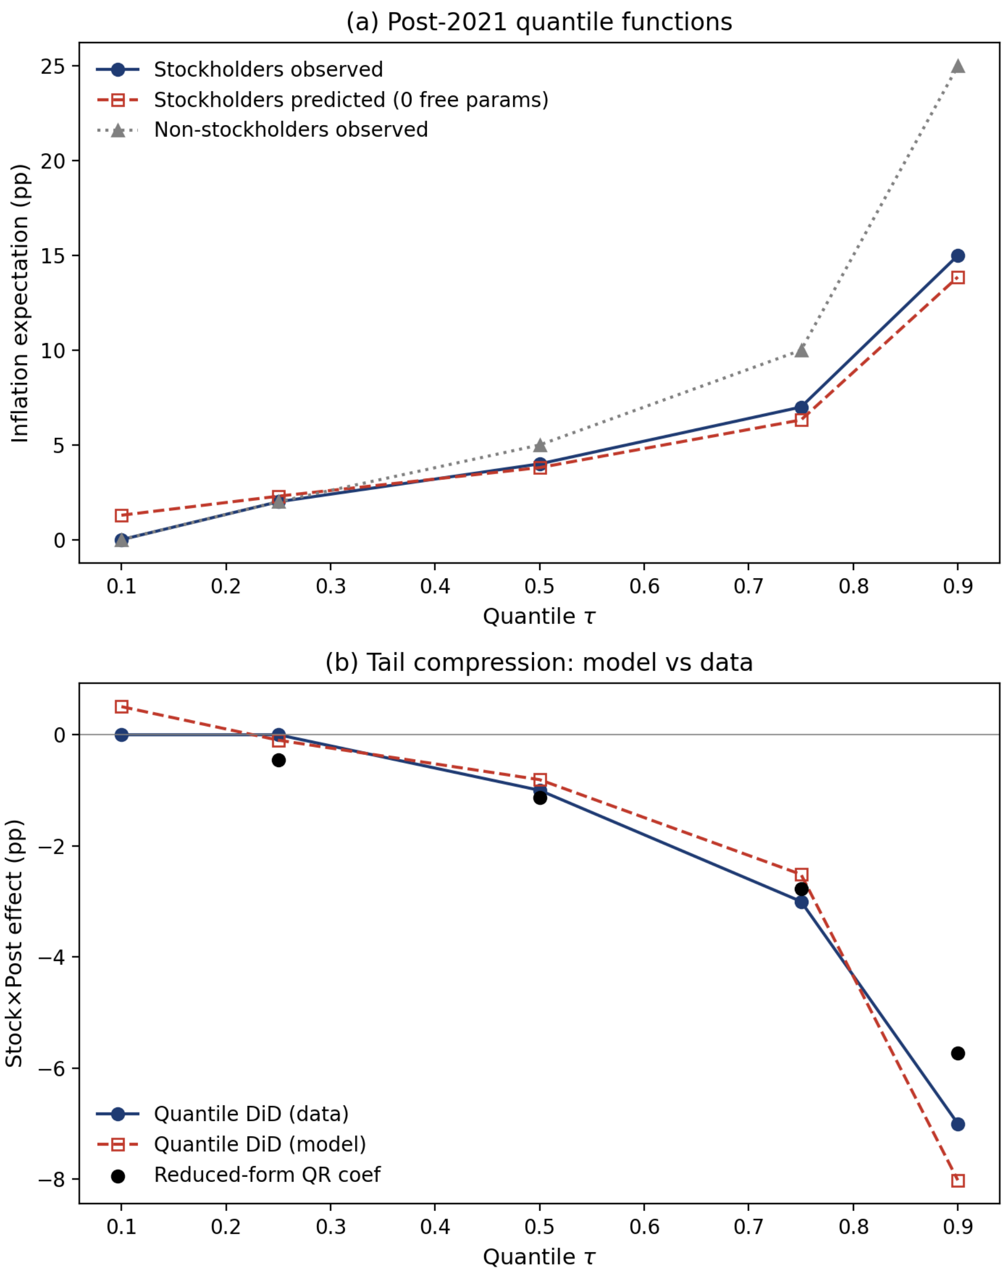

Figure 4: Tail compression as a structural prediction. Panel (a): the observed post-2021 stockholder quantile function (solid) against the zero-free-parameter prediction of (2.8.3) (dashed), with the non-stockholder quantiles (dotted) as the input distribution. Panel (b): the *Stock* × *Post* quantile difference-in-differences in the data (solid) and as implied by the model (dashed), with the reduced-form quantile-regression coefficients (points).

2.9 Conclusion

The gap between stockholders' and non-stockholders' inflation expectations was small and stable for two decades and then more than tripled during the 2021–2025 surge. In the most conservative specification stockholders' expectations rose 1.70 percentage points less than non-stockholders', and the divergence was an improvement in accuracy, not merely a lower number: stockholders' absolute forecast errors and upward bias both fell relative to non-stockholders'. The effect is concentrated in the upper tail, so it operates by preventing extreme expectations rather than shifting the distribution uniformly downward.

What separates 2021 from 2008 is duration, not magnitude. The two episodes carried almost identical peak experienced-price pressure, yet 2008 produced no gap. Only the sustained post-2021 episode kept salient food and energy prices elevated long enough for perceptual bias to accumulate. A two-signal rational-inattention model with an information-bundling discount on stockholders' attention cost organises the evidence: as the bias grows, stockholders raise their weight on the unbiased financial signal (the intensive margin) while non-stockholders, anchored to experienced prices, absorb the rising bias (amplification). A first-order-condition-disciplined moment-matching estimator recovers both margins. The bias rises from 1.2 to 2.1 points and stockholder attention from 0.39 to 0.50, with non-stockholder attention pinned at zero throughout. The widening is intensive-margin and amplification, not compositional entry, a reading four direct tests confirm. The heterogeneity of the effect, larger for lower-income households and women, flat across education and age, points to exposure to salient prices rather than analytical sophistication as the operative margin.

The findings carry a communication implication, and a diagnostic one before a prescriptive one. When inflation is brief, households that do and do not monitor financial signals form similar expectations. When it persists, and in particular during episodes where the experienced component of prices is prominent, their beliefs diverge, with the less-attended population drifting upward through the prices it is exposed to most. The contribution of the stock-ownership margin is to identify *who* drifts and *why*, households that read inflation off their own purchases rather than aggregate data, not merely to recommend that more of them be brought into the stock market. The advantage is a property of the attention that active monitoring brings, not of ownership as such: the bundling channel lowers the cost of tracking aggregate signals only for households that engage with their holdings, and ownership without engagement, as in default-enrolled or passively held accounts, carries none of it. The result showing how within-stockholder effect steepens with portfolio size rather than holding flat points to a graded attention margin, plausibly stakes-driven (Sims, 2003), rather than the binary fact of ownership. The same logic bounds the policy reading: since attention is allocated to stakes, the households where ownership and attention most plainly come apart are low-stakes, unengaged holders, small default-enrolled balances, which is precisely the margin at which broader participation would operate.

Anchoring expectations during a prolonged surge therefore depends less on the level of the headline figure than on its reach. The relevant lever is not broader participation, which would expand ownership without necessarily expanding attention, but the delivery of a credible aggregate-price signal to the households most prone to disanchoring, those who form beliefs from experienced prices. This is a segment conventional central-bank communication has had limited success in reaching (Coibion et al., 2022). Active monitoring is one channel that already does this incidentally, the policy problem is to reproduce its effect for the households that do not monitor at all.

Two limitations bound the interpretation. First, the analysis conditions on observed stockholder status and does not model participation. Selection into ownership is almost certainly heterogeneous, and although the 2008 placebo, the larger low-income effect, and the sensitivity analysis all cut against a selection-driven reading, the paper does not estimate the participation margin. Second, the structural exercise is a calibrated, corroborative complement to the reduced form rather than the headline result; its estimates fit four group means and, while the two margins are bootstrap-significant, the magnitudes are identified only moderately. What the data establish is sharp: a belief cleavage long treated as a fixed feature of who people are is in fact state-dependent and it opens only when an inflation episode lasts long enough for experienced-price bias to accumulate. The duration of inflation, not its level, is what divides expectations.

Appendix

2.A Model Derivations

This appendix derives the results stated in Section 2.3: the attention loss (2.3.4), the optimal weight (2.3.5) and its comparative statics, the bundling ranking $m_S^* > m_N^*$, the group means and gap (2.3.6), the duration decomposition, and the tail-compression prediction (2.8.3). Throughout, a household in group $g \in \{S, N\}$ has attention cost κ_g, and the signal noises $\varepsilon^F, \varepsilon^E$ are independent, mean zero, with variances σ_F^2, σ_E^2.

2.A.1 Forecast error and effective signal noise

Substituting the signals (2.3.1) into the forecast rule (2.3.2), the forecast error decomposes into a deterministic bias term and two noise terms,

$$E_i - \pi \ = \ (1-m)\, b_\tau \ + \ m\, \varepsilon_i^F \ + \ (1-m)\, \varepsilon_i^E. \tag{2.A.1}$$

The forecast is biased by $(1-m)b_\tau$: attending more to the financial signal (higher m) carries less of the experienced-price bias. Squaring (2.A.1) and taking expectations, the cross terms vanish because the noises are independent and mean zero and b_τ is deterministic, leaving the mean squared error

$$\mathbb{E}\big[(E_i - \pi)^2\big] \ = \ (1-m)^2 b_\tau^2 + m^2\sigma_F^2 + (1-m)^2\sigma_E^2 \ = \ (1-m)^2\Sigma_{E,\tau}^2 + m^2\sigma_F^2, \tag{2.A.2}$$

where $\Sigma_{E,\tau}^2 \equiv \sigma_E^2 + b_\tau^2$. The bias enters the loss exactly as extra variance: for forecasting, a more biased experienced signal is a noisier one. Adding the quadratic attention cost $\frac{\kappa_g}{2}m^2$ yields the objective (2.3.4).

2.A.2 Optimal attention weight

Write the objective as $O(m) = (1-m)^2\Sigma^2_{E,\tau} + m^2\sigma^2_F + \frac{\kappa_g}{2}m^2$. Its first derivative is

$$O'(m) \;=\; -2(1-m)\,\Sigma^2_{E,\tau} + 2m\,\sigma^2_F + \kappa_g\,m \;=\; 0, \tag{2.A.3}$$

which rearranges to $\Sigma^2_{E,\tau} = m\big(\sigma^2_F + \Sigma^2_{E,\tau} + \kappa_g/2\big)$ and hence to the optimal weight (2.3.5),

$$m^*_{g,\tau} \;=\; \frac{\Sigma^2_{E,\tau}}{\sigma^2_F + \Sigma^2_{E,\tau} + \kappa_g/2}. \tag{2.A.4}$$

The second derivative $O''(m) = 2\Sigma^2_{E,\tau} + 2\sigma^2_F + \kappa_g > 0$, so O is strictly convex and (2.A.4) is its unique minimiser. Since the numerator is positive and strictly below the denominator, $0 < m^*_{g,\tau} < 1$ for every finite κ_g. Two limits bound it: as $\kappa_g \to \infty$, $m^*_{g,\tau} \to 0$—the non-stockholder corner, where attention is too costly to use the financial signal at all—and as $\kappa_g \to 0$, $m^*_{g,\tau} \to \Sigma^2_{E,\tau}/(\sigma^2_F + \Sigma^2_{E,\tau})$, the precision weight of costless Bayesian updating.

2.A.3 Comparative statics and the bundling ranking

Let $D_{g,\tau} \equiv \sigma^2_F + \Sigma^2_{E,\tau} + \kappa_g/2$ be the denominator of (2.A.4). Differentiating,

$$\frac{\partial m^*_{g,\tau}}{\partial \kappa_g} = -\frac{\Sigma^2_{E,\tau}}{2\,D^2_{g,\tau}} < 0, \qquad \frac{\partial m^*_{g,\tau}}{\partial \Sigma^2_{E,\tau}} = \frac{\sigma^2_F + \kappa_g/2}{D^2_{g,\tau}} > 0. \tag{2.A.5}$$

Attention falls with its cost and rises as the experienced signal deteriorates. The bundling discount (2.3.3), $\kappa_S = \kappa_N - \delta$ with $\delta > 0$, then gives $m^*_{S,\tau} > m^*_{N,\tau}$ at every date, because m^* is decreasing in κ.

2.A.4 Group means and the cross-sectional gap

Taking expectations of (2.3.2) within a group (the noises are mean zero) gives the group mean forecast $\mu_{g,\tau} = \pi + (1 - m^*_{g,\tau})\,b_\tau$. The cross-sectional gap is therefore

$$\Delta_\tau \;\equiv\; \mu_{N,\tau} - \mu_{S,\tau} \;=\; \big(m^*_{S,\tau} - m^*_{N,\tau}\big)\,b_\tau, \tag{2.A.6}$$

which at the non-stockholder corner $m^*_{N,\tau} \approx 0$ collapses to $\Delta_\tau = m^*_{S,\tau}\, b_\tau$, equation (2.3.6): the gap is the product of stockholder attention and the experienced-price bias.

2.A.5 Duration: the two margins

Let the bias grow with cumulative salience, $b_\tau = b_0 + \gamma S_\tau$ (equation (2.3.7) at the group average), so $\mathrm{d}b_\tau/\mathrm{d}S_\tau = \gamma$. Differentiating the corner gap $\Delta_\tau = m^*_{S,\tau} b_\tau$ and applying the chain rule through $\Sigma^2_{E,\tau} = \sigma^2_E + b^2_\tau$ (so $\mathrm{d}\Sigma^2_{E,\tau}/\mathrm{d}S_\tau = 2b_\tau\gamma$),

$$\frac{\mathrm{d}\Delta_\tau}{\mathrm{d}S_\tau} = \underbrace{\frac{\partial m^*_{S,\tau}}{\partial \Sigma^2_{E,\tau}} 2b_\tau\gamma \cdot b_\tau}_{\text{intensive margin } (m_S\uparrow)} + \underbrace{m^*_{S,\tau}\,\gamma}_{\text{amplification } (b\uparrow)} = \gamma\left[\frac{2b^2_\tau\big(\sigma^2_F + \kappa_S/2\big)}{D^2_{S,\tau}} + m^*_{S,\tau}\right] > 0. \qquad (2.A.7)$$

Both terms are positive: salience widens the gap by raising stockholder attention (the first term, operating through $\Sigma^2_{E,\tau}$) and by amplifying the bias the gap multiplies (the second). The non-stockholder weight remains at its corner throughout, so no term reflects entry—the widening is intensive-margin attention plus bias amplification, never composition.

2.A.6 Tail compression

Now let attention be common within a group, m_g, while households draw idiosyncratic experienced signals: write household i's experienced signal as π^E_i (its own basket inflation) and let π^F be the common realised financial signal. The forecast is $E_i = m_g\pi^F + (1 - m_g)(\pi^E_i + b_\tau)$. For non-stockholders $m_N = 0$, so $E^N_i = \pi^E_i + b_\tau$: the observed non-stockholder distribution *is* the latent experienced-signal-plus-bias distribution. Stockholders face the same latent distribution but compress it, $E^S_i = m_S\pi^F + (1 - m_S)\, E^N_i$ in distribution. Because the map $x \mapsto m_S\pi^F + (1 - m_S)x$ is increasing and affine (as $0 < 1 - m_S < 1$), quantiles transform equivariantly,

$$Q_\tau\big(E^S\big) \;=\; m_S\,\pi^F + (1 - m_S)\, Q_\tau\big(E^N\big), \qquad (2.A.8)$$

which is (2.8.3). The quantile gap is therefore

$$Q_\tau(E^N) - Q_\tau(E^S) \; = \; m_S\big(Q_\tau(E^N) - \pi^F\big), \tag{2.A.9}$$

which increases in τ wherever $Q_\tau(E^N) > \pi^F$: the gap is widest in the upper tail and vanishes where non-stockholder beliefs meet the financial anchor. As the episode lengthens, m_S rises and the E^N distribution shifts up and spreads through b_τ, so the upper-tail gap widens most—the quantile difference-in-differences of Section 2.8.

2.A.7 Heterogeneity in γ

Finally restore individual bias sensitivities, $b_{i,\tau} = b_0 + \gamma_i S_\tau$. A household's contribution to the gap scales with its bias, and from (2.A.6) the marginal effect of sensitivity is $\partial\Delta_{i,\tau}/\partial\gamma_i = m_S^* S_\tau > 0$: households with greater exposure to salient prices (higher γ_i) accumulate bias faster and show larger gaps. Because γ_i reflects price *exposure* rather than analytical capacity, it should rise with shopping frequency and food–energy budget shares (lower-income households, women) but not with education or age—Prediction 3.

2.B Full Specification Coefficients

Table 10: Full coefficients, headline specification (+Demog.×Post)

	Coefficient
Stock × Post	-1.702^{***}
	(0.190)
Stockholder	-0.478^{***}
	(0.028)
Female	0.594^{***}
	(0.024)
Female × Post	0.996^{***}
	(0.195)
Age below 45	-0.165^{***}
	(0.026)
Age below 45 × Post	0.236^{*}
	(0.137)
College graduate	-0.404^{***}
	(0.027)
College × Post	-1.195^{***}
	(0.136)
High income	-0.465^{***}
	(0.024)
High income × Post	-0.571^{***}
	(0.140)
Constant	5.427^{***}
	(0.174)
Survey-year fixed effects	Yes
Observations	166,125

Notes: Coefficients from the headline +Demog.×Post specification (the fourth column of Table 2). Dependent variable: cleaned one-year-ahead inflation expectation. Survey-weighted estimates; standard errors clustered by survey month in parentheses. Survey-year fixed effects (1997–2025) are included but suppressed. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

2.C Heterogeneity: Full Coefficients

Table 11: Heterogeneity triple-difference estimates, full coefficients

	Education	Income	Gender	Age
Stockholder	-0.650^{***}	-0.602^{***}	-0.773^{***}	-0.799^{***}
	(0.034)	(0.032)	(0.036)	(0.036)
Stock × Post	-2.052^{***}	-2.320^{***}	-1.787^{***}	-2.543^{***}
	(0.247)	(0.240)	(0.193)	(0.232)
College graduate	-0.507^{***}			
	(0.046)			
College × Post	-1.511^{***}			
	(0.274)			
Stock × College	−0.031			
	(0.052)			
Stock × Post × College	+0.358			
	(0.326)			
High income		-0.724^{***}		
		(0.050)		
High income × Post		-1.829^{***}		
		(0.442)		
Stock × High income		$+0.120^{**}$		
		(0.058)		
Stock × Post × High income		$+1.056^{**}$		
		(0.415)		
Female			$+0.658^{***}$	
			(0.044)	
Female × Post			$+1.968^{***}$	
			(0.314)	
Stock × Female			−0.034	
			(0.052)	
Stock × Post × Female			-1.419^{***}	
			(0.302)	
Young (<45)				-0.171^{***}
				(0.050)
Young × Post				+0.035
				(0.287)
Stock × Young				−0.086
				(0.058)
Stock × Post × Young				+0.184
				(0.297)
Observations	166,125	166,125	166,125	166,125
Survey-year FE	Yes	Yes	Yes	Yes

Notes: Each column is a separate triple-difference regression adding the full set of two-way and three-way interactions of *Stock*, *Post*, and the demographic characteristic to the headline specification; the coefficient of interest is *Stock* × *Post*× characteristic. Dependent variable: one-year-ahead inflation expectation. Survey-weighted; standard errors clustered by survey month in parentheses. $^{**}p < 0.05$, $^{***}p < 0.01$.

2.D Forecast Accuracy: Full Coefficients

Table 12: Forecast accuracy and efficiency, full coefficients

	Signed bias	Efficiency	Abs. error
Stock × Post	-1.467^{***}	$+0.014$	-1.390^{***}
	(0.189)	(0.011)	(0.178)
Stockholder	-0.493^{***}	-0.014^{**}	-0.470^{***}
	(0.028)	(0.006)	(0.024)
Expectation (E_π)		$+1.002^{***}$	
		(0.001)	
Female	$+0.592^{***}$	-0.003	$+0.448^{***}$
	(0.024)	(0.003)	(0.020)
Female × Post	$+0.689^{***}$	-0.003	$+0.896^{***}$
	(0.187)	(0.006)	(0.159)
Young (<45)	-0.169^{***}	-0.003	-0.071^{***}
	(0.026)	(0.003)	(0.021)
Young × Post	$+0.194$	$+0.003$	$+0.007$
	(0.150)	(0.006)	(0.098)
College graduate	-0.400^{***}	$+0.005$	-0.361^{***}
	(0.028)	(0.005)	(0.022)
College × Post	-1.301^{***}	-0.019^{**}	-0.898^{***}
	(0.136)	(0.009)	(0.131)
High income	-0.457^{***}	$+0.009^{*}$	-0.299^{***}
	(0.025)	(0.005)	(0.022)
High income × Post	-0.519^{***}	$+0.004$	-0.554^{***}
	(0.161)	(0.010)	(0.153)
Constant	$+0.289^{***}$	-5.149^{***}	$+3.846^{***}$
	(0.090)	(0.175)	(0.168)
Observations	159,865	159,865	159,865
Survey-year FE	Yes	Yes	Yes

Notes: Full coefficients for the accuracy regressions summarised in Table 3. Dependent variable: signed forecast error, i.e. expectation minus realised inflation (cols. 1–2) and absolute forecast error (col. 3); the efficiency column adds the expectation E_π as a regressor. Survey-weighted; standard errors clustered by survey month in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

2.E Event-Study Coefficients

Table 13: Event-study coefficients (Stock × year), 2019 base

Year	Coefficient	SE
2007	-0.359^{**}	0.159
2008	-0.152	0.167
2009	-0.171	0.197
2010	-0.487^{***}	0.120
2011	-0.068	0.167
2012	-0.107	0.143
2013	-0.164	0.137
2014	-0.303^{*}	0.183
2015	-0.154	0.139
2016	-0.302^{***}	0.116
2017	-0.126	0.123
2018	-0.054	0.135
2019	— (base)	—
2020	$+0.016$	0.152
2021	$+0.343$	0.238
2022	-1.322^{***}	0.383
2023	-2.316^{***}	0.286
2024	-2.382^{***}	0.302

Notes: Coefficients on *Stock* × year interactions from the difference-in-differences specification, normalised to the 2019 base year; these are the estimates plotted in Figure 2. Survey-weighted; standard errors clustered by survey month. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

2.F Rejecting Extensive-Margin Entry: Supplementary Tests

Section 2.8 reports that the moment-matching corner ($m_N = 0$) and the direct informed-share heterogeneity test both reject extensive-margin entry. Two further tests, summarised here, corroborate that conclusion.

A structural mixture model. Rather than fixing non-stockholders at the corner, this specification lets an unobserved *informed share* among non-stockholders, households that weight the financial signal, vary freely across four sub-periods (pre-2021, 2021, 2022, 2023–25), and estimates that share by maximum likelihood from the group's expectation distribution. Entry predicts a share that climbs monotonically with duration. With food-and-energy as the experienced signal, the estimated share is instead non-monotone: 58, 79, 68, and 46 percent across the four windows, rising then falling rather than accumulating. The accompanying bias estimates swing from near zero to 17 pp, and their bootstrap confidence intervals exclude their own point estimates, the signature of a label-switching, misspecified mixture rather than a stable entry process. A headline-CPI variant produces a V-shaped share (88, 30, 34, 67 percent) with similarly unstable bias, confirming that the non-monotonicity is not an artifact of the signal choice.

Alternative informed-band definitions. The direct heterogeneity test classifies a household as 'informed' when its expectation falls within a fixed band of the professional signal. To show the result is not an artifact of that band, the test is repeated for four band widths and three demographic splits (income, gender, education). In every band the informed share falls post-2021, and the only difference-in-differences interactions that reach significance run the *wrong* way for entry. Women's informed share falls *more*, not less. No band or split produces the rising, duration-increasing informed share that entry requires.

Together with the moment-matching corner and the baseline informed-share test in the main text, these results leave no specification under which non-stockholders enter the informed group as the surge persists.

References

Ahn, H. J., Xie, S. 2024. Stock-driven household attention.Technical report, SSRN Working Paper No. 4902195, Available at SSRN: https://ssrn.com/abstract=4902195.

Andrade, P., Le Bihan, H. 2013. Inattentive professional forecasters. Journal of Monetary Economics. 60 (8), 967–982.

Andre, P., Pizzinelli, C., Roth, C., Wohlfart, J. 2022. Subjective models of the macroeconomy: Evidence from experts and representative samples. Review of Economic Studies. 89 (6), 2958–2991.

Armantier, O., Bruine de Bruin, W., Topa, G., van der Klaauw, W., Zafar, B. 2015. Inflation expectations and behavior: Do survey respondents act on their beliefs? International Economic Review. 56 (2), 505–536.

Bachmann, R., Berg, T. O., Sims, E. R. 2015. Inflation expectations and readiness to spend: Cross-sectional evidence. American Economic Journal: Economic Policy. 7 (1), 1–35.

Ball, L., Leigh, D., Mishra, P. 2022. Understanding US inflation during the Covid era. Brookings Papers on Economic Activity. Fall, 1–83.

Bernanke, B. S., Blanchard, O. 2023. What caused the US pandemic-era inflation? Brookings Papers on Economic Activity. Spring, 1–86.

Bordalo, P., Gennaioli, N., Ma, Y., Shleifer, A. 2020. Overreaction in macroeconomic expectations. American Economic Review. 110 (9), 2748–2782.

Callaway, B., Sant'Anna, P. H. C. 2021. Difference-in-differences with multiple time periods. Journal of Econometrics. 225 (2), 200–230.

Calvet, L. E., Campbell, J. Y., Sodini, P. 2007. Down or out: Assessing the welfare costs of household investment mistakes. Journal of Political Economy. 115 (5), 707–747.

Carroll, C. D. 2003. Macroeconomic expectations of households and professional forecasters. Quarterly Journal of Economics. 118 (1), 269–298.

Cavallo, A., Cruces, G., Perez-Truglia, R. 2017. Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. American Economic Journal: Macroeconomics. 9 (3), 1–35.

de Chaisemartin, C., D'Haultfœuille, X. 2020. Two-way fixed effects estimators with heterogeneous treatment effects. American Economic Review. 110 (9), 2964–2996.

Coibion, O., Gorodnichenko, Y. 2012. What can survey forecasts tell us about information rigidities? Journal of Political Economy. 120 (1), 116–159.

Coibion, O., Gorodnichenko, Y. 2015. Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review. 105 (8), 2644–2678.

Coibion, O., Gorodnichenko, Y., Kamdar, R. 2018. The formation of expectations, inflation, and the phillips curve. Journal of Economic Literature. 56 (4), 1447–1491.

Coibion, O., Gorodnichenko, Y., Weber, M. 2022. Monetary policy communications and their effects on household inflation expectations. Journal of Political Economy. 130 (6), 1537–1584.

Da, Z., Engelberg, J., Gao, P. 2011. In search of attention. Journal of Finance. 66 (5), 1461–1499.

D'Acunto, F., Malmendier, U., Ospina, J., Weber, M. 2021a. Exposure to grocery prices and inflation expectations. Journal of Political Economy. 129 (5), 1615–1639.

D'Acunto, F., Malmendier, U., Weber, M. 2021b. Gender roles produce divergent economic expectations. Proceedings of the National Academy of Sciences. 118 (21), e2008534118.

Das, S., Kuhnen, C. M., Nagel, S. 2020. Socioeconomic status and macroeconomic expectations. Review of Financial Studies. 33 (1), 395–432.

Engelberg, J., Manski, C. F., Williams, J. 2009. Comparing the point predictions and subjective probability distributions of professional forecasters. Journal of Business & Economic Statistics. 27 (1), 30–41.

Gabaix, X. 2014. A sparsity-based model of bounded rationality. Quarterly Journal of Economics. 129 (4), 1661–1710.

Guiso, L., Sodini, P. 2013. Household finance: An emerging field. In Constantinides, G. M., Harris, M., Stulz, R. M. eds. Handbook of the Economics of Finance. 2B, 1397–1532, Elsevier.

Haliassos, M., Bertaut, C. C. 1995. Why do so few hold stocks? Economic Journal. 105 (432), 1110–1129.

Jeong, D., Ma, S., Yang, C. 2025. Household attention and the transmission of monetary policy.Technical report, Finance and Economics Discussion Series 2025-084, Board of Governors of the Federal Reserve System.

Kacperczyk, M., Van Nieuwerburgh, S., Veldkamp, L. 2016. A rational theory of mutual funds' attention allocation. Econometrica. 84 (2), 571–626.

Koenker, R., Bassett, G. 1978. Regression quantiles. Econometrica. 46 (1), 33–55.

Kuchler, T., Zafar, B. 2019. Personal experiences and expectations about aggregate outcomes. Journal of Finance. 74 (5), 2491–2542.

Lamla, M. J., Lein, S. M. 2014. The role of media for consumers' inflation expectation formation. Journal of Economic Behavior & Organization. 106, 62–77.

Lusardi, A., Mitchell, O. S. 2014. The economic importance of financial literacy: Theory and evidence. Journal of Economic Literature. 52 (1), 5–44.

Maćkowiak, B., Matějka, F., Wiederholt, M. 2023. Rational inattention: A review. Journal of Economic Literature. 61 (1), 226–273.

Maćkowiak, B., Wiederholt, M. 2009. Optimal sticky prices under rational inattention. American Economic Review. 99 (3), 769–803.

Malmendier, U., Nagel, S. 2016. Learning from inflation experiences. Quarterly Journal of Economics. 131 (1), 53–87.

Mankiw, N. G., Reis, R. 2002. Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. Quarterly Journal of Economics. 117 (4), 1295–1328.

Mankiw, N. G., Reis, R., Wolfers, J. 2003. Disagreement about inflation expectations. In Gertler, M., Rogoff, K. eds. NBER Macroeconomics Annual 2003. 18, 209–248, MIT Press.

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

Manski, C. F. 2004. Measuring expectations. Econometrica. 72 (5), 1329–1376.

Patton, A. J., Timmermann, A. 2010. Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion in forecasts. Journal of Monetary Economics. 57 (7), 803–820.

Peress, J. 2014. The media and the diffusion of information in financial markets: Evidence from newspaper strikes. Journal of Finance. 69 (5), 2007–2043.

Rambachan, A., Roth, J. 2023. A more credible approach to parallel trends. Review of Economic Studies. 90 (5), 2555–2591.

Reis, R. 2006. Inattentive consumers. Journal of Monetary Economics. 53 (8), 1761–1800.

van Rooij, M., Lusardi, A., Alessie, R. 2011. Financial literacy and stock market participation. Journal of Financial Economics. 101 (2), 449–472.

Roth, C., Wohlfart, J. 2020. How do expectations about the macroeconomy affect personal expectations and behavior? Review of Economics and Statistics. 102 (4), 731–748.

Sicherman, N., Loewenstein, G., Seppi, D. J., Utkus, S. P. 2016. Financial attention. Review of Financial Studies. 29 (4), 863–897.

Sims, C. A. 2003. Implications of rational inattention. Journal of Monetary Economics. 50 (3), 665–690.

Vissing-Jorgensen, A. 2003. Perspectives on behavioral finance: Does "irrationality" disappear with wealth? evidence from expectations and actions. In Gertler, M., Rogoff, K. eds. NBER Macroeconomics Annual 2003. 18, 139–194, MIT Press.

Weber, M., D'Acunto, F., Gorodnichenko, Y., Coibion, O. 2022. The subjective inflation expectations of households and firms: Measurement, determinants, and implications. Journal of Economic Perspectives. 36 (3), 157–184.