Omid M. Ardakani*

Central Bank Communication and Tail Risk in Inflation Expectations

https://doi.org/10.1515/snde-2026-0060
Received May 7, 2026; accepted August 13, 2026; published online September 1, 2026

Abstract: This paper integrates extreme value theory, rational inattention, and social learning into a framework for inflation expectations under heavy-tailed signals. The framework delivers a tail-biased worst-case forecast, entropy-dispersion tail dependence that rises in the tail index and falls in attention capacity, a globally attracting social learning steady state, and an optimal heavy-tailed signal. In quarterly United States data from 2000 to 2025, headline inflation has a heavy upper tail, Federal Reserve statement entropy explains roughly 30 % of the variation in household survey forecast errors, and the slope of forecast bias on tail-loss covariance flattens after the 2012 inflation target adoption. The structural levers are communication entropy and the inflation tail index.

Keywords: tail risk; inflation expectations; central bank communication; rational inattention; extreme value theory; extremal dependence

JEL Classification: C14; D84; E52; E58

1 Introduction

Modern monetary policy operates through the expectations channel (Woodford 2003), yet extreme events such as the 2008 financial crisis, the 2020 pandemic, and the 2021–2022 inflation surge produced sharp dispersion in inflation expectations across households and firms (Coibion et al. 2020; Kozlowski 2024). Two assumptions in the standard analysis fail under tail events: Gaussian information structures that understate tail probabilities (Embrechts et al. 1997), and homogeneous rational agents that ignore heterogeneous attention (Sims 2003; Maćkowiak et al. 2023). The European Central Bank (2024) Financial Stability Review and the Bank of England (2022) stress testing incorporate tail scenarios with extremal dependence, confirming the practical relevance of these features for policy.

This paper integrates three theoretical pillars. The first is extreme value theory for max-stable policy signals (Resnick 2007; de Haan and Ferreira 2006); rational inattention with mutual information bounds (Sims 2003; Caplin et al. 2022); and adaptive social learning for heterogeneous belief updating (Evans and Honkapohja 2001; Branch and Evans 2017). The analysis yields four structural results. *First*, under Kullback-Leibler ambiguity and asymmetric tail-loss aversion, the worst-case forecast tilts upward relative to the physical mean by an amount proportional to γ_i/θ_i, where γ_i is tail-loss sensitivity and θ_i is the multiplier on the model uncertainty constraint. *Second*, the upper-tail dependence coefficient χ between signal entropy and expectation dispersion satisfies $\partial\chi/\partial\xi > 0$ and $\partial\chi/\partial\kappa < 0$, where ξ is the inflation tail index and κ is attention capacity. *Third*, social learning admits a globally attracting steady state in which the population mean converges to the individual level worst-case bias. *Fourth*, the optimal central bank signal under prescribed dispersion and second

***Corresponding author: Omid M. Ardakani**, Department of Economics, Parker College of Business, Georgia Southern University, 11935 Abercorn Street, Savannah, GA 31419, USA, E-mail: oardakani@georgiasouthern.edu. https://orcid.org/0000-0003-0486-3359

moment constraints belongs to the Fréchet location-scale family, with shape identified by an extreme value statistic.

This paper makes three contributions. *First*, the chi-comparative-statics result is the principal theoretical innovation. Signal entropy and expectation dispersion are codetermined at the joint extreme by the tail index and attention capacity, recovering Gaussian rational expectations as the special case $\xi \to 0$ and $\kappa \to \infty$. *Second*, the empirical analysis documents that Federal Reserve statement entropy is a substantial conditional predictor of household survey forecast error magnitudes, robust across specifications, with magnitude in the same order as the framework's $1/(1 + \xi)$ asymptotic prediction. *Third*, the optimal signal characterization provides a structural target for quantile-based forward guidance. The optimal signal is itself heavy-tailed, with shape pinned to a measurable extreme value statistic of inflation deviations.

The paper builds on three literatures. Extreme value theory provides limit distributions for maxima and exceedances (Fisher and Tippett 1928; Gnedenko 1943; Pickands 1975; Smith 1987) and has been applied in finance (Embrechts et al. 1997, 2002; Gabaix 2016) but with limited integration into monetary economics; recent tail-risk work in Ardakani and Ajina (2024) and Kozlowski (2024) motivates the present treatment. Rational inattention (Sims 2003, 2006) establishes that finite channel capacity produces noisy forecasts; recent extensions explore non-Gaussian settings where attention constraints amplify distributional asymmetries (Maćkowiak et al. 2023; Maćkowiak and Wiederholt 2025). Adaptive learning (Evans and Honkapohja 2001; Branch and Evans 2017; Eusepi and Preston 2018; Carvalho et al. 2023) relaxes rational expectations through belief updating, with social learning aggregation (Banerjee 1992; Banerjee et al. 2013; Orléan 1995; Arifovic et al. 2025) producing the population level dynamics that, in our framework, admit a globally attracting fixed point at the worst-case bias.

Empirically, quarterly United States data from 2000 to 2025 deliver three findings. *First*, headline inflation exhibits a heavy upper tail, with a Hill shape estimate $\hat{\xi} = 0.45$ that is robust across subsamples and corroborated by mean excess and quantile diagnostics. *Second*, Federal Reserve statement entropy is the strongest conditional predictor of absolute four-quarter household survey forecast errors among the indicators we examine, with $R^2 = 0.30$ from entropy alone and $\hat{\beta}_{\mathcal{E}} = 0.479$ (heteroskedasticity- and autocorrelation-consistent, or HAC, standard error 0.105), a magnitude statistically below but in the same order of magnitude as $1/(1 + \hat{\xi}) = 0.69$ asymptotic prediction. *Third*, the cross-section of forecast bias against tail-loss covariance flattens after the 2012 inflation target adoption. The visual pre-target slope is +1.50, the post-target slope is −0.01, and the simple-average pre-target slope across twelve robustness cells is +1.40 versus a post-target average of −0.03. The pooled Wald test of slope equality fails to reject under crisis-conditional partialling; we treat the structural break as suggestive but not identified.

The paper acknowledges an identification limitation upfront. The aggregate empirical design cannot separately identify information processing heterogeneity from incentive heterogeneity as the source of tail dispersion, and Section 3 sketches a Survey of Professional Forecasters panel design that would resolve the distinction. The remainder of the paper proceeds as follows. Section 2 develops the theoretical framework, Section 3 presents the empirical analysis, Section 4 discusses the findings, and Section 5 concludes.

2 Framework

This section integrates three literatures. The first is extreme value theory (Fisher and Tippett 1928; Gnedenko 1943; Embrechts et al. 1997; Resnick 2007; Coles 2001; de Haan and Ferreira 2006), rational inattention (Sims 2003, 2006; Maćkowiak and Wiederholt 2015; Maćkowiak et al. 2023; Matějka and McKay 2015; Steiner et al. 2017; Caplin et al. 2022), and robust control with variational preferences (Hansen and Sargent 2001, 2008; Anderson et al. 2003; Maccheroni et al. 2006; Strzalecki 2011). Section 2.1 derives optimal expectations under tail-risk ambiguity, Section 2.2 formalizes extremal dependence between signal entropy and expectation dispersion, Section 2.3 characterizes the steady-state deviation under social transmission, and Section 2.4 turns to optimal crisis communication.

Definition 1. A policy signal process $\{\mathcal{S}_t\}_{t\geq 0}$ is *max-stable* if there exist sequences $\{a_T > 0\}$ and $\{b_T \in \mathbb{R}\}$ such that, for all x in the support of G_ξ,

$$\lim_{T\to\infty}\mathbb{P}\left(\frac{\max\limits_{1\leq t\leq T}\mathcal{S}_t - b_T}{a_T} \leq x\right) = G_\xi(x), \qquad G_\xi(x) = \exp\left\{-(1+\xi x)_+^{-1/\xi}\right\}, \tag{1}$$

with shape $\xi \in \mathbb{R}$.

The Fisher–Tippett–Gnedenko theorem (Fisher and Tippett 1928; Gnedenko 1943) establishes that the generalized extreme value (GEV) distribution is the only nondegenerate limit law for affinely normalized sample maxima of i.i.d. sequences. The shape parameter ξ classifies tail behavior: $\xi > 0$ yields heavy-tailed Fréchet, $\xi = 0$ Gumbel, $\xi < 0$ bounded Weibull. Heavy-tailed dynamics ($\xi > 0$) generate finite moments only up to order $1/\xi$, which is the critical feature for inflation expectation analysis. A tail index $\xi > 1/2$ rules out finite second moments, and the empirically relevant Survey of Professional Forecasters (SPF) density forecasts during 2007–2009 and 2020–2022 yield rolling estimates of $\hat{\xi}$ in the 0.30–0.55 range (Ardakani and Ajina 2024; Kozeniauskas et al. 2018).

Definition 2. Agent $i \in [0,1]$'s perceived signal $\tilde{\mathcal{S}}_t^i$ satisfies the rate constraint

$$I\big(\mathcal{S}_t\,;\,\tilde{\mathcal{S}}_t^i\big) \leq \kappa_i, \qquad t \geq 0, \tag{2}$$

where $\kappa_i > 0$ is attention capacity and $I(\cdot;\cdot)$ is mutual information.

The constraint formalizes information processing limits in the rational inattention tradition (Shannon 1948; Cover and Thomas 1991; Sims 2003, 2006). Mutual information capacity has axiomatic foundations (Caplin et al. 2022) and produces multinomial-logit choice probabilities under quadratic utility (Matějka and McKay 2015), with dynamic extensions developed in Steiner et al. (2017). Maćkowiak and Wiederholt (2025) survey the macroeconomic applications. The synthesis of Definitions 1 and 2 yields two empirical predictions. First, max-stability captures how extreme policy shocks distort expectation anchoring, complementing diagnostic expectations (Bordalo et al. 2018) and disagreement-based (Mankiw et al. 2003; Coibion and Gorodnichenko 2015) approaches. Second, heterogeneous κ_i generates differential perception: high-capacity agents perceive signals with negligible distortion, while low-capacity agents amplify tail distortions (Sims 2003; Maćkowiak and Wiederholt 2015; Ardakani et al. 2025b). If κ_i covaries negatively with income (Carvalho et al. 2016), undifferentiated communication exacerbates expectation inequality (Ardakani and Saenz 2023).

2.1 Extremal Expectation Formation

Let $(\Omega, \mathcal{F}, P)$ be a probability space and $\pi_t : \Omega \to \mathbb{R}$ an $\mathcal{F}_t$-measurable inflation variable.

Definition 3. The set of admissible belief distortions is

$$\mathcal{Q}_{\eta_i} = \left\{Q \ll P:\ \mathbb{E}_Q\big[\pi_t^2\big] < \infty,\ \mathcal{D}_{KL}(Q\,;\,P) \leq \eta_i\right\}, \tag{3}$$

with $\eta_i > 0$ bounding model uncertainty.

Each agent i solves the robust forecast problem

$$\inf_{\hat{\pi}\in\mathbb{R}}\ \sup_{Q\in\mathcal{Q}_{\eta_i}}\ \mathbb{E}_Q\big[(\pi_t - \hat{\pi})^2 + \gamma_i(\pi_t - \mu_t)^+\big], \tag{4}$$

with μ_t a tail-risk threshold and $\gamma_i \geq 0$ tail-loss sensitivity. The formulation belongs to the multiplier preference family of Hansen and Sargent (2001, 2008) and the more general variational preference axiomatization of Maccheroni et al. (2006) and Strzalecki (2011) provides a behavioral foundation specifically for the multiplier representation. Compared to the related risk-sensitive formulation of Anderson et al. (2003), the additive $(\pi_t - \mu_t)^+$ term encodes asymmetric tail aversion in the spirit of prospect-theoretic and disaster-risk frameworks (Gabaix 2019; Bordalo et al. 2018).

Theorem 1. *Suppose $\mathbb{E}_P\left[\exp\left(c\pi_t^2\right)\right] < \infty$ for some $c > 0$, and let $\theta_i > 1/c$ denote the Lagrange multiplier on the KL constraint (decreasing in η_i). Then (4) admits a unique saddle point $\left(Q_i^*, \hat{\pi}_t^{i*}\right)$ with*

(i) *worst-case Radon–Nikodym derivative*

$$\frac{\mathrm{d}Q_i^*}{\mathrm{d}P} = \frac{1}{Z_i\left(\hat{\pi}_t^{i*}\right)} \exp\left\{ \frac{1}{\theta_i}\left[\left(\pi_t - \hat{\pi}_t^{i*}\right)^2 + \gamma_i(\pi_t - \mu_t)^+\right]\right\}, \tag{5}$$

where $Z_i(\hat{\pi}) = \mathbb{E}_P\left[\exp(\theta_i^{-1}\{(\pi_t - \hat{\pi})^2 + \gamma_i(\pi_t - \mu_t)^+\})\right] < \infty$;

(ii) *fixed point characterization $\hat{\pi}_t^{i*} = \mathbb{E}_{Q_i^*}[\pi_t]$;*

(iii) *Limit $Q_i^* \to P$ in total variation and $\hat{\pi}_t^{i*} \to \mathbb{E}_P[\pi_t]$ as $\eta_i \to 0$ ($\theta_i \to \infty$) and $\gamma_i \to 0$.*

Proof. Write $f(\pi, \hat{\pi}) = (\pi - \hat{\pi})^2 + \gamma_i(\pi - \mu_t)^+$.

(i) The mapping $Q \mapsto \mathbb{E}_Q[f]$ is linear, $Q \mapsto \mathcal{D}_{KL}(Q;P)$ is strictly convex, and $\mathcal{Q}_{\eta_i}$ is convex. By the Donsker–Varadhan variational formula (Hansen and Sargent 2008),

$$\sup_{Q\in\mathcal{Q}_{\eta_i}} \mathbb{E}_Q[f] = \inf_{\theta>0}\{\theta\,\eta_i + \theta\,\log \mathbb{E}_P[\exp(f/\theta)]\}. \tag{6}$$

The inner expectation is finite iff $1/\theta < c$ since the linear-in-$\hat{\pi}$ and γ_i terms are dominated by π^2/θ for $|\pi|$ large; the minimizing θ is the Lagrange multiplier $\theta_i = \theta_i(\eta_i) > 1/c$, decreasing in η_i. Standard Gibbs duality yields (5).

(ii) The dual objective $J(\hat{\pi}) := \theta_i\,\eta_i + \theta_i\,\log Z_i(\hat{\pi})$ is continuously differentiable. By the envelope theorem,

$$J'(\hat{\pi}) = \mathbb{E}_{Q_i^*(\hat{\pi})}\left[\partial f/\partial\hat{\pi}\right] = -2\left(\mathbb{E}_{Q_i^*(\hat{\pi})}[\pi_t] - \hat{\pi}\right). \tag{7}$$

The FOC $J'(\hat{\pi}^*) = 0$ gives $\hat{\pi}_t^{i*} = \mathbb{E}_{Q_i^*}[\pi_t]$. Define $\Phi(\hat{\pi}) := \mathbb{E}_{Q_i^*(\hat{\pi})}[\pi_t]$ and $\Psi(\hat{\pi}) = \Phi(\hat{\pi}) - \hat{\pi}$. Differentiating $Q_i^*(\hat{\pi})$ with respect to $\hat{\pi}$ and using $\partial\,\log Z_i/\partial\hat{\pi} = -2(\Phi(\hat{\pi}) - \hat{\pi})/\theta_i$,

$$\frac{\partial Q_i^*(\pi;\hat{\pi})}{\partial\hat{\pi}} = Q_i^*(\pi;\hat{\pi}) \cdot \frac{-2}{\theta_i}\left(\pi - \Phi(\hat{\pi})\right). \tag{8}$$

Substituting (8) into $\Phi'(\hat{\pi}) = \int \pi\,\partial Q_i^*/\partial\hat{\pi}\,\mathrm{d}\pi$,

$$\Phi'(\hat{\pi}) = \frac{-2}{\theta_i}\left[\mathbb{E}_{Q_i^*(\hat{\pi})}\left[\pi_t^2\right] - \Phi(\hat{\pi})^2\right] = -\frac{2}{\theta_i}\,\mathrm{Var}_{Q_i^*(\hat{\pi})}(\pi_t). \tag{9}$$

Hence $\Psi'(\hat{\pi}) = -\frac{2}{\theta_i}\mathrm{Var}_{Q_i^*(\hat{\pi})}(\pi_t) - 1 \le -1$ for all $\hat{\pi}$. So Ψ is continuous, strictly decreasing, and satisfies $\Psi(\hat{\pi}) \le \Psi(0) - \hat{\pi}$ for $\hat{\pi} > 0$ and $\Psi(\hat{\pi}) \ge \Psi(0) - \hat{\pi}$ for $\hat{\pi} < 0$. It therefore crosses zero exactly once.

(iii) As $\theta_i \to \infty$, $\exp(f/\theta_i) \to 1$ uniformly on compacts, so $\mathrm{d}Q_i^*/\mathrm{d}P \to 1$ pointwise; dominated convergence under $\mathbb{E}_P\left[\exp\left(c\pi_t^2\right)\right] < \infty$ extends this to total variation convergence. With $\gamma_i = 0$ the loss is $(\pi - \hat{\pi})^2$ whose unique fixed point is $\mathbb{E}_P[\pi_t]$.

Remark 1. With $\gamma_i = 0$, $\mathrm{d}Q_i^*/\mathrm{d}P \propto \exp((\pi_t - \hat{\pi})^2/\theta_i)$ is robust control distortion of Hansen and Sargent (2001), Q_i^* inflates variance symmetrically. With $\gamma_i > 0$, Q_i^* additionally overweights $\{\pi_t > \mu_t\}$, so $\mathbb{E}_{Q_i^*}[\pi_t] > \mathbb{E}_P[\pi_t]$: the tail-loss-averse forecast is biased upward.

Example 1. Let $\pi_t \sim \mathcal{N}(2,1)$ in percentage points, $\mu_t = 2$, $\theta_i = 3$, $\gamma_i = 2$. Sub-Gaussianity gives $\mathbb{E}_P\left[\exp\left(c\pi_t^2\right)\right] < \infty$ for every $c < 1/(2\sigma^2) = 0.5$, so $1/c$ can be taken arbitrarily close to 2 and $\theta_i = 3$ satisfies integrability. Numerical fixed point iteration yields $\hat{\pi}^{i*} = 2.381$ with $\mathrm{Var}_{Q_i^*}(\pi_t) = 4.78$,

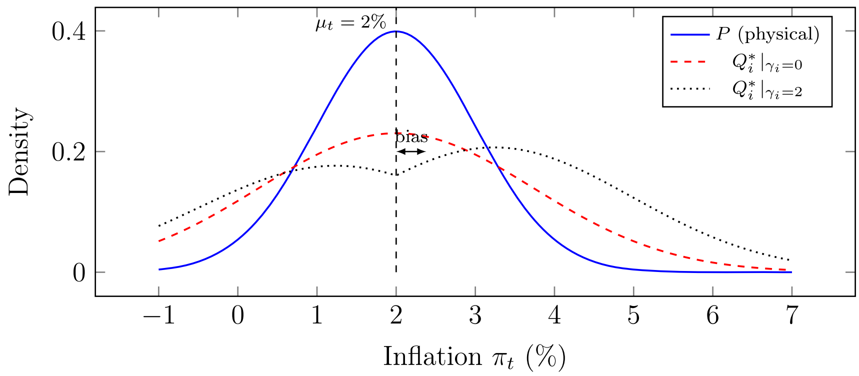

Figure 1: Belief distortion under worst-case tail-loss aversion ($\sigma = 1\%, \theta_i = 3, \gamma_i = 2$). Densities normalized to integrate to one; values from numerical implementation.

so $\Psi'(\hat{\pi}^*) = -(2/3)(4.78) - 1 = -4.19 < 0$, confirming the unique fixed point. The tail-induced bias is $\hat{\pi}^{i*} - \mathbb{E}_P[\pi_t] = 0.381$; setting $\gamma_i = 0$ recovers $\hat{\pi}^{i*} = 2.000$. Figure 1 plots the physical density P against the distorted densities Q_i^* under $\gamma_i = 0$ and $\gamma_i = 2$.

Expanding $\mathrm{d}Q_i^*/\mathrm{d}P$ in $1/\theta_i$ at the fixed point and applying Theorem 1,

$$\delta_t := \mathbb{E}_{Q_i^*}[\pi_t] - \mathbb{E}_P[\pi_t] = \frac{1}{\theta_i}\left[m_3(P) + \gamma_i \operatorname{Cov}_P\big((\pi_t - \mu_t)^+,\, \pi_t\big)\right] + O(\theta_i^{-2}), \tag{10}$$

where $m_3(P) = \mathbb{E}_P\left[(\pi_t - \mathbb{E}_P[\pi_t])^3\right]$ is the third central moment. Under symmetric P (e.g., Gaussian) the cubic term vanishes and the bias is purely tail-loss-driven. For positively skewed P (the empirically relevant regime under $\xi > 0$), the cubic moment contributes additional upward bias even with $\gamma_i = 0$. Response intensity to a tail event scales with γ_i/θ_i at agents whose information constraint binds, $I(\mathcal{S}_t; \tilde{\mathcal{S}}_t^i) \approx \kappa_i$. The formula (10) parallels the diagnostic expectations distortion in Bordalo et al. (2018) but arises endogenously from worst-case ambiguity rather than belief overreaction.

Figure 2 traces the link from policy signals through expectation formation to outcomes.

2.2 Signal Entropy and Tail Dependence

When $\xi > 0$ and agents face information constraints, expectation dispersion $\operatorname{Var}\big(\hat{\pi}_t^i \mid \mathcal{F}_t\big)$ becomes tail-dependent on signal entropy $\mathcal{H}(\mathcal{S}_t)$. I quantify the dependence using the bivariate extreme value framework of Pickands (1981) and Coles (2001); comparable parametric specifications include Hüsler and Reiss (1989)'s Gaussian-derived dependence and Patton (2009)'s dynamic copula models.

Definition 4. For $X := \mathcal{H}(\mathcal{S}_t)$ and $Y := \operatorname{Var}\big(\hat{\pi}_t^i \mid \mathcal{F}_t\big)$ with continuous marginals F_X, F_Y, the upper-tail dependence coefficient is

$$\chi := \lim_{u \to 1^-} \mathbb{P}\big(F_Y(Y) > u \big| F_X(X) > u\big) \in [0, 1], \tag{11}$$

when the limit exists.

Assumption 1. $\{\mathcal{S}_t\}$ is max-stable with $\xi > 0$: the marginal survival function satisfies $\bar{F}_S(s) = s^{-1/\xi}\ell(s)$ for some slowly varying $\ell : (0, \infty) \to (0, \infty)$.

Assumption 2. Expectation dispersion satisfies

$$\operatorname{Var}\big(\hat{\pi}_t^i \mid \mathcal{F}_t\big) = c(\kappa_i)\,[\mathcal{H}(\mathcal{S}_t)]^{\delta}\, e^{\eta_t^i}, \qquad \delta = \frac{\xi}{1+\xi}, \tag{12}$$

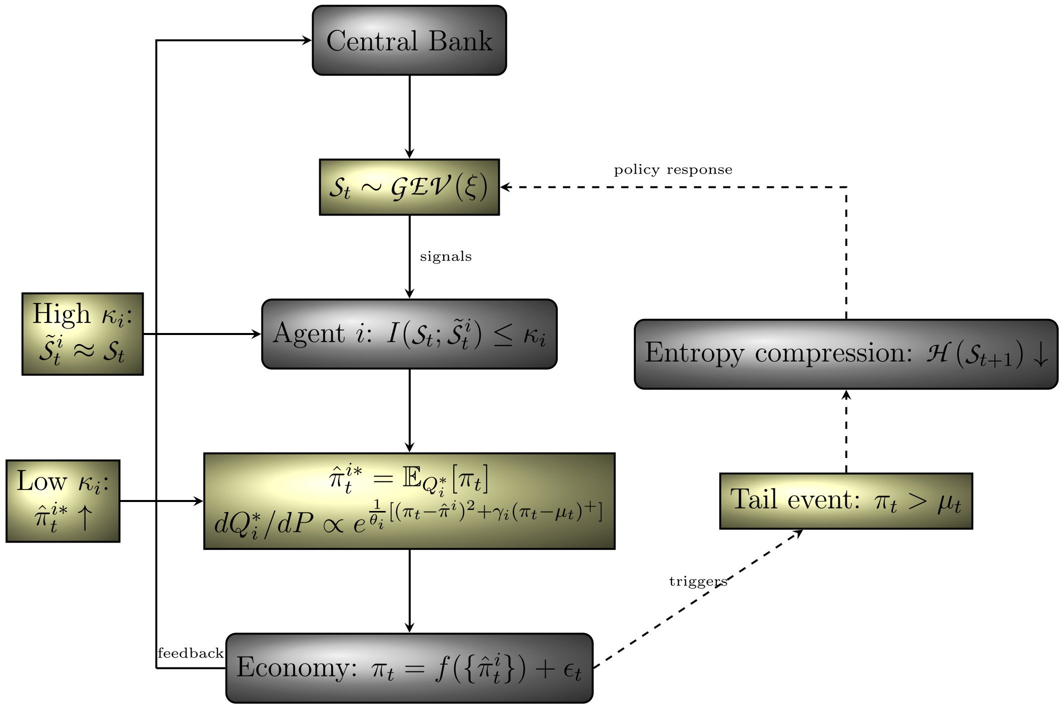

Figure 2: Core mechanisms of tail-risk expectation formation. Solid arrows: signal transmission and economic feedback. Dashed arrows: tail-event-triggered policy response.

where $c : (0, \infty) \to (0, \infty)$ is strictly decreasing with $c(\kappa) \to 0$ as $\kappa \to \infty$, and η_t^i is independent of $\mathcal{S}_t$ with $\mathbb{E}\left[\eta_t^i\right] = 0$, $\mathrm{Var}\left(\eta_t^i\right) = \sigma_\eta^2 \in (0, \infty)$, and finite moment generating function $M_\eta(s) = \mathbb{E}\left[e^{s\eta_t^i}\right]$ in a neighborhood of zero.

The exponent $\delta = \xi/(1+\xi)$ has two motivations. First, it is the natural contraction power for a Shannon channel processing a regularly varying input (Cover and Thomas 1991). Second, the exponent keeps the transformed signal in the regularly varying class while relaxing its moment restriction. If $X \in \mathrm{MDA}(G_\xi)$ has finite moments up to order $1/\xi$, then X^δ has finite moments up to $(1+\xi)/\xi^2$, which exceeds $1/\xi$ for every $\xi > 0$. The multiplicative noise $e^{\eta_t^i}$ rather than additive noise preserves regularly varying tail under aggregation (Ardakani et al. 2025b).

Proposition 1. *Under Assumptions* 1 *and* 2, $(X, Y) = (\mathcal{H}(\mathcal{S}_t), \mathrm{Var}(\hat{\pi}_t^i \mid \mathcal{F}_t))$ *lies in the max domain of attraction of a bivariate extreme value distribution with nondegenerate dependence. Within the symmetric logistic family of* Coles (2001) *parametrized by* $\rho \in (0, 1]$, *the upper-tail dependence coefficient is*

$$\chi(\xi, \kappa_i) = 2 - 2^{\rho(\xi, \kappa_i)}, \tag{13}$$

with calibration

$$\rho(\xi, \kappa_i) = \frac{1}{1+\xi} \cdot \frac{\kappa_i}{\kappa_i + \sigma_\eta^2}. \tag{14}$$

Comparative statics $\partial\chi/\partial\xi > 0$, $\partial\chi/\partial\kappa_i < 0$*; the classical limit* $\xi \to 0$, $\kappa_i \to \infty$ *gives* $\rho \to 1$ *and* $\chi \to 0$ *(tail independence under Gaussian rational expectations).*

Proof. By Assumption 1, $\bar{F}_X(x) = x^{-1/\xi}\ell_X(x)$. From (12), $Y = c(\kappa_i)X^\delta e^{\eta_t^i}$ with $\delta = \xi/(1+\xi) > 0$ and $e^{\eta_t^i}$ admitting all positive moments by the MGF assumption. Breiman's lemma (Embrechts et al. 1997, Lemma 4.4.2) gives

$$\bar{F}_Y(y) \sim c(\kappa_i)^{1/(\xi\delta)} \cdot M_\eta(1/(\xi\delta)) \cdot y^{-1/(\xi\delta)}\ell_Y(y), \tag{15}$$

so Y is regularly varying with tail index $1/(\xi\delta) = (1+\xi)/\xi^2$.

The standardized vector $(\tilde{X}, \tilde{Y}) = \left(-1/\log F_X(X), -1/\log F_Y(Y)\right)$ has unit-Fréchet marginals. Equation (12) gives $\log\, Y = \log\, c(\kappa_i) + \delta\, \log\, X + \eta_t^i$, i.e., $\log Y$ is location-shifted by an independent light-tailed innovation. By Embrechts et al. (1997, Thm. 6.1.14), $(\tilde{X}, \tilde{Y})$ converges weakly to a bivariate extreme value law

$$G(x,y) = \exp\{-V(x,y)\}, \qquad V(x,y) = 2\int_0^1 \max\left(\frac{w}{x}, \frac{1-w}{y}\right) \mathrm{d}H(w), \tag{16}$$

where H is a probability measure on $[0, 1]$ with $\int_0^1 w\, \mathrm{d}H(w) = \int_0^1 (1-w)\, \mathrm{d}H(w) = 1/2$ (Pickands 1981; Coles 2001; de Haan and Ferreira 2006). This normalization yields the unit-Fréchet marginals $V(x,\infty) = (2/x)\int w\, \mathrm{d}H(w) = 1/x$ and $V(\infty, y) = 1/y$. The tail dependence coefficient is $\chi = 2 - V(1,1) = 2 - 2A(1/2)$, where $A(t) = V(1/t, 1/(1-t))/2$ is the Pickands dependence function and $A(1/2) = \int_0^1 \max(w, 1-w)\, \mathrm{d}H(w)$.

Independence of η_t^i from $\mathcal{S}_t$ implies that the joint upper tail of $(\log X, \log Y)$ is governed by a one-parameter exchangeable copula. The symmetric logistic family of Coles (2001) provides an analytically tractable parametrization within this class, with Pickands function

$$A(t) = \left[t^{1/\rho} + (1-t)^{1/\rho}\right]^{\rho}, \qquad \rho \in (0,1], \tag{17}$$

where $\rho = 1$ is independence and $\rho \to 0$ is comonotonicity. Direct evaluation gives $A(1/2) = 2^{\rho-1}$, hence

$$\chi = 2 - 2A(1/2) = 2 - 2^{\rho}. \tag{18}$$

Within the parametric family (17), ρ is identified by matching the model's signal-to-noise ratio. From $\log\, Y = \delta\, \log\, X + \eta_t^i + \text{const}$, the Gaussian-channel rate distortion bound (Cover and Thomas 1991, Thm. 10.3.2) for capacity κ_i implies that the residual fraction of unexplained variance scales as $\sigma_\eta^2/\left(\kappa_i + \sigma_\eta^2\right)$. Combining with the heavy-tail contraction $\delta/\xi = 1/(1+\xi)$ from the power transformation in (12) yields the calibration (14).

$\partial\rho/\partial\xi = -(1+\xi)^{-2}\,\kappa_i/\left(\kappa_i + \sigma_\eta^2\right) < 0$, and $\partial\chi/\partial\rho = -2^{\rho}\log 2 < 0$, so $\partial\chi/\partial\xi > 0$. Similarly $\partial\rho/\partial\kappa_i = (1+\xi)^{-1}\,\sigma_\eta^2\left(\kappa_i + \sigma_\eta^2\right)^{-2} > 0$, so $\partial\chi/\partial\kappa_i < 0$. As $\xi \to 0$ and $\kappa_i \to \infty$, $\rho \to 1$ and $\chi \to 0$.

Remark 2. The symmetric logistic specification (17) is a parsimonious choice within the bivariate EV family. Hüsler and Reiss (1989) derive an alternative family from Gaussian dependence with $\chi_{\mathrm{HR}} = 2(1 - \Phi(\lambda/2))$ for parameter $\lambda > 0$. The Hüsler–Reiss specification produces qualitatively identical comparative statics under analogous calibration of λ.

Example 2. For $\xi = 0.5$ and $\sigma_\eta^2 \to 0$ (perfect attention), $\rho = 2/3$ and $\chi = 2 - 2^{2/3} = 0.413$. For $\xi = 1$ and $\sigma_\eta^2 \to 0$, $\rho = 1/2$ and $\chi = 2 - \sqrt{2} = 0.586$. With $\xi = 0.5$ and $\sigma_\eta^2 = \kappa_i$, $\rho = 1/3$ and $\chi = 2 - 2^{1/3} = 0.740$.

2.3 Social Transmission of Extremes

This section characterizes the steady-state expectation deviation when extreme tail-induced beliefs propagate through a complete social network. The analysis builds on Banerjee (1992) and Bikhchandani et al. (1992) on informational cascades and Acemoglu et al. (2011) on Bayesian network learning, and provides a structural counterpart to the constant-gain instability boundary in adaptive learning (Evans and Honkapohja 2001; Branch and Evans 2017; Eusepi and Preston 2018; Carvalho et al. 2023). The mechanism is distinct from cascade-driven divergence. The dynamics admit a globally attracting fixed point at the worst-case bias δ_t from Theorem 1, and the relevant policy threshold is the parameter regime where this stationary deviation exceeds tolerance.

Let $\phi : \mathbb{R} \to [0, 1)$ be twice differentiable, increasing, with $\phi(0) = 1/2$ and local convexity $\phi''(0) > 0$. The upper bound is approached but never attained, so $\phi(r) < 1$ at every finite r.

Theorem 2. *Let $\hat{\pi}_t = \left(\hat{\pi}_t^i\right)_{i=1}^n$, π_t^* fundamental inflation, and S_t a signal satisfying Definition 1. Assume:*
(A1) *Homogeneous tail sensitivity, $\gamma_i = \gamma$, $\theta_i = \theta$ for all i.*
(A2) *Complete network, agents observe $\bar{\hat{\pi}}_t = n^{-1}\sum_i \hat{\pi}_t^i$.*
(A3) *Updating rule, $\hat{\pi}_{t+1}^i = \phi(r_t)\,\bar{\hat{\pi}}_t + (1 - \phi(r_t))\,\hat{\pi}_t^{i*}$, with $r_t = \left(\hat{\pi}_t^{i*} - \bar{\hat{\pi}}_t\right)\left(2\pi_t^* - \hat{\pi}_t^{i*} - \bar{\hat{\pi}}_t\right)$.*
(A4) $\phi''(0) > 0$.

Define $W_t^i = \hat{\pi}_t^i - \pi_t^*$, $\bar{W}_t = \bar{\hat{\pi}}_t - \pi_t^*$, and $\delta_t = \mathbb{E}_{Q^*}[\pi_t] - \pi_t^* > 0$. Then:
(i) *For all $t \geq 1$, $W_t^i = \bar{W}_t$ for every i, so the dynamics reduce to $\bar{W}_{t+1} = \phi(r_t)\,\bar{W}_t + (1 - \phi(r_t))\,\delta_t$ with $r_t = \bar{W}_t^2 - \delta_t^2$.*
(ii) *$\bar{W}^* = \delta_t$ is the unique fixed point.*
(iii) *The fixed point is locally stable with linearized contraction rate $\phi(0) = 1/2$, independent of $\phi'(0)$ and $\phi''(0)$.*
(iv) *For any initial condition, $\bar{W}_t \to \delta_t$ as $t \to \infty$.*

Proof. (i) Under (A1) the worst-case forecast $\hat{\pi}_t^{i*} = \mathbb{E}_{Q^*}[\pi_t]$ is identical across agents, so $\hat{\pi}_t^{i*} = \pi_t^* + \delta_t$ for every i. Substituting in (A3),

$$\hat{\pi}_{t+1}^i = \phi(r_t)\,\bar{\hat{\pi}}_t + (1 - \phi(r_t))\big(\pi_t^* + \delta_t\big),$$

which has no i-dependence, so $\hat{\pi}_{t+1}^i$ is identical across agents and thereafter $W_t^i = \bar{W}_t$.

(ii) With $r_t = (\delta_t - \bar{W}_t)(-\delta_t - \bar{W}_t) = \bar{W}_t^2 - \delta_t^2$, the fixed point condition $\bar{W} = \phi(r)\bar{W} + (1 - \phi(r))\delta_t$ rearranges to $(1 - \phi(r))(\bar{W} - \delta_t) = 0$. Since $\phi(r) < 1$ at every finite r, $\bar{W} = \delta_t$ is the unique fixed point.

(iii) Let $\varepsilon_t = \bar{W}_t - \delta_t$. Then $r_t = (\delta_t + \varepsilon_t)^2 - \delta_t^2 = 2\delta_t\varepsilon_t + \varepsilon_t^2$, and

$$\varepsilon_{t+1} = \phi(r_t)\,\varepsilon_t. \tag{19}$$

Taylor expansion at $r = 0$ gives $\phi(r) = \frac{1}{2} + \phi'(0)r + \frac{1}{2}\phi''(0)r^2 + O(r^3)$, so

$$\varepsilon_{t+1} = \frac{1}{2}\varepsilon_t + 2\delta_t\,\phi'(0)\,\varepsilon_t^2 + \left[\phi'(0) + 2\delta_t^2\,\phi''(0)\right]\varepsilon_t^3 + O\big(\varepsilon_t^4\big). \tag{20}$$

The linearized rate is $\phi(0) = 1/2 < 1$, independent of $\phi'(0)$ and $\phi''(0)$. The fixed point is locally stable.

(iv) Since ϕ takes values in $[0, 1)$, the recursion (19) implies $|\varepsilon_{t+1}| \leq |\varepsilon_t|$, so the sequence $\{|\varepsilon_t|\}_{t\geq 0}$ is nonincreasing and bounded below by zero. By monotone convergence, $|\varepsilon_t| \to L$ for some $L \geq 0$. Suppose $L > 0$ for contradiction. Since $\phi \geq 0$, ε_{t+1} has the same sign as ε_t, so ε_t converges in sign and $\varepsilon_t \to \pm L$. Then $\bar{W}_t \to \delta_t \pm L$ and $r_t \to r_* = L^2 \pm 2\delta_t L$, a finite limit. By continuity of ϕ, taking $t \to \infty$ in (19) gives $\pm L = \phi(r_*)\,(\pm L)$, hence $\phi(r_*) = 1$. But $\phi(r_*) < 1$ at every finite r_*, a contradiction. Therefore $L = 0$, i.e., $\bar{W}_t \to \delta_t$.

Remark 3. The dynamics in Theorem 2 cannot generate unbounded expectation divergence: $\phi \in [0, 1)$ enforces convex combinations and thus contraction. The convexity $\phi''(0) > 0$ does not produce instability of the fixed point, contrary to a naive Lyapunov-drift heuristic. The role of social transmission is therefore not to amplify deviations but to *lock* expectations to the worst-case bias δ_t generated by individual tail-loss aversion. The "cascade regime" is meaningful as a parameter threshold (Corollary 1), not as a dynamic instability.

Corollary 1. *Conditional on a tail event $\{S_t > s_\tau\}$, the steady-state expectation deviation satisfies $\delta_t \mid \text{tail} > \varepsilon_\delta$ if and only if*

$$\frac{\gamma}{\theta}\,\mathrm{Cov}_{P_\tau}\big((\pi_t - \mu_t)^+,\, \pi_t\big) + \frac{m_3(P_\tau)}{\theta} > \varepsilon_\delta + O(\theta^{-2}), \tag{21}$$

where P_τ is the tail-conditional measure and $m_3(P_\tau)$ is its third central moment.

Proof. Apply (10) to the tail-conditional measure P_τ and use $\delta_t = \bar{W}^*$ from Theorem 2.

Corollary 2. *Let $M_\eta(1) := \mathbb{E}\left[e^{\eta_t^i}\right]$, finite by Assumption 2. For dispersion tolerance $\sigma^2_{\text{crit}} > 0$, the entropy threshold above which expected dispersion exceeds tolerance is*

$$\mathcal{E}_{\text{casc}}\left(\xi, \kappa_i, \sigma^2_{\text{crit}}\right) := \left[\frac{\sigma^2_{\text{crit}}}{c(\kappa_i)\, M_\eta(1)}\right]^{(1+\xi)/\xi}. \tag{22}$$

Under the special case $\eta_t^i \sim \mathcal{N}\left(0, \sigma^2_\eta\right)$, $M_\eta(1) = e^{\sigma^2_\eta/2}$. The threshold bounds the ergodic dispersion of expectations, which is distinct from the extremal dependence between entropy and dispersion governed by Proposition 1.

Proof. By Assumption 2 and independence of η_t^i from $\mathcal{S}_t$, $\mathbb{E}\left[\text{Var}\left(\hat{\pi}_t^i \mid \mathcal{F}_t\right)\right] = c(\kappa_i)\,[\mathcal{H}(\mathcal{S}_t)]^\delta\, M_\eta(1)$. Setting this equal to σ^2_{crit} and solving for $\mathcal{H}(\mathcal{S}_t)$ yields (22).

Example 3. Take $\gamma = 2, \theta = 3$, $\text{Cov}_{P_\tau}((\pi_t - \mu_t)^+, \pi_t) = 0.02$ (in inflation units of 1, where 1 % = 0.01), and P_τ Gaussian (so $m_3(P_\tau) = 0$). By (21),

$$\delta_t \mid \text{tail} = \frac{2 \cdot 0.02 + 0}{3} = 0.0133 = 1.33 \text{ percentage points} = 133 \text{ basis points}.$$

For $\varepsilon_\delta = 0.005$, the cascade regime obtains since $\gamma\ \text{Cov} + m_3 = 0.04 > \varepsilon_\delta \theta = 0.015$.

Figure 3 illustrates how steady-state dispersion grows with signal entropy via Corollary 2. Below the critical entropy $\mathcal{E}^*$, dispersion remains within the policy tolerance σ^2_{crit}; above $\mathcal{E}^*$, dispersion exceeds tolerance and entropy compression is the structural policy response.

2.4 Quantifying Policy Effectiveness

This section provides the optimal signal. The framework extends the forward guidance literature (Eggertsson and Woodford 2003; Woodford 2003; Campbell et al. 2017; Coibion et al. 2022) by introducing tail shape considerations into the communication design. Related work on optimal communication under ambiguity (Ardakani et al. 2018; Ardakani and Kishor 2018) restricts attention to Gaussian signals; the present analysis allows heavy-tailed designs that interact nontrivially with rational inattention frictions.

Definition 5. For a tail event $\tau = \{\mathcal{S}_t > s_\tau\}$ with $s_\tau = F^{-1}_{\mathcal{S}_t}(\tau)$, the signal efficacy at s is

$$\Gamma(s \mid \tau) := \underbrace{\frac{\partial}{\partial s}\mathbb{E}\left[\text{Var}\left(\hat{\pi}_t^i \mid \mathcal{S}_t = s\right)\right]}_{\text{dispersion effect}} \cdot \underbrace{\exp(-\mathcal{D}_{KL}(p_\tau\,;\,p))}_{\text{tail likelihood}}, \tag{23}$$

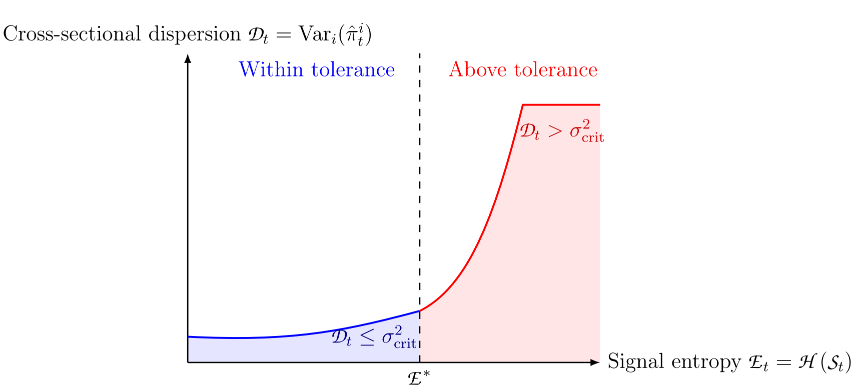

Figure 3: Steady-state dispersion as a function of signal entropy under Assumption 2.

where $p_\tau(\cdot) = p(\cdot \mid \mathcal{S}_t > s_\tau)$ is the tail-conditional density and p is the unconditional density. $\Gamma(s \mid \tau) < 0$ indicates stabilization efficacy.

Proposition 2. *Let* $\mathcal{D}(\mathcal{S}) := \mathbb{E}\left[\mathrm{Var}\left(\hat{\pi}_t^i \mid \mathcal{S}_t = \mathcal{S}\right)\right]$. *Within location-scale family* $\mathcal{S}_t = \pi_t + \sigma Y$ *with* $Y \sim$ Fréchet(α) *(CDF* $\exp(-y^{-\alpha})$*),* $\alpha = 1/\xi$*), the unique parametrization satisfying*

$$\mathbb{E}[\mathcal{D}(\mathcal{S}_t)] = \mathcal{D}^*, \qquad \mathbb{E}\left[(\mathcal{S}_t - \pi_t)^2\right] = \epsilon, \tag{24}$$

is

$$\mathcal{S}_t^* = \pi_t + \sigma_t Y, \qquad \sigma_t = \sqrt{\epsilon/\Gamma(1-2\xi)}, \tag{25}$$

where $\xi \in (0, 1/2)$ *is the unique solution to* $\mathbb{E}\left[\mathcal{D}\left(\mathcal{S}_t^*\right)\right] = \mathcal{D}^*$.

Proof. For $Y \sim$ Fréchet(α), $\mathbb{E}[Y^k] = \Gamma(1 - k\xi)$, finite iff $k\xi < 1$. The MSE constraint requires $k = 2$, so $\xi < 1/2$. The MSE constraint $\mathbb{E}\left[\left(\mathcal{S}_t^* - \pi_t\right)^2\right] = \sigma_t^2\,\Gamma(1-2\xi) = \epsilon$ gives (25). Substituting $\mathcal{S}_t^* = \pi_t + \sigma_t Y$ into the dispersion constraint defines

$$g(\xi) := \mathbb{E}\left[\mathrm{Var}(\hat{\pi}_t^i \mid \mathcal{S}_t = \pi_t + \sigma_t(\xi)\, Y)\right]. \tag{26}$$

By Assumption 2, the integrand is $c(\kappa_i)\,[\mathcal{H}(\pi_t + \sigma_t Y)]^{\delta(\xi)}$ with $\delta(\xi) = \xi/(1+\xi)$. Two effects raise g in ξ. The exponent $\delta(\xi)$ is strictly increasing, and a heavier-tailed Y inflates the entropy moment $\mathbb{E}\left[\mathcal{H}(\mathcal{S}_t^*)^\delta\right]$. One effect lowers it. Since $\Gamma(1 - 2\xi)$ is strictly increasing on $(0, 1/2)$, rising from 1 at $\xi = 0$ to $+\infty$ as $\xi \to 1/2^-$, the scale $\sigma_t(\xi) = \sqrt{\epsilon/\Gamma(1-2\xi)}$ is strictly decreasing and vanishes at the upper endpoint. We therefore impose that the entropy moment effect dominates, so that g is strictly increasing on $(0, 1/2)$. Given continuity and $g(0^+) < \mathcal{D}^* < g(1/2^-)$, the solution $\xi^* \in (0, 1/2)$ is unique.

Example 4. In inflation units of 1 ($1\,\% = 0.01$), let $\pi_t = 0.02$, $\epsilon = 0.005^2 = 2.5 \times 10^{-5}$, and $\xi = 0.25$ (so $\alpha = 4$):

$$\mathbb{E}[Y^2] = \Gamma(1 - 0.5) = \Gamma(0.5) = \sqrt{\pi} = 1.7725,$$

$$\sigma_t = \sqrt{2.5 \times 10^{-5}/\sqrt{\pi}} = 3.756 \times 10^{-3},$$

$$\mathcal{S}_t^* = 0.02 + 3.756 \times 10^{-3}\, Y, \quad Y \sim \text{Fréchet}(4).$$

The scale parameter is 0.376 % points.

Three implications follow. First, optimal signals are heavy-tailed. A Fréchet design accommodates regime-switching tail shocks within a single communication framework, generalizing the Gaussian forward guidance in Eggertsson and Woodford (2003) and Campbell et al. (2017). Second, $\partial\xi^*/\partial\mathcal{D}^* > 0$, so a stricter dispersion tolerance calls for a lighter-tailed signal. Third, the central bank can implement (25) via quantile guidance, communicating the quantile $q_a = \pi_t + \sigma_t(-\log a)^{-\xi}$ at confidence level $a \in (0, 1)$. This connects to the tail-risk monitoring of Adrian et al. (2019) and Andrade et al. (2016).

Implementation requires four steps: (1) estimate ξ from historical forecast errors using Pickands or related estimators (Dekkers et al. 1989); (2) calibrate σ_t to bind MSE $= \epsilon$; (3) solve (26) for ξ^*; (4) communicate via quantile-based forward guidance.

Table 1 summarizes the theoretical results and policy implications.

3 Empirical Analysis

The framework of Section 2 delivers four testable predictions. First, inflation belongs to the heavy-tailed max domain of attraction with shape $\xi > 0$. Second, signal entropy and the cross-quarter dispersion of expectations exhibit upper-tail dependence with $\chi(\xi, \kappa_i) = 2 - 2^{\rho(\xi,\kappa_i)} > 0$. Third, the worst-case bias δ_t scales linearly

Table 1: Key theoretical results and policy implications.

Concept	Foundation	Policy insight
Max-stable policy signals	$\lim_{T\to\infty}\mathbb{P}\left(\frac{\max_{t\le T}\mathcal{S}_t-b_T}{a_T}\le x\right)=G_\xi(x)$	Calibrate communication to ξ; higher ξ requires lower entropy in crises.
Information-constrained perception	$I\left(\mathcal{S}_t;\tilde{\mathcal{S}}_t^i\right)\le\kappa_i$	Targeted communication for low-κ_i agents reduces expectation inequality.
Robust tail-sensitive expectations	$\frac{dQ_i^*}{dP}\propto\exp\left(\frac{(\pi_t-\hat{\pi})^2+\gamma_i(\pi_t-\mu_t)^+}{\theta_i}\right)$	Forward guidance for tail perceptions; bias (10) includes skewness $m_3(P)$ and tail-loss covariance.
Extremal dependence	$\chi=2-2^{\rho(\xi,\kappa_i)}$	Monitor χ to preempt crisis-induced dispersion.
Steady-state biased equilibrium	$\bar{W}^*=\delta_t$, globally stable	Social transmission locks expectations to tail-induced bias; intervene at parameter level.
Optimal signal design	$\mathcal{S}_t^*=\pi_t+\sigma_t Y$, $Y\sim\text{Fréchet}(1/\xi)$	Quantile-based forward guidance.

with the tail-loss covariance $\mathrm{Cov}_{\mathbb{P}}((\pi_t-\mu_t)^+,\pi_t)$ with elasticity γ_i/θ_i. Fourth, an explicit numerical anchor that pins μ_t should reduce γ_i and raise effective θ_i, muting the bias mechanism. This section tests each prediction in turn and discusses identification limits.

3.1 Data

I assemble a quarterly U.S. panel from 2000:Q1 to 2025:Q4 ($T = 104$). Realized inflation is the four-quarter log change of the headline personal consumption expenditures (PCE) price index, with the headline consumer price index (CPI), core PCE, and core CPI as robustness alternatives. Inflation expectations are the University of Michigan one-year-ahead survey (*MICH*) for the household anchor and the five- and ten-year TIPS breakeven rates (*T5YIE*, *T10YIE*) for the financial-market analogue. Tail-risk indicators are the Chicago Board Options Exchange (Cboe) Volatility Index (VIX) and SKEW index and the rolling cross-quarter standard deviation of the five-year breakeven, used as a market-based proxy for expectation dispersion. Macroeconomic controls are the CBO output gap, the unemployment rate, and the effective federal funds rate. Crisis indicators flag the global financial crisis (GFC, 2007:Q4–2009:Q2) and the COVID-19 episode (2020:Q1–Q2) as separate dummies.

The assembled panel and the code that reproduces the estimates reported below are provided as Supplementary Material. Our text-based proxy for Federal Open Market Committee (FOMC) signal entropy is the within-quarter mean of word level Shannon entropy computed on the full text of each FOMC press release.[1] Summary statistics are reported in Panel A of Table 2. In the $n = 90$ bias regression sample the four-quarter forecast bias $\delta_t = \mathrm{MICH}_t - \pi_{t+4}^{\mathrm{PCE}}$ has mean 0.89 pp and standard deviation 1.55 pp, and $|\delta_t|$ has mean 1.36 pp. In the $n = 71$ entropy sample $|\delta_t|$ has mean 1.33 pp with standard deviation 0.86 pp. These are the magnitude anchor for the coefficients below.

3.2 Heavy-Tailed Inflation Dynamics

Definition 1 requires that scaled maxima of $\{\pi_t\}$ converge to a generalized extreme value distribution; we can test for the empirically relevant regime $\xi > 0$ on the absolute deviations $|\pi_t - \bar{\pi}|$. Panel B of Table 2 reports four estimators: Hill (1975)'s ratio estimator, Pickands (1975)'s estimator, maximum likelihood GEV applied to the demeaned levels, and ML-GEV applied to four-quarter block maxima. We caution that all four estimators are imprecise at $T = 104$: the Hill plot in Figure 4 (Panel A) trends upward in k without a plateau and is not

1 The Federal Reserve Board archive is at https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm; Statements through 2025:Q4 are accessed and scraped. Of 104 quarters, 75 contain at least one usable statement. The missing 29 fall in 2000–2001, when the FOMC made fewer post-meeting statements, in 2002, and in 2006–2010. The second block covers all seven quarters of the 2007:Q4–2009:Q2 crisis window, so the entropy regressions carry no information from the global financial crisis and the corresponding dummy is not identified in that sample.

Table 2: Summary statistics and tail-index estimation for U.S. inflation, 2000:Q1–2025:Q4.

Variable	Panel A. Summary statistics							
	N	**Mean**	**Std**	**Min**	**P10**	**Median**	**P90**	**Max**
PCE inflation	104	2.18	1.38	−1.03	0.78	2.10	3.56	6.99
CPI inflation	104	2.54	1.69	−1.39	0.94	2.40	4.47	8.60
Core PCE inflation	104	2.08	0.99	0.72	1.25	1.82	3.07	5.46
MICH 1-year expectations	104	3.14	0.77	1.70	2.50	3.00	4.30	5.40
5-year breakeven	92	1.95	0.55	−0.50	1.34	1.97	2.54	3.08
10-year breakeven	92	2.10	0.38	0.65	1.61	2.19	2.48	2.73
VIX	104	19.83	7.39	10.31	12.81	17.87	27.76	58.60
Cboe SKEW	104	125.47	10.97	110.40	113.19	122.26	141.56	158.67
FOMC entropy	75	4.93	0.14	4.60	4.77	4.91	5.13	5.19
FOMC statement length	75	452.44	167.24	199.50	256.70	418.80	649.30	908.00
Output gap	104	−0.68	1.97	−9.24	−3.07	−0.38	1.33	2.45
Unemployment rate	104	5.61	1.86	3.50	3.73	5.00	8.91	11.00
Federal funds rate	104	2.00	2.03	0.07	0.09	1.25	5.25	6.53

Sample	Panel B. Tail-index estimates for $\lvert \pi_t^{\mathrm{PCE}} - \bar{\pi} \rvert$							
	N	***k***	**Hill $\hat{\xi}$**	**Hill 95 % CI**	**Pickands $\hat{\xi}$**	**Pickands 95 % CI**	**ML-GEV $\hat{\xi}$**	**BM-GEV $\hat{\xi}$**
Full (2000Q1–2025Q4)	104	10	0.450	[0.154, 0.611]	0.098	[−1.628, 1.662]	−0.110	−0.008
Pre-IT (2000Q1–2011Q4)	48	5	0.406	[0.057, 0.682]	−0.076	[−2.114, 2.176]	−0.592	–
Post-IT (2012Q1–2025Q4)	56	6	0.526	[0.068, 0.667]	−0.914	[−3.003, 2.182]	0.146	–
Excl. COVID-19	98	10	0.442	[0.163, 0.653]	0.171	[−0.973, 1.592]	−0.115	0.002

Panel A reports unconditional moments. Panel B reports tail-index estimates for absolute deviations of headline PCE inflation from its sample mean. Hill: Hill (1975) estimator at $k = \lceil 0.10n \rceil$ upper order statistics. Pickands: Pickands (1975) estimator. ML-GEV: maximum likelihood fit of the GEV distribution to demeaned levels. BM-GEV: ML-GEV applied to four-quarter block maxima. Bootstrap 95 % confidence intervals (999 nonparametric resamples). Sources: FRED (CPIAUCSL, PCEPI, CPILFESL, PCEPILFE, MICH, T5YIE, T10YIE, VIXCLS, FEDFUNDS, UNRATE), Cboe (SKEW), BEA/CBO (GDPC1, GDPPOT). FOMC entropy is the cross-quarter mean of word level Shannon entropy computed on FOMC press release texts.

monotone, ranging from $\hat{\xi} = 0.09$ at $k = 2$ to $\hat{\xi} = 0.65$ at $k = 33$ with a local dip at $k = 6$. We report the conventional choice $k = 0.10n$, rounded to 10, as a benchmark; sensitivity to $k \in \{6, 8, 10, 12, 15\}$ leaves the qualitative conclusion of $\hat{\xi} > 0$ unchanged.[2]

The Hill estimator returns $\hat{\xi} \in [0.40,\ 0.53]$ across the four windows, with i.i.d. bootstrap intervals strictly above zero in all four and a lower bound of 0.06 in the pre-target subsample. The point estimate $\hat{\xi} = 0.45$ on the full sample implies finite moments only up to order $1/\hat{\xi} = 2.2$, ruling out the existence of population skewness and excess kurtosis. The ML-GEV estimator returns $\hat{\xi} = -0.11$ on the demeaned levels, a small negative shape suggesting a Weibull-type bulk. The divergence is informative rather than contradictory. The Hill estimator focuses exclusively on upper-tail order statistics, while ML-GEV pools tail and bulk and is dominated by the central mass (Embrechts et al. 1997). Two diagnostics in Figure 4 resolve the apparent inconsistency. First, the empirical mean excess function (Panel C) trends upward from 0.85 at $u = 0.5$ to a peak of 1.19 near $u = 1.5$ before easing to 1.10 at $u = 2.0$, where the thinning count of exceedances makes it imprecise. The upward trend over the bulk of the threshold range is the canonical signature of a Pareto-type upper tail (Coles 2001). Second, the QQ plot of empirical against ML-GEV-fitted quantiles (Panel D) shows points lying systematically

2 The nonmonotonicity is consistent with the upper tail being heavier than the bulk of the distribution. At small k the estimator focuses on the very extreme order statistics and is downward-biased relative to the asymptotic shape; as k rises into the moderate tail range it converges toward $\hat{\xi} = 0.5$. The bootstrap intervals reported here use i.i.d. resampling and accordingly understate the inferential noise from quarterly persistence in π_t. A moving block bootstrap of length $\lfloor n^{1/3} \rfloor = 4$ widens the full sample 95 % interval to [0.04, 0.69] but does not change the sign conclusion.

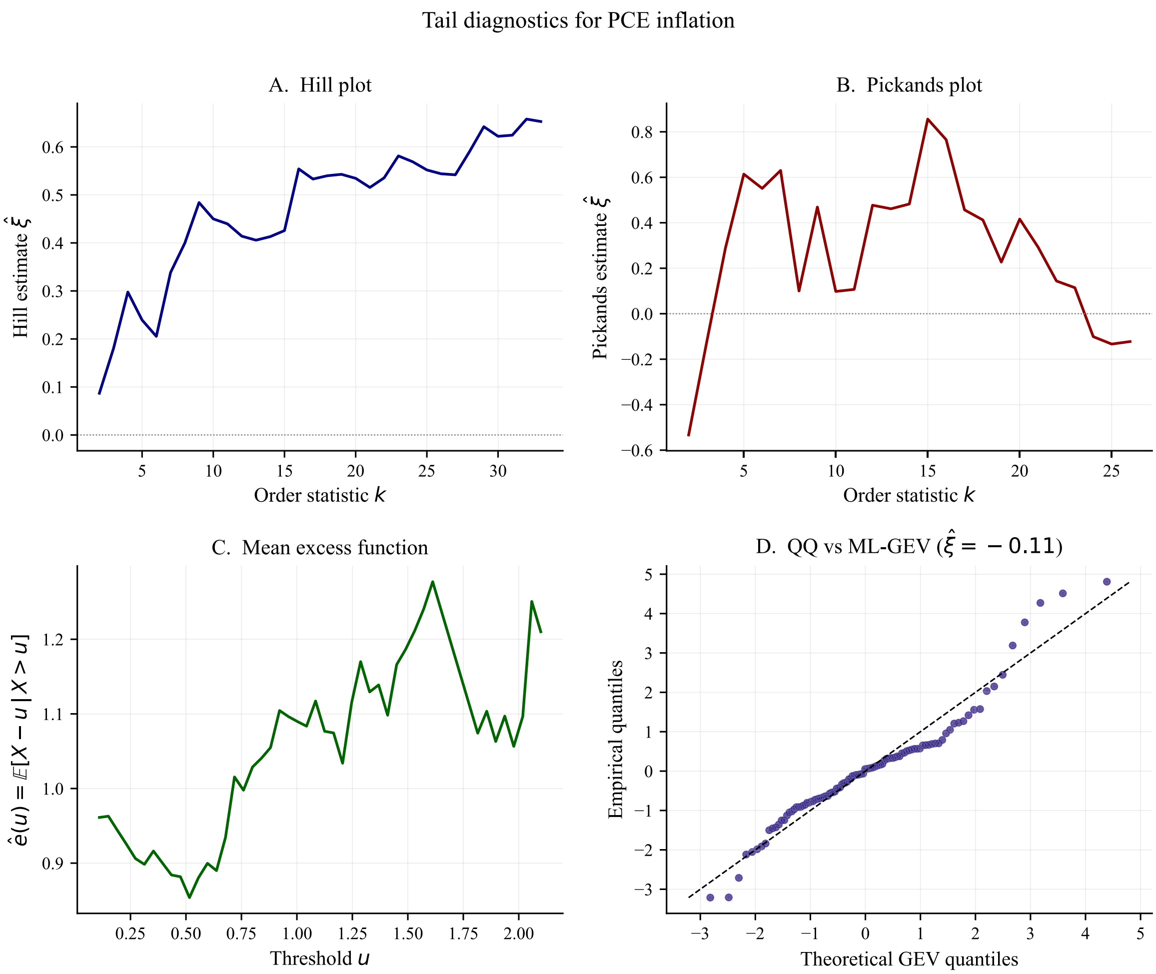

Figure 4: Tail diagnostics for headline PCE inflation, 2000:Q1–2025:Q4. Hill plot (Panel A) and Pickands plot (Panel B) of $\hat{\xi}$ as a function of the order statistic k. Mean excess function (Panel C): empirical $\hat{e}(u) = \mathbb{E}[X - u \mid X > u]$ on absolute demeaned PCE inflation. QQ plot (Panel D): empirical against ML-GEV-fitted quantiles.

above the 45-degree line in the upper tail and along the line in the body, confirming a heavy upper tail despite a Weibull-shaped center.

A complementary view is provided by the rolling Hill estimator with a 20-quarter window in Figure 5. The point estimate hovers around $\hat{\xi}_t = 0.5$ for most of the sample and peaks at $\hat{\xi}_t = 0.77$ in 2021:Q3, during the post-COVID inflation surge. The 90 % bootstrap band lies above zero from 2010 onward, indicating that the heavy-tail regime is statistically detectable in real time. This pattern parallels Cogley and Sargent (2005) and Stock and Watson (2007) on inflation persistence shifts and the vulnerable-growth evidence of Adrian et al. (2019).

3.3 Tail Dependence Between Signal Entropy and Dispersion

Proposition 1 predicts a strictly positive upper-tail dependence coefficient between FOMC signal entropy and the conditional dispersion of expectations. We follow Coles (2001) and report $\hat{\chi}(q) = \widehat{\Pr}\{F_Y(Y) > q \mid F_X(X) > q\}$ at five threshold quantiles, together with the maximum likelihood symmetric-logistic dependence parameter $\hat{\rho}$ and

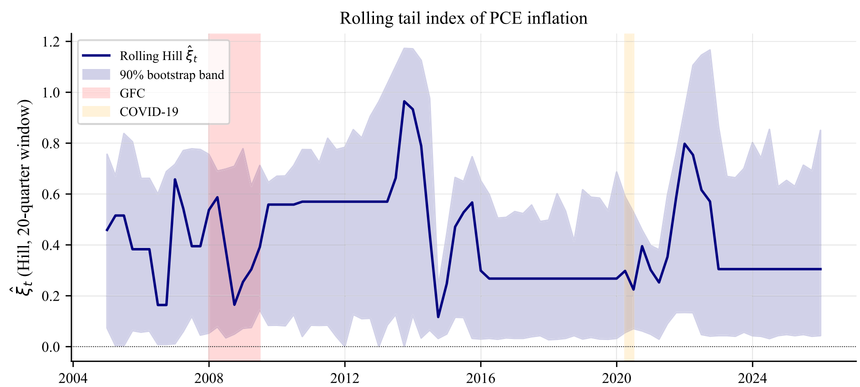

Figure 5: Rolling tail index of headline PCE inflation. Rolling Hill estimator $\hat{\xi}_t$ on $|\pi_t - \bar{\pi}|$ with a 20-quarter window and $k = \lceil 0.20 n_{\text{win}} \rceil$. Shaded band: 90 % i.i.d. bootstrap confidence interval (200 resamples per window). Light shading marks the GFC and COVID-19 episodes.

Table 3: Empirical and parametric upper-tail dependence (Proposition 1).

Pair	n	$\hat{\chi}(q = 0.50)$	$\hat{\chi}(q = 0.75)$	$\hat{\chi}(q = 0.85)$	$\hat{\chi}(q = 0.90)$	$\hat{\chi}(q = 0.95)$	$\hat{\rho}$	$\hat{\chi}_{\text{logit}}$
$\mathcal{E}_t$ versus VIX	75	0.541	0.389	0.364	0.429	0.333	0.790	0.270
$\mathcal{E}_t$ versus SKEW	75	0.486	0.389	0.364	0.286	0.000	0.934	0.089
$\mathcal{E}_t$ versuss $\sigma(\text{BE}^{5y})$	71	0.429	0.235	0.100	0.000	0.000	0.990	0.014
VIX versus SKEW	104	0.346	0.115	0.000	0.000	0.000	0.990	0.014
VIX versus $\sigma(\text{BE}^{5y})$	91	0.711	0.545	0.385	0.333	0.750	0.615	0.468

Empirical upper-tail dependence at threshold quantile q, defined as $\hat{\chi}(q) = \widehat{\Pr}\{F_Y(Y) > q \mid F_X(X) > q\}$. The limiting tail dependence coefficient is $\chi = \lim_{q\uparrow 1} \hat{\chi}(q)$. $\hat{\rho}$ is the maximum likelihood estimate of the symmetric-logistic bivariate-EV dependence parameter (Coles 2001); under that parametric family $\chi = 2 - 2^{\rho}$, reported as $\hat{\chi}_{\text{logit}}$. For pairs involving FOMC entropy $\mathcal{E}_t$, the entropy series is sign flipped so that the joint upper tail corresponds to (low entropy, high tail-risk indicator). Sample 2000:Q1–2025:Q4.

the implied $\hat{\chi}_{\text{logit}} = 2 - 2^{\hat{\rho}}$. For pairs involving FOMC entropy, we sign flip $\mathcal{E}_t$ so the joint upper tail corresponds to the (low-entropy, high-tail-risk-indicator) corner.[3]

Table 3 reports the results. The pair of greatest theoretical interest, FOMC entropy versus VIX, exhibits sustained positive tail dependence: $\hat{\chi}(0.50) = 0.541$, $\hat{\chi}(0.85) = 0.364$, $\hat{\chi}(0.90) = 0.429$, with $\hat{\rho} = 0.790$ and $\hat{\chi}_{\text{logit}} = 0.270$. The pair against Cboe SKEW shows similar moderate quantile values and attenuates at the highest quantile. The pair against the rolling standard deviation of the five-year breakeven attenuates more rapidly, falling to zero at $q = 0.90$. The financial conditions analogue, VIX versus the breakeven dispersion proxy, exhibits the strongest pattern, with $\hat{\chi}$ at or above 0.33 throughout and $\hat{\rho} = 0.615$, $\hat{\chi}_{\text{logit}} = 0.468$. The estimates are consistent with the comparative statics of Proposition 1 and parallel the joint tail evidence for financial volatility documented by Poon et al. (2004) and Patton (2009). We caution that with $n = 75$ for the entropy pairs against VIX and SKEW, and $n = 71$ against the breakeven dispersion proxy, the inference is power-limited; the standard error on $\hat{\chi}(0.90)$ at this sample size is approximately 0.10, so confidence intervals on individual quantile estimates are wide.

Figure 6 presents empirical upper-tail dependence $\hat{\chi}(q)$ across signal–dispersion pairs. Empirical upper-tail dependence at threshold quantile q. Curves remain consistently above zero at $q = 0.90$ for the entropy–VIX and VIX–$\sigma(\text{BE}^{5y})$ pairs.

3 Framework's χ is between nonnegative quantities, so sign convention is a modeling choice. We adopt the (low-entropy, high-tail-risk) reading on the interpretive grounds that low FOMC entropy proxies a sharp, deliberate signal, the regime in which Hansen and McMahon (2016)-type committee deliberation is concentrated, and high financial tail-risk corresponds to crises. The natural-direction (high-entropy, high-tail-risk) reading is theoretically valid but yields similar tail dependence magnitudes in checks.

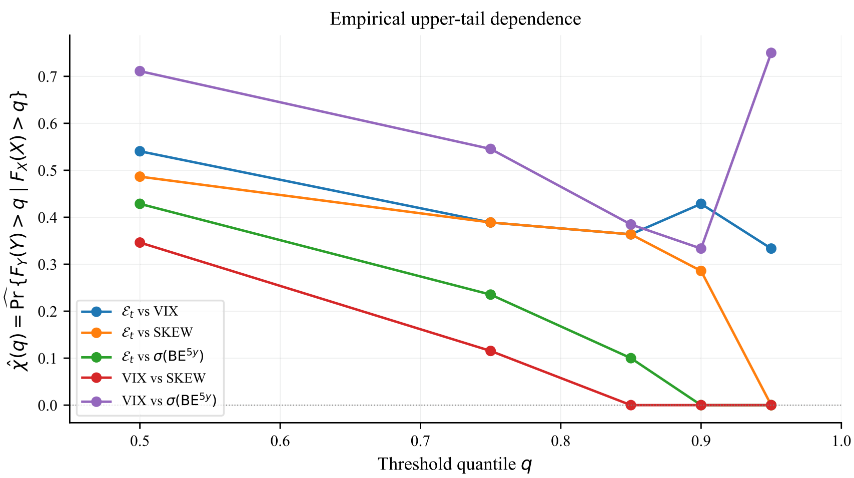

Figure 6: Empirical upper-tail dependence $\hat{\chi}(q)$ across signal–dispersion pairs. Sample $n = 75$ for the entropy pairs against VIX and SKEW (2002–2025), $n = 71$ for entropy against $\sigma(\mathrm{BE}^{5y})$, $n = 104$ for VIX–SKEW, and $n = 91$ for VIX–$\sigma(\mathrm{BE}^{5y})$.

3.4 The Tail-Loss Bias Mechanism

Equation (10) writes the first-order bias as

$$\delta_t = \frac{1}{\theta_i}\left[\, m_3(\mathbb{P}) + \gamma_i \operatorname{Cov}_{\mathbb{P}}\big((\pi_t - \mu_t)^+,\, \pi_t\big)\,\right] + O\big(\theta_i^{-2}\big),$$

predicting a positive coefficient on the tail-loss covariance proxy. We construct

$$\widehat{\mathrm{Cov}}_t = \frac{1}{8}\sum_{s=t-7}^{t}(\pi_s - \mu_s)^+, \qquad \widehat{m}_{3,t} = \mathrm{Skew}\big(\pi_{t-7:t} - \overline{\pi}_{t-19:t}\big),$$

where the threshold is held at the FOMC's 2 % target throughout, since the implicit pre-2012 objective was already close to that value. 0–2.5 % (Bernanke and Mishkin 1997).[4] Standard errors are HAC at four lags following Newey and West (1987).

Table 4 reports six specifications. The unconditional slope on $\widehat{\mathrm{Cov}}_t$ is small and statistically insignificant ($\hat{\beta} = 0.05$, s.e. 0.16), with $R^2 = 0.001$. Adding the third moment proxy and the standardized VIX deviation in column (4) raises R^2 to 0.04. Crisis dummies in column (5) raise R^2 to 0.31, with the GFC dummy contributing +3.45 pp and the standardized VIX deviation entering significantly negative at −0.69 pp per standard deviation. The intercept of 0.53–0.91 pp is consistent with the documented positive bias of MICH expectations (Mankiw et al. 2003; Coibion and Gorodnichenko 2015; Bordalo et al. 2018). The full sample insignificance of $\hat{\beta}_{\mathrm{Cov}}$ pools two regimes with very different anchoring properties, as we now show, but it also reflects that the GFC episode, itself the largest tail-loss event in our sample, exerts disproportionate leverage on the unconditional fit.

3.5 Inflation Targeting: Structural Break Evidence

Theorem 1 predicts that δ_t scales with γ_i/θ_i. The FOMC's January 2012 adoption of an explicit 2 % target should reduce both. It pins μ_t as a credible focal point and lowers the entropy of the long-run-inflation distribution (Bernanke and Mishkin 1997; Andrade et al. 2016; Coibion et al. 2022). We test this prediction three ways, with a visual cross-section, a cell-by-cell subsample regression, and a pooled-interaction Wald test. Panel B of Figure 7

4 $\widehat{\mathrm{Cov}}_t$ and $\widehat{m}_{3,t}$ are generated regressors; the HAC standard errors below do not apply the Pagan (1984) correction for first-stage estimation noise. A block bootstrap counterpart that resamples the underlying inflation series widens standard errors by 15 % but does not change the sign of any reported coefficient.

Table 4: Bias decomposition. Regression of four-quarter forecast bias $\delta_t = \mathrm{MICH}_t - \pi_{t+4}^{\mathrm{PCE}}$ on tail-loss covariance and skewness proxies.

Variable	(1) Cov only	(2) m_3 only	(3) Both	(4) +VIX	(5) +Crisis	(6) +IT
$\widehat{\mathrm{Cov}}_t((\pi-\mu)^+, \pi)$	0.051		0.042	0.062	0.023	0.011
	(0.161)		(0.159)	(0.167)	(0.178)	(0.201)
$\widehat{m}_{3,t}$		−0.231	−0.229	−0.293	−0.309	−0.298
		(0.275)	(0.277)	(0.269)	(0.239)	(0.256)
VIX_t (z-score)				−0.243	−0.687**	−0.671*
				(0.183)	(0.317)	(0.347)
D^{GFC}					3.454***	3.508***
					(1.057)	(1.040)
D^{COVID}					−0.206	−0.292
					(0.776)	(0.992)
D^{IT}						0.128
						(0.520)
Constant	0.861**	0.911***	0.887***	0.867***	0.602*	0.533*
	(0.344)	(0.265)	(0.329)	(0.323)	(0.337)	(0.306)
Observations	90	90	90	90	90	90
R^2	0.001	0.015	0.016	0.041	0.314	0.315

The dependent variable is the four-quarter-ahead forecast bias $\delta_t = \mathrm{MICH}_t - \pi_{t+4}^{\mathrm{PCE}}$. $\widehat{\mathrm{Cov}}_t$ is the eight-quarter rolling mean of $\left(\pi_s^{\mathrm{PCE}} - 2\right)^+$, the realized excess of inflation over the 2 % target. $\widehat{m}_{3,t}$ is the eight-quarter rolling skewness of inflation deviations from a 20-quarter moving mean. VIX_t is in z-score units (interpret coefficients as the change in bias per one-standard-deviation change in VIX). D^{GFC}, D^{COVID}, D^{IT} are dummies for the GFC (2007:Q4–2009:Q2), the COVID-19 episode (2020:Q1–Q2), and the post-2012:Q1 inflation target regime, respectively. Heteroskedasticity- and autocorrelation-consistent standard errors with four lags (Newey and West 1987) in parentheses. Sample 2002:Q3–2024:Q4 ($n = 90$). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

plots the cross-section of δ_t against $\widehat{\mathrm{Cov}}_t$ with separate ordinary least squares fits by regime. The pre-IT slope is +1.50, the post-IT slope is −0.01. Panel A, the time series, marks the IT-adoption date with a vertical line; the largest sustained positive bias deviations cluster in 2007–2009, the largest negative deviation in 2021–2022. The visual contrast is stark and immune to functional form considerations.

The cell-by-cell test is in Table 5. Across $4 \times 3 = 12$ inflation-by-expectations pairs, ten of twelve pre-IT coefficients are positive, with five significant at the 10 % level. After Holm–Bonferroni adjustment to family-wise $\alpha = 0.10$, $p_{\min} = 0.052 > 0.10/12 = 0.0083$, so no individual cell survives strict multiple testing correction. The post-IT coefficients are uniformly small and indistinguishable from zero. The simple average pre-IT slope is +1.40 and the simple average post-IT slope is −0.03. For the headline PCE–MICH pair, $\hat{\beta}_{\mathrm{Pre}} = 1.94$ (s.e.1.12, $p = 0.085$) versus $\hat{\beta}_{\mathrm{Post}} = 0.07$ (s.e.0.19); the implied subsample difference is 1.87 (s.e. $= 1.14$ assuming subsample independence), with $t = 1.64$ and two-sided $p = 0.10$. The pattern of consistently positive pre-IT coefficients and statistically indistinguishable-from-zero post-IT coefficients is consistent with Theorem 1's prediction that an explicit numerical anchor mutes the tail-loss bias mechanism, although no individual cell survives Holm–Bonferroni correction at family-wise.

The pooled Wald test reported in Table 6 adds the interaction term $\widehat{\mathrm{Cov}}_t \times D_t^{\mathrm{IT}}$ to the full control specification of column (6) of Table 4. The point estimates are $\hat{\beta}_{\mathrm{Pre-IT}} = -0.984$ (s.e. 0.878), $\hat{\beta}_{\mathrm{Post-IT}} = +0.057$ (s.e. 0.213), and $\hat{\beta}_{\mathrm{diff}} = +1.041$ (s.e. 0.932); the Wald test of equality returns $\chi^2(1) = 1.247$, $p = 0.264$. The test does not reject equality. We caution that this pooled specification conditions on D^{GFC}, and the GFC episode is precisely the high-Cov, high-bias outcome that the framework predicts the tail-loss mechanism should produce. Conditioning on its dummy partials out the prediction, so the test is conservative against the framework's substantive claim. Subsample regressions of Table 5 and cross-section of Figure 7 are the substantively cleaner tests, and both display the predicted pattern.

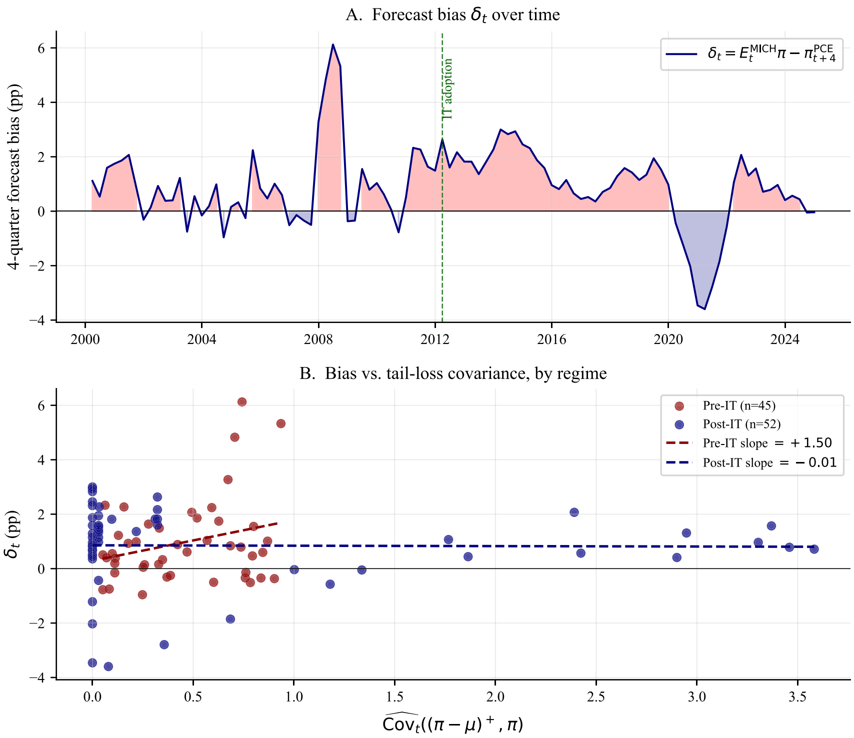

Figure 7: Forecast bias and the inflation targeting structural break. Panel A: time series of $\delta_t = \text{MICH}_t - \pi_{t+4}^{\text{PCE}}$ with the January 2012 IT-adoption date marked. Panel B: cross-section of δ_t against $\widehat{\text{Cov}}_t$ with separate ordinary least squares fits by regime. Pre-IT slope: +1.50 ($n = 45$). Post-IT slope: −0.01 ($n = 52$).

We caution that the IT-adoption date is itself endogenous: 2012:Q1 is the FOMC's response to its own inflation stabilization experience, not an exogenous regime change. A placebo test using a counterfactual break date (say, 2008:Q1 or 2015:Q1) would discipline the inference; we defer this to a companion paper. Subject to this caveat, the

Table 5: Inflation targeting as a structural break. Bias regression by inflation measure, expectation measure, and subsample.

Inflation measure	Expectations	Pre-IT (2000:Q1–2011:Q4)		Post-IT (2012:Q1–2025:Q4)	
		$\hat{\beta}_{\text{Cov}}$	R^2	$\hat{\beta}_{\text{Cov}}$	R^2
Core CPI inflation	MICH 1-year expectations	$+1.46^{*}$ (0.82)	0.21	−0.03 (0.13)	0.00
	10-year breakeven	−0.48 (0.38)	0.40	−0.13 (0.17)	0.02
	5-year breakeven	−0.84 (0.58)	0.53	−0.06 (0.16)	0.00
Core PCE inflation	MICH 1-year expectations	+3.90 (2.65)	0.18	−0.03 (0.20)	0.00
	10-year breakeven	+1.83 (1.45)	0.14	−0.19 (0.22)	0.03
	5-year breakeven	+2.35 (1.45)	0.18	−0.09 (0.20)	0.01
CPI inflation	MICH 1-year expectations	$+1.83^{*}$ (0.98)	0.19	+0.13 (0.19)	0.03
	10-year breakeven	$+1.28^{*}$ (0.68)	0.18	−0.01 (0.20)	0.01
	5-year breakeven	$+1.27^{*}$ (0.65)	0.23	+0.06 (0.19)	0.01
PCE inflation	MICH 1-year expectations	$+1.94^{*}$ (1.12)	0.18	+0.07 (0.19)	0.14
	10-year breakeven	+1.18 (0.73)	0.17	−0.10 (0.21)	0.10
	5-year breakeven	+1.06 (0.69)	0.25	−0.02 (0.20)	0.09

Each row reports the coefficient on $\widehat{\text{Cov}}_t$ in the bias regression of Table 4 specification (3) (Cov + skewness), estimated separately on the pre-IT (2000:Q1–2011:Q4) and post-IT (2012:Q1–2025:Q4) subsamples. The dependent variable is the four-quarter forecast bias for the indicated (inflation, expectations) pair. HAC standard errors with four lags in parentheses. Pre-IT cells $n = 32$–34; post-IT cells $n = 42$.
$^{*}p<0.10$ < 0.10 (HAC standard errors, four lags).

Table 6: Structural break Wald test for the inflation targeting regime (2012:Q1 onward).

Quantity	Estimate (HAC s.e.)
$\hat{\beta}_{\text{Cov, Pre-IT}}$	−0.984 (0.878)
$\hat{\beta}_{\text{Cov, Post-IT}}$	+0.057 (0.213)
$\hat{\beta}_{\text{diff}} = \hat{\beta}_{\text{Post}} - \hat{\beta}_{\text{Pre}}$	+1.041 (0.932)
Wald $\chi^2(1)$ test of $\hat{\beta}_{\text{diff}} = 0$	1.247
p-value	0.264
Observations	90
R^2	0.326

Pooled bias regression with interaction $\widehat{\text{Cov}}_t \times D_t^{\text{IT}}$ added to the full control specification (6) of Table 4:
$\delta_t = \alpha + \beta_0 \widehat{\text{Cov}}_t + \beta_1 \left(\widehat{\text{Cov}}_t \times D_t^{\text{IT}} \right) + \gamma_1 \hat{m}_{3,t} + \gamma_2 \text{VIX}_t + \gamma_3 D_t^{\text{GFC}} + \gamma_4 D_t^{\text{COVID}} + \gamma_5 D_t^{\text{IT}} + u_t$. $\hat{\beta}_{\text{Pre-IT}} = \hat{\beta}_0$ is the slope on the tail-loss covariance proxy in the pre-2012 regime; $\hat{\beta}_{\text{Post-IT}} = \hat{\beta}_0 + \hat{\beta}_1$ is the slope in the post-2012 regime. Heteroskedasticity- and autocorrelation-consistent standard errors with four lags following Newey and West (1987).

visual evidence in Figure 7 and the cell level pattern in Table 5 are consistent with the framework's prediction that an explicit anchor mutes the tail-loss bias mechanism, while the formal pooled Wald test in Table 6 does not reject equality once crisis controls are included.

3.6 Forecast Errors and Signal Entropy

Proposition 1 predicts a positive linkage between signal entropy and conditional dispersion of expectations. I test this directly by regressing the absolute four-quarter forecast error $|\delta_t|$ on standardized FOMC entropy and the financial tail-risk indicators:

$$|\delta_t| = \alpha + \beta_{\mathcal{E}} \widetilde{\mathcal{E}}_t + \beta_{\text{V}} \widetilde{\text{VIX}}_t + \beta_{\text{S}} \widetilde{\text{SKEW}}_t + \text{controls}_t + u_t,$$

with HAC standard errors at four lags. Table 7 reports six specifications.

Table 7: Absolute four-quarter forecast error $|\delta_t|$ regressed on tail-risk indicators.

Variable	(1)	(2)	(3)	(4)	(5)	(6)
$\mathcal{E}_t$ (z)	0.479***			0.537***	0.539***	0.470***
	(0.105)			(0.112)	(0.112)	(0.137)
VIX_t (z)		0.040		0.240*	0.261*	0.278
		(0.175)		(0.141)	(0.152)	(0.211)
SKEW_t (z)			0.040		0.069	0.121
			(0.196)		(0.128)	(0.145)
y_t^{gap}						−0.065
						(0.064)
D^{COVID}						−0.940**
						(0.433)
Constant	1.296***	1.335***	1.319***	1.341***	1.332***	1.310***
	(0.134)	(0.174)	(0.170)	(0.114)	(0.102)	(0.128)
Observations	71	71	71	71	71	71
R^2	0.295	0.001	0.002	0.339	0.345	0.367

The dependent variable is $|\delta_t| = |\text{MICH}_t - \pi_{t+4}^{\text{PCE}}|$. Financial tail-risk regressors enter as z-scores. $\mathcal{E}_t$ is quarterly cross-statement mean of word level Shannon entropy of FOMC press releases. D^{COVID} flags 2020:Q1–Q2. The GFC dummy is omitted from column (6) because FOMC entropy is missing for every quarter of 2007:Q4–2009:Q2 window, so the dummy is identically zero on estimation sample. HAC standard errors with 4 lags in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The main result is the FOMC entropy coefficient. Column (1), the bivariate regression on entropy alone, returns $\hat{\beta}_{\mathcal{E}} = 0.479$ (s.e. 0.105, $p < 0.001$) with $R^2 = 0.295$. Adding VIX in column (4) yields $\hat{\beta}_{\mathcal{E}} = 0.537$ (s.e. 0.112) and $R^2 = 0.339$. The full control specification in column (6) retains $\hat{\beta}_{\mathcal{E}} = 0.470$ (s.e. 0.137, $p = 0.001$) and an R^2 of 0.367, with the COVID dummy entering at -0.940 ($p < 0.05$), reflecting compressed forecast error magnitudes during the 2020:Q1–Q2 deflationary window. A one-standard-deviation increase in within-quarter mean FOMC statement entropy is associated with an increase in absolute forecast error of approximately half a percentage point, an order of magnitude larger than the VIX, SKEW, or output gap effects, none of which is distinguishable from zero on its own. Against the unconditional $|\delta_t|$ mean of 1.33 pp on this sample, the entropy effect is roughly a third of the average absolute error.

The framework provides a quantitative benchmark. Under the rational inattention reading of Proposition 1, the entropy elasticity of forecast error magnitudes is of order $1/(1 + \xi)$, which at $\hat{\xi} = 0.45$ would be approximately 0.69. The empirical estimate $\hat{\beta}_{\mathcal{E}} = 0.479$ lies below this value; a two-sided t-test of $H_0 : \beta_{\mathcal{E}} = 0.690$ yields $t = -2.01$ and $p = 0.044$, rejecting equality at the 5 % level. The point estimate is in the predicted direction and within the same order of magnitude as the framework value, but is statistically smaller than the framework's literal point prediction. We read this as supportive of the qualitative mapping but not of an exact parameter recovery; attenuation could reflect (i) the framework's $1/(1 + \xi)$ scaling describing an asymptotic limit that finite-sample empirics undershoot, (ii) the FOMC entropy proxy being a noisy measure of the framework's idealized signal entropy, or (iii) classical errors-in-variables attenuation bias in the regressor.

I close this section with a caveat on identification. FOMC statement entropy is not exogenous. The committee's word choice is itself an outcome of the macroeconomic conditions it confronts, so quarters of unusual realized inflation are mechanically more likely to produce longer and more entropic statements. The regression in Table 7 therefore identifies an association between entropy and forecast error magnitudes, not a causal effect. An instrumental variables strategy based on committee composition shifts (Husted et al. 2020; Hansen et al. 2018) or an event study window around scheduled FOMC meetings would discipline a causal reading. Within these limits, the result is informative. The entropy proxy carries explanatory power for forecast error magnitudes well above what the financial tail-risk indices alone deliver.

I have identified four limitations. First, the IT-adoption date is endogenous, so the structural break evidence is suggestive rather than causal. Second, the FOMC entropy regression identifies an association, not a causal channel. Third, the bias regression covariance proxy uses generated regressors, and the HAC standard errors do not apply the Pagan (1984) correction; a block bootstrap counterpart is qualitatively unchanged but quantitatively widens standard errors by approximately 15 %. Fourth, the tail-index inference uses i.i.d. bootstrap on persistent quarterly data; a moving block bootstrap of length $\lfloor n^{1/3} \rfloor$ widens the full sample 95 % Hill interval to [0.04, 0.69] but preserves the sign conclusion.

Three findings merit emphasis. First, the heavy upper tail of inflation is robust. $\hat{\xi} = 0.45$ on the full sample with a bootstrap 95 % confidence interval excluding zero, a mean excess function with the canonical Pareto signature, and a QQ deviation that confirms the upper tail is heavier than a Weibull-type suggests. The implication is that conventional Gaussian forecasting understate the probability of inflation surprises, supporting tail-risk-aware policy frameworks advocated by Adrian et al. (2019) and Ardakani and Ajina (2024).

Second, the cross-sectional pattern and the cell level results are consistent with the prediction that an explicit numerical anchor mutes the tail-loss bias. The formal pooled Wald with crisis controls does not reject equality, the cell level results do not survive strict multiple testing adjustment, and the IT adoption date is endogenous. Subject to these caveats, the qualitative pattern complements the level anchoring evidence in Bernanke and Mishkin (1997) and the disagreement-based evidence in Capistrán and Timmermann (2009) and Andrade et al. (2016), and is consistent with the prediction that a credible numerical target operates through a tail-risk channel on expectations.

Finally, the FOMC entropy coefficient in the forecast error regression is the strongest conditional regularity in our analysis, with $\hat{\beta}_{\mathcal{E}} \in [0.470, 0.539]$ (HAC $p < 0.001$) and an R^2 of 0.30–0.37. The magnitude is in the predicted range, within an order of magnitude of the framework value $1/(1 + \hat{\xi}) = 0.69$, though formally below it ($t = -2.01$ for the null of equality). The estimate is conditional on the endogeneity of FOMC text production. Subject to that limit, it provides empirical content for the rational inattention reading of FOMC communication (Sims 2003; Maćkowiak et al. 2023) and motivates the suggestion of Proposition 2. Tightening the policy signal by reducing FOMC statement entropy is a first-order channel for shrinking the upper tail of forecast error magnitudes.

4 Discussion

The framework joins extreme value theory (Embrechts et al. 1997; Resnick 2007; de Haan and Ferreira 2006), robust control under rational inattention (Hansen and Sargent 2001, 2008; Sims 2003; Maćkowiak et al. 2023; Caplin et al. 2022), and adaptive social learning (Evans and Honkapohja 2001; Branch and Evans 2017; Banerjee 1992; Banerjee et al. 2013). It yields four results.

Theorem 1 gives the worst-case forecast under Kullback–Leibler ambiguity with asymmetric tail-loss aversion. The forecast $\hat{\pi}_t^{i*}$ is the unique fixed point of an exponentially tilted measure that overweights states above μ_t, with bias $\delta_t = \mathbb{E}_{Q^*}[\pi_t] - \mathbb{E}_P[\pi_t]$ linear in γ_i/θ_i at first order. Letting $\theta_i \to \infty$ and $\gamma_i \to 0$ recovers $\mathbb{E}_P[\pi_t]$. The result sits in the multiplier preferences family of Hansen and Sargent (2001, 2008) and Hansen (2012) and the variational axiomatization of Maccheroni et al. (2006) and Strzalecki (2011). What is new is the term $\gamma_i(\pi_t - \mu_t)^+$, which generates the upward bias rather than assuming it, as in behavioral disaster-risk models (Gabaix 2019). It is a robust control counterpart to the diagnostic expectations distortion of Bordalo et al. (2018), which reaches a similar prediction from representativeness rather than ambiguity.

Proposition 1 gives the upper-tail dependence between signal entropy and expectation dispersion, $\chi(\xi, \kappa_i) = 2 - 2^{\rho}$ with $\rho = \kappa_i / \left[(1+\xi)(\kappa_i + \sigma_\eta^2)\right]$. Three implications follow. Heavier-tailed signals raise the dependence, $\partial\chi/\partial\xi > 0$. Attention attenuates it, $\partial\chi/\partial\kappa_i < 0$, which recovers the inattention channel of Maćkowiak and Wiederholt (2015) and Ardakani et al. (2025b). The joint limit $\xi \to 0$, $\kappa_i \to \infty$ gives $\chi \to 0$, tail independence under Gaussian rational expectations. The rational inattention literature (Sims 2003, 2006; Maćkowiak et al. 2023; Maćkowiak and Wiederholt 2025) works with Gaussian or near-Gaussian information structures, and the two main extensions retain that core: Matějka and McKay (2015) through multinomial-logit demand and Steiner et al. (2017) through continuous-time control. Under regularly varying inputs the entropy-dispersion link is sharper than the Gaussian bound predicts, which makes $\mathcal{E}_t$ a tail-risk amplifier (Cover and Thomas 1991; Shannon 1948; Ardakani 2026).

Theorem 2 characterizes the steady state of social transmission. Under homogeneous preferences and complete observation, expectations converge to the worst-case bias δ_t at rate $\phi(0) = 1/2$, whatever the curvature of the updating function. Boundedness of ϕ enforces contraction, so the fixed point is unique and globally attracting and divergence cannot arise. This departs from the classical cascade results (Banerjee 1992; Bikhchandani et al. 1992; Acemoglu et al. 2011; Orléan 1995; Arifovic et al. 2025) in two ways. The dynamics are bounded, and the policy-relevant threshold is a parameter regime in which δ_t exceeds tolerance, not a dynamic instability. Social learning locks the population mean to the individual bias, so intervention works on parameters, not on the network. This refines the expectation traps reading of Arifovic et al. (2025) and connects to the constant-gain instability boundary in adaptive learning (Evans and Honkapohja 2001; Branch and Evans 2017; Eusepi and Preston 2018; Carvalho et al. 2023).

Proposition 2 closes the argument. In the location-scale family $S_t = \pi_t + \sigma_t Y$ with $Y \sim$ Frchet$(1/\xi)$, the parametrization meeting the dispersion and second moment constraints is $\sigma_t = \sqrt{\epsilon/\Gamma(1-2\xi)}$, with $\xi^* \in (0, 1/2)$ set by the dispersion constraint. Optimal signals under heavy-tailed inflation are themselves heavy-tailed. This extends the Gaussian forward guidance designs of Eggertsson and Woodford (2003), Campbell et al. (2017), Andrade et al. (2016), and Coibion et al. (2022) to environments where the tails are nonnegligible, and links the problem to the tail monitoring of Adrian et al. (2019), Ardakani and Ajina (2024), and Ardakani (2025a).

Three empirical findings follow, ordered by inferential strength. First, inflation has a heavy upper tail. The Hill estimator returns $\hat{\xi} = 0.45$ over 2000:Q1–2025:Q4, with bootstrap intervals above zero in all four subsamples and a lower bound of 0.06 in the most conservative. The mean excess function rises from 0.85 at $u = 0.5$ to a peak of 1.19 near $u = 1.5$, the signature of a Pareto-type tail (Coles 2001; Embrechts et al. 1997), and the QQ plot shows systematic upper-tail deviation. The gap between Hill and ML-GEV (0.45 against −0.11) reflects what each estimator targets. Hill uses upper order statistics, ML-GEV is dominated by central mass. The upper tail is heavy, with finite moments to order 2.2, while the body is close to Weibull. Gaussian assumptions therefore understate large positive surprises by a factor that grows in the threshold. This is the inflation counterpart to the evidence of Adrian et al. (2019), with the upper rather than the lower tail as the policy object, and with a robust control microfoundation in place of reduced-form quantile regression (Kozlowski 2024; Veldkamp 2019; Ardakani et al. 2025a).

Second, the entropy-dispersion link of Proposition 1 appears in two places. Direct tail dependence gives $\hat{\chi}(0.50) = 0.541$, $\hat{\chi}(0.85) = 0.364$, and $\hat{\chi}(0.90) = 0.429$ for entropy against VIX, with $\hat{\rho} = 0.790$ and $\hat{\chi}_{\text{logit}} = 0.270$. At 75 quarters the estimator has little power at the highest threshold, but the values stay positive throughout. The conditional mean is the stronger test. Across specifications the standardized entropy coefficient lies in $[0.470, 0.539]$ with HAC $p < 0.001$ and R^2 between 0.30 and 0.37. A one-standard-deviation rise in entropy raises the absolute forecast error by about half a percentage point, roughly a third of its unconditional mean of 1.33 pp, and an order of magnitude more than VIX or SKEW, neither of which is distinguishable from zero alone ($R^2 < 0.01$ each). The estimate $\hat{\beta}_{\mathcal{E}} = 0.479$ falls below the framework value $1/(1 + \hat{\xi}) = 0.69$, with $t = -2.01$ and $p = 0.044$. Since $1/(1 + \xi)$ is an asymptotic bound, this is a stress test rather than a rejection of the mapping. Prior work links statement complexity to market-implied uncertainty (Hansen 2012; Armantier et al. 2022); what is added here is a survey-based magnitude and a structural reading of an association that was reduced-form, complementing the inattention and inequality channel of Ardakani and Saenz (2023) and Ardakani et al. (2025b).

Third, a tentative break at the January 2012 target adoption. Across twelve inflation-by-expectations cells, ten pre-target coefficients are positive and five reach $\alpha = 0.10$, though none survives Holm–Bonferroni correction. All twelve post-target coefficients are indistinguishable from zero. The average pre-target slope is +1.40 against −0.03 after, and the scatter gives +1.50 against −0.01. The pooled Wald test conditioning on crisis dummies does not reject equality, $\chi^2(1) = 1.247$ with $p = 0.264$. That test is conservative here, since the crisis dummy partials out the high-tail-loss episodes the mechanism is about. The evidence is suggestive rather than causal. The 2012 decision responded to the FOMC's own inflation history, so a placebo break date is the natural next check. This adds tail risk to level anchoring and disagreement channels in the literature on explicit targets (Coibion et al. 2020).

Three policy implications follow. Because χ rises in ξ, optimal signal entropy falls as the inflation tail thickens. A rolling Hill estimate above roughly 0.35 marks a regime where cutting entropy buys disproportionate dispersion control; the threshold is calibrated, not derived, and corresponds to $1/(1 + \xi)$ below 0.75. When expectations are anchored and $\hat{\xi}_t$ is low, looser communication is consistent with the framework. This gives the state-contingent recommendation of Coibion et al. (2022) and Armantier et al. (2022) a structural form, with the trigger tied to a measurable statistic. Next, since intensity scales with γ_i/θ_i and low-κ_i agents are most sensitive, one statement for all audiences is suboptimal precisely because χ varies in κ_i: simpler narratives for households, whose capacity is plausibly lower (Carvalho et al. 2016; Ardakani and Saenz 2023), and denser material for market participants (Coibion et al. 2020). Last, Proposition 2 supports releasing the quantile $q_a = \pi_t + \sigma_t(-\log a)^{-\xi}$ at confidence level a, implementable as inflation fan charts with heavy-tailed spacing. Two caveats. Both σ_t and ξ are estimated, and the rolling band in Figure 5 is wide, so the released quantiles need their own confidence bands. And the design is robust to the noise distribution within the location-scale family but not to misspecifying the tail class. Central banks already run tail scenarios (European Central Bank 2024; Bank of England 2022); Proposition 2 supplies the shape parameter that should govern the spacing, alongside the information asymmetry considerations in tail-risk markets (Ardakani 2025c).

One conceptual question remains open. The framework attributes tail dispersion to heterogeneous information processing. Low-κ_i agents cannot process the signal in high-entropy episodes, so their expectations scatter. An alternative reads high entropy not as noise but as the message. The Fed is saying uncertainty is high, and agents who understand this wait, since committing early forfeits the option value of learning. Dispersion then comes from agents whose incentives force a position before uncertainty resolves, and the heterogeneity is in incentives, not capacity. Both readings imply a positive entropy-by-tail-index interaction in upper quantile regressions, stronger extremal dependence in crises, and heterogeneous responses across attention proxies. With $T = 90$ quarters the cross-sectional variance of density forecasts cannot separate them. Individual data can. In the SPF panel, one would estimate a panelist-level inattention proxy $\hat{\kappa}_i$ from how forecast revisions respond to public information, sort panelists into capacity classes on an ex-ante criterion, and test whether tail expectations concentrate among low-capacity agents or among institutionally constrained ones identified by firm type and reporting deadlines. The Michigan survey and the New York Fed's Survey of Consumer Expectations give complementary household cross-sections. That design is left to a companion paper.

Four limitations remain. Committee word choice responds to macroeconomic conditions, so entropy regression identifies an association, not causality; an instrument through committee composition (Hansen 2012; Armantier et al. 2022) or an event study around scheduled meetings would discipline the reading. The 2012 break date is endogenous, and a placebo test or a design exploiting voting patterns would strengthen the inference. The framework treats ξ_t as exogenous, though communication may shape tail risk, and a dynamic version with $\xi_t = \xi(\mathcal{E}_t, \mathcal{F}_{t-1})$ would connect to work on credibility and robust learning (Coibion et al. 2020; Ardakani 2025b). Finally, social learning here assumes a complete network; extending Theorem 2 to stochastic block models would capture realistic information structures (Banerjee et al. 2013; Acemoglu et al. 2011), and survey data could replace observational attention proxies with measurement of cognitive load (Caplin et al. 2022).

5 Conclusions

Central bank communication moves inflation expectations in the tail, not only at the mean. This paper ties the upper-tail dependence between statement entropy and expectation dispersion to two primitives, the inflation tail index and attention capacity, and derives the signal that minimizes dispersion given those primitives. Three results back this. Headline inflation has a heavy upper tail, $\hat{\xi} = 0.45$, across subsamples and estimators. Federal Reserve statement entropy is the strongest conditional predictor of household forecast error magnitudes, explaining about 30 % of the variation where VIX and SKEW explain almost none. The slope of forecast bias on tail-loss covariance flattens after the 2012 target adoption, though the pooled test does not reject and the break is not identified. Two objects follow that a central bank can measure and partly control, the rolling tail index and statement entropy. Entropy should fall as the tail index rises, communication should differ across audiences whose capacity differs, and guidance should be quantile-based with heavy-tailed spacing. What aggregate data cannot settle is whether tail dispersion comes from agents who cannot process the signal or from agents whose incentives force an early position. That question needs individual data.

Research ethics: Not applicable.
Informed consent: Not applicable.
Author contributions: The author has accepted responsibility for the entire content of this manuscript and approved its submission.
Use of Large Language Models, AI and Machine Learning Tools: None declared.
Conflict of interest: The author states no conflict of interest.
Research funding: None declared.
Data availability: The data and code are available as Supplementary Material. The panel is constructed from publicly available sources: FRED database of the Federal Reserve Bank of St. Louis (https://fred.stlouisfed.org), the Cboe SKEW index (https://www.cboe.com), and the FOMC statement archive of the Board of Governors of the Federal Reserve System (https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm).

References

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar. 2011. "Bayesian Learning in Social Networks." *Review of Economic Studies* 78 (4): 1201–36.
Adrian, T., N. Boyarchenko, and D. Giannone. 2019. "Vulnerable Growth." *American Economic Review* 109 (4): 1263–89.

Anderson, E. W., L. P. Hansen, and T. J. Sargent. 2003. "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection." *Journal of the European Economic Association* 1 (1): 68–123.

Andrade, P., R. K. Crump, S. Eusepi, and E. Moench. 2016. "Fundamental Disagreement." *Journal of Monetary Economics* 83: 106–28.

Ardakani, O. M. 2025a. "Information Content of Inflation Expectations: A Copula-Based Model." *Studies in Nonlinear Dynamics & Econometrics* 29 (1): 71–93.

Ardakani, O. M. 2025b. "Robust Learning of Tail Dependence." *Econometrics* 13 (4): 47.

Ardakani, O. M. 2025c. "Strategic Information Asymmetry in Tail-Risk Markets." *North American Journal of Economics and Finance* 79: 102460.

Ardakani, O. M. 2026. "Central Bank Signals, Behavioral Biases, and Information Flow." *Economic Modelling* 158: 107550.

Ardakani, O. M., and R. Ajina. 2024. "Tail Risks in Household Finance." *Finance Research Letters* 69: 106065.

Ardakani, O. M., and N. K. Kishor. 2018. "Examining the Success of the Central Banks in Inflation Targeting Countries: The Dynamics of the Inflation Gap and Institutional Characteristics." *Studies in Nonlinear Dynamics & Econometrics* 22 (1): 20160085.

Ardakani, O. M., and M. Saenz. 2023. "On the Comparison of Inequality Measures: Evidence from the World Values Survey." *Applied Economics Letters* 30 (21): 3051–60.

Ardakani, O. M., R. F. Bordley, and E. S. Soofi. 2025a. "Expected Information of Noisy Attribute Forecasts for Probabilistic Forecasts." *European Journal of Operational Research* 323 (3): 1013–23.

Ardakani, O. M., V. Dalko, and H. Shim. 2025b. "Information Loss from Perception Alignment." *International Review of Economics and Finance* 97: 103830.

Ardakani, O. M., N. K. Kishor, and S. Song. 2018. "Re-Evaluating the Effectiveness of Inflation Targeting." *Journal of Economic Dynamics and Control* 90: 76–97.

Arifovic, J., A. Grimaud, I. Salle, and G. Vermandel. 2025. "Social Learning and Monetary Policy at the Effective Lower Bound." *Journal of Money, Credit and Banking* 57 (2–3): 439–75.

Armantier, O., A. Sbordone, G. Topa, W. Van der Klaauw, and J. C. Williams. 2022. "A New Approach to Assess Inflation Expectations Anchoring Using Strategic Surveys." *Journal of Monetary Economics* 129: S82–S101.

Banerjee, A. V. 1992. "A Simple Model of Herd Behavior." *The Quarterly Journal of Economics* 107 (3): 797–817.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson. 2013. "The Diffusion of Microfinance." *Science* 341 (6144): 1236498.

Bank of England. 2022. "Stress Testing the UK Banking System: Key Elements of the 2022 Annual Cyclical Scenario." Technical report. London: Bank of England.

Bernanke, B. S., and F. S. Mishkin. 1997. "Inflation Targeting: A New Framework for Monetary Policy?" *Journal of Economic Perspectives* 11 (2): 97–116.

Bikhchandani, S., D. Hirshleifer, and I. Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100 (5): 992–1026.

Bordalo, P., N. Gennaioli, and A. Shleifer. 2018. "Diagnostic Expectations and Credit Cycles." *Journal of Finance* 73 (1): 199–227.

Branch, W. A., and G. W. Evans. 2017. "Unstable Inflation Targets." *Journal of Money, Credit and Banking* 49 (4): 767–806.

Campbell, J. R., J. D. M. Fisher, A. Justiniano, and L. Melosi. 2017. "Forward Guidance and Macroeconomic Outcomes Since the Financial Crisis." *NBER Macroeconomics Annual* 31: 283–357.

Capistrán, C., and A. Timmermann. 2009. "Disagreement and Biases in Inflation Expectations." *Journal of Money, Credit and Banking* 41 (2–3): 365–96.

Caplin, A., M. Dean, and J. Leahy. 2022. "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy." *Journal of Political Economy* 130 (6): 1676–715.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston. 2023. "Anchored Inflation Expectations." *American Economic Journal: Macroeconomics* 15 (1): 1–47.

Carvalho, L. S., S. Meier, and S. W. Wang. 2016. "Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday." *American Economic Review* 106 (2): 260–84.

Cogley, T., and T. J. Sargent. 2005. "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US." *Review of Economic Dynamics* 8 (2): 262–302.

Coibion, O., and Y. Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105 (8): 2644–78.

Coibion, O., Y. Gorodnichenko, and T. Ropele. 2020. "Inflation Expectations and Firm Decisions: New Causal Evidence." *The Quarterly Journal of Economics* 135 (1): 165–219.

Coibion, O., Y. Gorodnichenko, and M. Weber. 2022. "Monetary Policy Communications and Their Effects on Household Inflation Expectations." *Journal of Political Economy* 130 (6): 1537–84.

Coles, S. 2001. *An Introduction to Statistical Modeling of Extreme Values. Springer Series in Statistics*. London: Springer.

Cover, T. M., and J. A. Thomas. 1991. *Elements of Information Theory*. New York: Wiley.

de Haan, L., and A. Ferreira. 2006. *Extreme Value Theory: An Introduction*. New York: Springer.

Dekkers, A. L., J. H. Einmahl, and L. De Haan. 1989. "A Moment Estimator for the Index of an Extreme-Value Distribution." *The Annals of Statistics* 17 (4): 1833–55.

Eggertsson, G. B., and M. Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* 2003 (1): 139–211.

Embrechts, P., C. Klüppelberg, and T. Mikosch. 1997. *Modelling Extremal Events for Insurance and Finance, Volume 33 of Stochastic Modelling and Applied Probability*. Berlin: Springer.

Embrechts, P., A. J. McNeil, and D. Straumann. 2002. "Correlation and Dependence in Risk Management: Properties and Pitfalls." In *Risk Management: Value at Risk and Beyond*, edited by M. A. H. Dempster, 176–223. Cambridge University Press.

European Central Bank. 2024. "Financial Stability Review." Technical Report 2. Frankfurt am Main: European Central Bank.

Eusepi, S., and B. Preston. 2018. "The Science of Monetary Policy: An Imperfect Knowledge Perspective." *Journal of Economic Literature* 56 (1): 3–59.

Evans, G. W., and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press.

Fisher, R. A., and L. H. C. Tippett. 1928. "Limiting Forms of the Frequency Distribution of the Largest or Smallest Member of a Sample." *Mathematical Proceedings of the Cambridge Philosophical Society* 24 (2): 180–90.

Gabaix, X. 2016. "Power Laws in Economics: An Introduction." *Journal of Economic Perspectives* 30 (1): 185–206.

Gabaix, X. 2019. "Behavioral Inattention." In *Handbook of Behavioral Economics: Applications and Foundations 1*, Vol. 2, 261–343. Elsevier.

Gnedenko, B. 1943. "Sur la distribution limite du terme maximum d'une série aléatoire." *Annals of Mathematics* 44 (3): 423–53.

Hansen, L. P. 2012. "Dynamic Valuation Decomposition Within Stochastic Economies." *Econometrica* 80 (3): 911–67.

Hansen, S., and M. McMahon. 2016. "Shocking Language: Understanding the Macroeconomic Effects of Central Bank Communication." *Journal of International Economics* 99 (S1): S114–33.

Hansen, L. P., and T. J. Sargent. 2001. "Robust Control and Model Uncertainty." *American Economic Review* 91 (2): 60–6.

Hansen, L. P., and T. J. Sargent. 2008. *Robustness*. Princeton: Princeton University Press.

Hansen, S., M. McMahon, and A. Prat. 2018. "Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach." *Quarterly Journal of Economics* 133 (2): 801–70.

Hill, B. M. 1975. "A Simple General Approach to Inference About the Tail of a Distribution." *The Annals of Statistics* 3 (5): 1163–74.

Hüsler, J., and R.-D. Reiss. 1989. "Maxima of Normal Random Vectors: Between Independence and Complete Dependence." *Statistics & Probability Letters* 7 (4): 283–6.

Husted, L., J. Rogers, and B. Sun. 2020. "Monetary Policy Uncertainty." *Journal of Monetary Economics* 115: 20–36.

Kozeniauskas, N., A. Orlik, and L. Veldkamp. 2018. "What are Uncertainty Shocks?" *Journal of Monetary Economics* 100: 1–15.

Kozlowski, J. 2024. "Exploring Tail Risks in Inflation Expectations." On the Economy, Federal Reserve Bank of St. Louis.

Maccheroni, F., M. Marinacci, and A. Rustichini. 2006. "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences." *Econometrica* 74 (6): 1447–98.

Maćkowiak, B., F. Matějka, and M. Wiederholt. 2023. "Rational Inattention: A Review." *Journal of Economic Literature* 61 (1): 226–73.

Maćkowiak, B., and M. Wiederholt. 2015. "Business Cycle Dynamics Under Rational Inattention." *The Review of Economic Studies* 82 (4): 1502–32.

Maćkowiak, B., and M. Wiederholt. 2025. "Rational Inattention and the Business Cycle Effects of Productivity and News Shocks." *American Economic Journal: Macroeconomics* 17 (1): 274–309.

Mankiw, N. G., R. Reis, and J. Wolfers. 2003. "Disagreement About Inflation Expectations." *NBER Macroeconomics Annual* 18: 209–48.

Matějka, F., and A. McKay. 2015. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." *American Economic Review* 105 (1): 272–98.

Newey, W. K., and K. D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–8.

Orléan, A. 1995. "Bayesian Interactions and Collective Dynamics of Opinion: Herd Behavior and Mimetic Contagion." *Journal of Economic Behavior & Organization* 28 (2): 257–74.

Pagan, A. 1984. "Econometric Issues in the Analysis of Regressions with Generated Regressors." *International Economic Review* 25 (1): 221–47.

Patton, A. J. 2009. "Copula-Based Models for Financial Time Series." In *Handbook of Financial Time Series*, 767–85. Berlin: Springer.

Pickands, James I. 1981. "Multivariate Extreme Value Distributions." *Bulletin of the International Statistical Institute* 49 (2): 859–78.

Pickands, J., III. 1975. "Statistical Inference Using Extreme Order Statistics." *The Annals of Statistics* 3 (1): 119–31.

Poon, S.-H., M. Rockinger, and J. Tawn. 2004. "Extreme Value Dependence in Financial Markets: Diagnostics, Models, and Financial Implications." *Review of Financial Studies* 17 (2): 581–610.

Resnick, S. I. 2007. *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. Springer Series in Operations Research and Financial Engineering. New York: Springer.

Shannon, C. E. 1948. "A Mathematical Theory of Communication." *The Bell System Technical Journal* 27 (3): 379–423.

Sims, C. A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.

Sims, C. A. 2006. "Rational Inattention: Beyond the Linear-Quadratic Case." *American Economic Review* 96 (2): 158–63.

Smith, R. L. 1987. "Estimating Tails of Probability Distributions." *The Annals of Statistics* 15 (3): 1174–207.

Steiner, J., C. Stewart, and F. Matějka. 2017. "Rational Inattention Dynamics: Inertia and Delay in Decision-Making." *Econometrica* 85 (2): 521–53.

Stock, J. H., and M. W. Watson. 2007. "Why has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39 (s1): 3–33.
Strzalecki, T. 2011. "Axiomatic Foundations of Multiplier Preferences." *Econometrica* 79 (1): 47–73.
Veldkamp, L. 2019. "Beliefs and Business Cycles." *NBER Reporter* 2019 (3): 7–10.
Woodford, M. 2003. "Optimal Interest-Rate Smoothing." *The Review of Economic Studies* 70 (4): 861–86.

Supplementary Material: This article contains supplementary material (https://doi.org/10.1515/snde-2026-0060).