Investor Learning about Monetary-Policy Transmission and the Stock Market

Daniel Andrei* Michael Hasler†

July 14, 2025

Abstract

We model how investor learning about monetary-policy transmission impacts asset prices. In an asset-pricing model, investors learn from realized inflation surprises how effectively monetary policy steers future inflation. Downward revisions in perceived effectiveness raise expected inflation persistence, increasing return volatility and risk premia. These effects intensify when policy deviates significantly from neutral or monetary-transmission uncertainty is high. We estimate the model using U.S. macro and policy data from 1954 to 2023. The resulting dynamics align with observed patterns in equity returns and volatility. Empirical tests support the model's core prediction: investor learning turns central-bank credibility into a priced risk factor.

Keywords: Asset Pricing, Learning, Inflation, Monetary Policy

JEL: D51, D53, G12, G13

Dimitris Papanikolaou was the editor for this article. We thank an anonymous referee, Hengjie Ai, Justin Balthrop, Tony Berrada, Shradha Bindal, Joanne Chen, Gjergji Cici, Alexandre Corhay, Alexander David, Jerome Detemple, Alessandro Dovis, Bjørn Eraker, Mohammad Ghaderi, Umit Gurun, Dirk Hackbarth, Leyla Han, Christian Heyerdahl-Larsen, Ted Juhl, Min Kim, Steven Kou, Niklas Kroner, Mayank Kumar, Wenhao Li, Thomas Mertens, Felix Meschke, Atanas Mihov, Dmitry Orlov, Kevin Pisciotta, Michael Rebello, Max Reppen, Alejandro Rivera, Maxime Sauzet, Ivan Shaliastovich, Angel Tengulov, Andrea Vedolin, Ole Wilms, Jide Wintoki, Hao Xing, and Fernando Zapatero for helpful comments. We also thank participants at conferences and seminars: Boston University, ESSEC, Finance Innovations & AI Seminar at CIBC, University of Hamburg, University of Kansas, University of Wisconsin-Madison, 2023 EFA, 2023 HEC-McGill Workshop, 2023 SFS Cavalcade, 2023 SITE Session "New Frontiers in Asset Pricing", 2024 Adam Smith Workshop, 2024 EFMA, 2024 UBC Summer Conference, and 2024 NFA. Financial support from McGill University and the University of Neuchâtel is gratefully acknowledged.

*McGill University, Desautels Faculty of Management, 1001 Sherbrooke Street West, Office 549, Montréal, Québec H3A 1G5, Canada, daniel.andrei@mcgill.ca, danielandrei.info.

†University of Neuchâtel, Institute of Financial Analysis, Rue A.-L. Breguet 2, 2000 Neuchâtel, Switzerland, michael.hasler@unine.ch, https://sites.google.com/michaelhasler.

1 Introduction

This paper builds an asset-pricing model in which investors learn how effectively interest-rate moves pass through to inflation. We show that uncertainty about this monetary-transmission channel—and the investor learning that follows—drives equity risk premia and stock market volatility.

Supply-chain disruptions, geopolitical tensions, and trade frictions have clouded the inflation outlook—making the task of central banks more uncertain. We model this uncertainty by treating the *effectiveness* of monetary-policy transmission as a hidden coefficient. Although this coefficient is central to how policy affects inflation, it is still poorly understood (Cochrane, 2024), and its asset-pricing implications remain largely unexplored. Our approach starts from a simple premise: financial markets observe inflation and policy in real time and update their beliefs about how policy shapes inflation. These shifting beliefs can alter perceptions of central-bank credibility and feed back into asset prices.

We formalize learning about monetary policy in an asset-pricing framework. A hidden transmission coefficient measures how effectively the policy stance (the gap between the nominal rate and its neutral level) passes through to future inflation. Investors learn about this coefficient by observing policy moves and realized inflation. This creates a belief-driven loop: realized inflation affects policy; investors update their views; and those shifting views feed back into expectations, the economy, and asset prices.

Our model's key result is that monetary policy is riskiest at its extremes: learning about transmission effectiveness creates a U-shaped relationship between risk and the policy stance. When policy moves away from neutral, investors demand higher compensation for inflation risk. Inflation risk is priced because positive inflation surprises lower expected consumption growth—consistent with the "bad news" channel in Piazzesi and Schneider (2006). Crucially, learning amplifies the pricing of inflation risk in two ways. First, when investors believe transmission is weak, inflation appears more persistent, raising long-run risk and strengthening the "bad news" channel. Second, as policy moves away from neutral, learning itself *creates* risk by making investor beliefs more sensitive to inflation news.

These learning dynamics yield three implications for asset prices. First, both volatility and the equity risk premium follow a U-shaped pattern with the size of the policy stance, whether accommodative or restrictive. Second, weaker perceived transmission amplifies return volatility and risk premia by making inflation appear more persistent. Third, volatility and risk premia rise further when monetary-transmission uncertainty is high. Together, these patterns imply that central-bank credibility is itself a priced risk factor.

The core results of the paper follow directly from the model's learning dynamics. Without investor learning about policy transmission—or without uncertainty in that process—the U-shaped patterns in volatility and risk premia would not occur. The model's learning process mirrors recent empirical evidence that perceptions of monetary policy evolve systematically across policy cycles (Bauer, Pflueger, and Sunderam, 2024). In particular, inflation surprises trigger asymmetric belief updates: under restrictive policy, positive surprises lower perceived effectiveness; under accommodative policy, they raise it. These belief updates feed back into expectations and asset prices, making learning itself a source of endogenous risk.

To isolate these learning-driven effects, we work within a *monetary-learning economy*—an asset-pricing framework in which investors face uncertainty and learn how policy moves pass through to future inflation. The model contains a reduced-form "monetary block" where policy follows a Taylor rule responding to inflation and to the output gap, a well-established asset-pricing predictor (Cooper and Priestley, 2009). This structure isolates the effects of investor learning about monetary-policy transmission, produces belief dynamics consistent with observed policy perceptions, and remains tractable for deriving equilibrium asset prices and generating testable predictions.

We estimate the model's parameters via maximum likelihood using U.S. macroeconomic data (real GDP, the Federal funds rate, inflation, and the output gap) from 1954 to 2023. Notably, the estimation does not use asset price data. Yet the model matches key asset-pricing moments well, including an average real interest rate of about 1%, an average nominal rate of 4.4%, a market risk premium of 5.7%, and a return volatility of 11.8%.

We test the model's predictions using model-implied outcomes, asset-market data, and model-free proxies. We construct time series for the equity risk premium, return volatility, the price-dividend ratio, the real interest rate, and expected output growth, and examine how they move with inflation, the output gap, perceived policy transmission effectiveness, and the policy stance. The results confirm the model's core predictions. In both the model and the data, the equity risk premium and return volatility decline with perceived transmission effectiveness and rise with the squared policy stance, consistent with the model's U-shaped prediction. These effects are economically sizable and statistically significant. The data also reveal asymmetric responses to inflation surprises under tight policy, in line with the model's learning process. Predictive regressions show that perceived transmission effectiveness and the squared policy stance forecast future excess returns. Together, these results support the view that investor learning about monetary-policy transmission shapes risk and return in asset markets.

Our analysis relates to several lines of research on belief formation, monetary policy, and asset pricing. A number of studies have examined how belief formation around monetary policy shapes macroeconomic outcomes. Orphanides and Williams (2004, 2007) analyze how learning (by both private agents and policymakers) about hidden policy parameters forces central banks to adjust their rules to maintain stability. Cogley, Matthes, and Sbordone (2015) show that as agents gradually infer the Fed's reaction-function parameters, the economy's policy response shifts and reshapes policymakers' trade-offs—and that learning can amplify inflation persistence. Sargent (1982, 1999) document that shifts in central-bank credibility produce abrupt re-anchoring of inflation expectations and large real economic adjustments.[1] We extend these insights by shifting the focus from policy design to asset pricing. In our model, investors' evolving beliefs about monetary transmission drive volatility and equity risk premia. The mechanism echoes the persistence channel in Cogley et al. (2015), but works through perceived inflation risk and its pricing, not through output or policy design.

Our work contributes to the literature on asset pricing under uncertainty and learning (e.g., Veronesi, 1999; Xiong and Yan, 2010), on monetary-policy perceptions and macro-finance (e.g., Piazzesi and Schneider, 2006; Pflueger and Rinaldi, 2022; Cieslak and Povala, 2016; Bauer

[1]Additional related literature includes research on asset pricing in environments with stochastic and regime-dependent inflation dynamics (e.g., Bansal and Shaliastovich, 2013; Gallmeyer, Hollifield, Palomino, and Zin, 2007; Gil de Rubio Cruz, Osambela, Palazzo, Palomino, and Suarez, 2022; Bonelli, Palazzo, and Yamarthy, 2024), as well as models of expectation formation with information frictions (e.g., Coibion and Gorodnichenko, 2012, 2015).

et al., 2024; Ghaderi, Seo, and Shaliastovich, 2024), as well as on inflation-news processing and investor attention (Kroner, 2025). We extend this research by endogenizing perceived policy credibility: investors learn about monetary-policy transmission over time by observing inflation and interest-rate data. This belief-updating process creates a feedback loop between inflation surprises, perceived transmission effectiveness, and asset prices. The result is a new mechanism linking risk premia and volatility to central-bank credibility, and capturing how monetary-policy perceptions evolve. Our model's U-shaped asset-pricing implications align with recent findings from Ghaderi et al. (2024), who find that investors prefer moderate inflation and view both low and high inflation as risky.

The rest of the paper proceeds as follows. Section 2 introduces the monetary-learning economy. Section 3 estimates the model and evaluates its fit. Section 4 derives asset-pricing implications and presents empirical evidence. Section 5 concludes.

2 Economic Environment

We study asset pricing in a continuous-time, infinite-horizon economy with a representative investor. Our goal is to examine how equilibrium asset prices respond to investor learning about monetary-policy transmission. Instead of building a full New-Keynesian model with price and wage frictions, we use a reduced-form *monetary-learning economy*: we directly model stochastic processes for the output gap, inflation, and the interest-rate rule—our "monetary block"—and treat them as exogenous drivers of asset prices.

2.1 Monetary Block

The monetary block consists of three exogenous state variables—output gap y_t, inflation π_t, and the (hidden) monetary-transmission coefficient a_t. The nominal rate $r_{N,t}$ is determined endogenously by a Taylor-type rule (Taylor, 1993) in response to y_t and π_t.

The output gap measures the deviation of actual output from potential output. We include the output gap as a state variable due to its documented relevance for asset pricing (Cooper and Priestley, 2009). We assume the output gap is exogenous to focus solely on the investors' learning mechanism. The output gap follows a mean-reverting process:

$$dy_t = -\lambda_y y_t dt + \sigma_y dB_{y,t}. \tag{1}$$

Here, $\lambda_y > 0$ is the constant speed of mean reversion towards zero, $\sigma_y > 0$ is the constant instantaneous volatility, and $B_{y,t}$ denotes a standard one-dimensional Brownian motion.

Monetary policy follows a Taylor rule, with the nominal rate responding to current inflation π_t and the output gap y_t:

$$r_{N,t} = \bar{r}_N + \beta_\pi(\pi_t - \bar{\pi}) + \beta_y y_t. \tag{2}$$

Here, $r_{N,t}$ adjusts around its neutral level $\bar{r}_N$. The coefficient $\beta_\pi \geq 0$ governs the response to inflation deviations from the target $\bar{\pi}$, and $\beta_y \geq 0$ governs the response to the output gap.

We define the *policy stance* as the gap between the nominal rate and its neutral level:

$$\phi_t := r_{N,t} - \bar{r}_N. \tag{3}$$

It is positive when policy is restrictive, negative when accommodative, and zero when neutral. The stance ϕ_t is the key channel through which policy affects inflation.

The third element of the monetary block is inflation, which follows:

$$d\pi_t = \lambda_\pi(\pi_t^{\mathcal{A}} - \pi_t)dt + \sigma_\pi dB_{\pi,t}, \tag{4}$$

where $\lambda_\pi > 0$ is the mean-reversion rate to the (time-varying) anchor $\pi_t^{\mathcal{A}}$, $\sigma_\pi > 0$ is the constant inflation volatility, and $B_{\pi,t}$ is a standard Brownian motion independent of $B_{y,t}$.[2]

The key innovation is that the long-run anchor adjusts with policy:

$$\pi_t^{\mathcal{A}} = \bar{\pi} - a_t\phi_t, \tag{5}$$

where $\bar{\pi}$ is the central bank's target and a_t is the time-varying *transmission effectiveness* measuring how strongly policy shifts the anchor toward which inflation mean-reverts.

A large $|a_t|$ means policy transmits more forcefully into expected inflation; a small $|a_t|$ means policy has little influence. When $a_t > 0$, a restrictive stance ($\phi_t > 0$) pulls the anchor

[2]This specification for inflation dynamics can be derived from a standard New-Keynesian Phillips curve combined with the Taylor rule. The transmission coefficient a_t emerges naturally as the parameter governing how policy-rate deviations affect the long-run inflation anchor $\pi_t^{\mathcal{A}}$. See Appendix A for details.

$\pi_t^{\mathcal{A}}$ below the target $\bar{\pi}$ by $a_t\phi_t$; an easing stance ($\phi_t < 0$) pushes it above. A negative a_t implies counterproductive effects—e.g., tightening that raises the inflation anchor rather than lowering it. Importantly, while the inflation shocks $dB_{\pi,t}$ are exogenous, the *future path* of inflation is shaped endogenously by policy ϕ_t and its transmission effectiveness a_t.

The central assumption of our model is that the transmission effectiveness coefficient a_t is *unobservable*. The representative investor in this economy fully understands the structure described in (1)–(5), but never observes the true value of a_t in real time. Instead, the investor infers a_t from the history of realized inflation and nominal interest rates, $\{\pi_s, r_{N,s}\}_{s\leq t}$. This learning process is motivated by empirical evidence (e.g., Bauer et al., 2024), showing that market participants systematically revise their perceptions about monetary policy in response to observed macroeconomic outcomes and policy actions.

To allow for shifts in transmission effectiveness—whether from structural changes in the economy or shifts in the policy regime—we model a_t as a stochastic, mean-reverting process:

$$da_t = -\lambda_a a_t dt + \sigma_a dB_{a,t}, \tag{6}$$

where $\lambda_a > 0$ is the speed of mean reversion, $\sigma_a > 0$ its volatility, and $B_{a,t}$ a Brownian motion independent of $B_{y,t}$ and $B_{\pi,t}$.

Together, (1)–(6) define our reduced-form monetary block. We treat these processes as exogenous inputs in the asset-pricing analysis. What evolves endogenously is the investor's expectation of future inflation, shaped by learning about policy transmission.

2.2 Investor Learning about Monetary-Policy Transmission

Since the transmission coefficient a_t in (6) is unobservable, the investor infers it from the history of realized inflation and nominal rates, $\mathcal{F}_t^{\pi,r_N} = \{\pi_s, r_{N,s}\}_{s\leq t}$. Let $\widehat{a}_t := \mathbb{E}[a_t \mid \mathcal{F}_t^{\pi,r_N}]$ denote the posterior mean, and $\nu_{a,t} := \mathbb{E}[(a_t - \widehat{a}_t)^2 \mid \mathcal{F}_t^{\pi,r_N}]$ the posterior variance. The posterior mean $\widehat{a}_t$—the investor's *perceived transmission effectiveness*—captures the investor's belief about policy's effect on inflation. The variance $\nu_{a,t}$ measures uncertainty in that belief; we refer to it as *monetary-transmission uncertainty*.

Invoking the Kalman-Bucy filter in a linear-Gaussian setting (Liptser and Shiryaev, 2001), the investor's posterior mean and variance evolve as

$$d\widehat{a}_t = -\lambda_a \widehat{a}_t dt - \frac{\phi_t \lambda_\pi \nu_{a,t}}{\sigma_\pi} d\widehat{B}_{\pi,t}, \tag{7}$$

$$d\nu_{a,t} = \left[\sigma_a^2 - 2\lambda_a \nu_{a,t} - \left(\tfrac{\phi_t \lambda_\pi \nu_{a,t}}{\sigma_\pi}\right)^2\right] dt. \tag{8}$$

Here, $\phi_t = r_{N,t} - \bar{r}_N$ is the policy stance defined in (3), and the innovation

$$d\widehat{B}_{\pi,t} := dB_{\pi,t} + \frac{\lambda_\pi}{\sigma_\pi}(\widehat{a}_t - a_t)\phi_t dt \tag{9}$$

is a Brownian motion under $\mathcal{F}_t^{\pi,r_N}$ capturing *inflation surprises* against the investor's forecast.

Equation (7) shows that inflation surprises push $\widehat{a}_t$ away from its mean-reverting path, with updates scaled by the policy stance ϕ_t and uncertainty $\nu_{a,t}$. Equation (8) is a deterministic Riccati equation that governs the evolution of monetary-transmission uncertainty, $\nu_{a,t}$.

From the investor's perspective, inflation evolves according to:

$$d\pi_t = \lambda_\pi\left(\widehat{\pi}_t^{\mathcal{A}} - \pi_t\right)dt + \sigma_\pi d\widehat{B}_{\pi,t}, \tag{10}$$

where the *perceived* long-term inflation drift is

$$\widehat{\pi}_t^{\mathcal{A}} = \bar{\pi} - \widehat{a}_t \phi_t. \tag{11}$$

This drift, shaped by the investor's belief $\widehat{a}_t$ and the policy stance ϕ_t, drives inflation expectations. A higher $\widehat{a}_t$ implies stronger perceived transmission and greater central-bank credibility; a lower $\widehat{a}_t$ signals weaker transmission and eroding credibility.

The learning rule (7) implies that investors interpret inflation news differently depending on the policy stance. A positive inflation surprise ($d\widehat{B}_{\pi,t} > 0$) lowers $\widehat{a}_t$ when policy is restrictive ($\phi_t > 0$), but raises it when policy is accommodative ($\phi_t < 0$). This asymmetric updating aligns with recent empirical evidence from Bauer et al. (2024), who show that perceptions of monetary policy vary systematically over the policy cycle, with different interpretations of policy actions across tightening and easing phases.

Investor learning amplifies the impact of inflation news. From (7), the size of the belief update $d\widehat{a}_t$ in response to a surprise $d\widehat{B}_{\pi,t}$ scales with the product $\phi_t \nu_{a,t}$. When policy is far from neutral ($|\phi_t|$ large) and uncertainty $\nu_{a,t}$ is high, beliefs respond more sharply to news.

2.3 Investor Preferences and Consumption Growth

We embed the investor's learning process into an equilibrium asset-pricing model to study its implications for asset prices. The representative investor has Kreps-Porteus preferences (Epstein and Zin, 1989; Weil, 1990), characterized by a subjective discount rate $\rho > 0$, relative risk aversion $\gamma > 0$, and elasticity of intertemporal substitution $\psi > 0$. Following Duffie and Epstein (1992), the investor's indirect utility J_t satisfies

$$J_t = \mathbb{E}_t\Big[\int_t^{\infty} h\Big(C_s, J_s\Big) ds\Big], \tag{12}$$

where C_s is investor's consumption at time s and

$$h(C, J) = \frac{\rho}{1 - 1/\psi}\Big(\frac{C^{1-1/\psi}}{[(1-\gamma)J]^{1/\theta - 1}} - (1-\gamma)J\Big), \quad \theta = \frac{1-\gamma}{1-1/\psi}. \tag{13}$$

The restrictions $\gamma > 1/\psi$ (preference for early resolution of uncertainty) and $\psi > 1$ are both necessary to generate long-run risk effects in asset prices (Bansal and Yaron, 2004). Finally, goods-market clearing implies $C_t = \delta_t$, where δ_t is the aggregate output.

Aggregate output δ_t evolves as

$$\frac{d\delta_t}{\delta_t} = \mu_{\delta,t} dt + \sigma_\delta dB_{\delta,t}, \tag{14}$$

where $\sigma_\delta > 0$ is the constant consumption-growth volatility, and $B_{\delta,t}$ is a standard Brownian motion independent of the monetary-block shocks ($B_{y,t}$, $B_{\pi,t}$, $B_{a,t}$).

Unlike in a standard pure-exchange economy with exogenous $\mu_{\delta,t}$, our monetary-learning model determines $\mu_{\delta,t}$ endogenously to ensure equilibrium consistency.[3] The mechanism is as follows: the Taylor rule (2) sets the nominal interest rate based on y_t and π_t; this nominal rate—via the Fisher equation—must align with the real rate implied by the consumption Euler

[3]Models in the long-run risk literature (Bansal and Yaron, 2004) typically treat expected consumption growth ($\mu_{\delta,t}$) as an exogenous, persistent process. In contrast, we endogenize $\mu_{\delta,t}$ through equilibrium. This parallels Campbell and Cochrane (1999), who take consumption growth as exogenous but specify the surplus-consumption ratio's dynamics to satisfy equilibrium conditions. It allows us to capture how monetary policy and learning shape expectations of future growth.

equation, which in equilibrium implies:[4]

$$r_{R,t} = \rho + \frac{\mu_{\delta,t}}{\psi} - \frac{\gamma(1+\psi)}{2\psi}\sigma_\delta^2 - \frac{1-\theta}{2}\Big(\sigma_{W,t}^2 - \sigma_\delta^2\Big). \tag{15}$$

Here, $\sigma_{W,t}^2$ is the variance of returns on total wealth (claim to aggregate consumption) and is itself endogenous. The first two terms capture discounting and intertemporal substitution; the last two capture precautionary saving motives driven by consumption and wealth risk.

Substituting the Fisher equation, $r_{R,t} = r_{N,t} - \pi_t$, into the equilibrium condition (15) yields the equilibrium expected consumption growth rate:

$$\mu_{\delta,t} = \underbrace{\psi(r_{N,t} - \pi_t - \rho)}_{\substack{\text{Real-rate}\\ \text{channel}}} + \frac{\gamma(1+\psi)}{2}\sigma_\delta^2 + \underbrace{\psi\frac{1-\theta}{2}(\sigma_{W,t}^2 - \sigma_\delta^2)}_{\substack{\text{Wealth-risk}\\ \text{channel}}}. \tag{16}$$

Expected consumption growth $\mu_{\delta,t}$ adjusts endogenously through two channels. First, the *real-rate channel* captures intertemporal substitution: changes in real rates shift the trade-off between current and future consumption, steepening or flattening the expected path. Second, the *wealth-risk channel* reflects precautionary saving: shifts in aggregate-wealth risk ($\sigma_{W,t}^2$) alter saving motives and tilt the path accordingly. These channels embed monetary policy and uncertainty directly into $\mu_{\delta,t}$.

Elaborating on the *real-rate channel*, a positive output gap raises the nominal rate via the Taylor rule, which lifts the real rate and shifts consumption toward the future—steepening the expected path. Conversely, when inflation rises and $\beta_\pi < 1$ (as confirmed in Section 3, Table 3), the central bank offsets inflation only partially, the real rate falls, and saving is discouraged—flattening the expected path.[5]

Turning to the *wealth-risk channel*, when the variance of total-wealth returns $\sigma_{W,t}^2$ rises, precautionary motives strengthen and push up $\mu_{\delta,t}$.[6] That is, greater uncertainty steepens the

[4]Derived from the Euler equation and market clearing. See Appendix B.

[5]While a lower real rate is often viewed as short-run stimulus (Bernanke and Kuttner, 2005), in our long-run risk framework the key fact is that surprise inflation predicts lower future consumption growth—and hence higher marginal utility (Piazzesi and Schneider, 2006). Any mechanism that induces this "bad-news" sign suffices for our asset-pricing results.

[6]The final term in (16) corresponds to what Galí (2015, Ch. 3) calls a *"discount-rate shock"*, which in standard monetary models enters as an exogenous preference shifter. Here, this shock is endogenous—driven by the excess variance $\sigma_{W,t}^2 - \sigma_\delta^2$. This channel vanishes under CRRA preferences (i.e., when $\theta = 1$).

expected consumption path.

In sum, equation (16) pins down the unique expected consumption-growth rate consistent with the monetary-policy stance and investor beliefs. By embedding a monetary block in an endowment economy, the model endogenizes $\mu_{\delta,t}$: nominal policy, perceived transmission effectiveness, and transmission uncertainty jointly shape the expected consumption path. Monetary policy and beliefs about its effectiveness thus have real, nonneutral effects in equilibrium.

2.4 State-Price Density and Market Prices of Risk

The central link between the real economy and financial markets is the state-price density ξ_t. Following Duffie and Epstein (1992), it is given by

$$\xi_t = \exp\left[\int_0^t h_J(C_s, J_s)ds\right]h_C(C_t, J_t), \tag{17}$$

where $h_J(\cdot)$ and $h_C(\cdot)$ are the partial derivatives of the aggregator $h(\cdot)$ in (13) with respect to J and C.

We adopt the value function proposed by Benzoni, Collin-Dufresne, and Goldstein (2011):

$$J(C, x_t) = \frac{C^{1-\gamma}}{1-\gamma}\left[\rho e^{I(x_t)}\right]^{\theta}, \tag{18}$$

where $I(x_t)$ is the log wealth-consumption ratio and $x_t := [\pi_t \ y_t \ \widehat{a}_t \ \nu_{a,t}]^\top$ is the state vector.

The log wealth-consumption ratio $I(x_t)$ depends on the state vector x_t because the four state variables $\{\pi_t, y_t, \widehat{a}_t, \nu_{a,t}\}$ shape the investor's views of future consumption. Inflation π_t and the output gap y_t determine the real interest rate, anchoring the drift of consumption growth via (16). The perceived effectiveness of policy transmission $\widehat{a}_t$ affects the persistence of inflation shocks, while its posterior variance $\nu_{a,t}$ captures the investor's uncertainty about that persistence. Since wealth is the value of a claim to the future consumption stream, $I(x_t)$ must vary with all four state variables that shape that stream.

Applying Itô's lemma to the SDF yields its dynamics, which reveal the real risk-free rate

$r_{R,t}$ and the market prices of risk m_t:

$$\frac{d\xi_t}{\xi_t} = -r_{R,t}dt - m_t^\top d\widehat{B}_t. \tag{19}$$

The drift term corresponds to the negative of the real risk-free rate $r_{R,t}$, as derived from the consumption Euler equation in (15). The diffusion term identifies the vector of market prices of risk m_t for the shocks perceived by the investor.

The relevant shocks are innovations to consumption growth ($dB_{\delta,t}$), inflation ($d\widehat{B}_{\pi,t}$), and the output gap ($dB_{y,t}$), which form the risk vector $\widehat{B}_t := [B_{\delta,t}\ \widehat{B}_{\pi,t}\ B_{y,t}]^\top$. The corresponding market prices of risk $m_t := [m_{\delta,t}\ m_{\pi,t}\ m_{y,t}]^\top$ are:[7]

$$m_{\delta,t} = \gamma\sigma_\delta, \tag{20}$$

$$m_{\pi,t} = (1-\theta)\Big(\sigma_\pi I_\pi - \tfrac{\lambda_\pi \nu_{a,t}}{\sigma_\pi}\phi_t I_{\widehat{a}}\Big), \tag{21}$$

$$m_{y,t} = (1-\theta)\sigma_y I_y, \tag{22}$$

where $I_z := \partial I/\partial z$ is the partial derivative of the log wealth-consumption ratio $I(x_t)$ with respect to $z \in \{\pi, \widehat{a}, y\}$, and ϕ_t is the monetary policy stance from (3).

The market price of consumption risk, $m_{\delta,t} = \gamma\sigma_\delta$, reflects standard compensation for fluctuations in aggregate consumption, scaled by risk aversion γ. More importantly, in our monetary-learning model, the prices of inflation and output gap risk ($m_{\pi,t}$ and $m_{y,t}$) are nonzero. These arise because the monetary block (1)–(6), together with investor beliefs (7)–(8), jointly shape real outcomes and asset prices. Both $m_{\pi,t}$ and $m_{y,t}$ are scaled by $(1-\theta)$, linking the pricing of inflation shocks ($d\widehat{B}_{\pi,t}$) and output gap shocks ($dB_{y,t}$) to Epstein-Zin preferences ($\theta \neq 1$), where the timing of uncertainty resolution matters—as in long-run risk models (Bansal and Yaron, 2004).

Inflation and output gap shocks affect real outcomes and thus marginal utility, making them priced sources of risk ($m_{\pi,t} \neq 0$ and $m_{y,t} \neq 0$). Inflation shocks move the real interest rate ($r_{N,t} - \pi_t$), altering the expected path of consumption growth $\mu_{\delta,t}$ in (16); they also affect beliefs about monetary-policy transmission via (7). These two effects give rise to the I_π and $I_{\widehat{a}}$ terms

[7] See Appendix C for the full derivation. The unobserved Brownian motion $B_{a,t}$, which drives a_t, does not enter the investor's pricing kernel because it is filtered into the observed inflation surprise process $\widehat{B}_{\pi,t}$.

in (21). Consistent with the "bad news" channel of Piazzesi and Schneider (2006)—in which unexpected inflation signals weaker future consumption growth—we expect $m_{\pi,t} < 0$.[8] This implies positive covariance between inflation shocks and marginal utility. Assets that hedge inflation—by performing well when inflation rises unexpectedly—provide insurance and earn negative risk premia; those that perform poorly when inflation rises unexpectedly are risky and require positive compensation.

Output gap shocks act through the same real-rate channel in (16). A positive shock lifts the real rate via the Taylor rule, steepening expected consumption growth. This raises the wealth-consumption ratio, so $I_y > 0$, and hence $m_{y,t} > 0$ in (22). These shocks are negatively correlated with marginal utility. Assets that gain when the gap increases are risky and earn a positive premium.

2.5 Stock Market Volatility and Equity Risk Premium

We model the aggregate equity claim with a real dividend process D_t, following the long-run risk literature (Bansal and Yaron, 2004). This setup captures two key features: dividends are more volatile and more procyclical than consumption. The instantaneous growth rate of dividends is:

$$\frac{dD_t}{D_t} = \Big[(1-\alpha)\bar{\mu}_\delta + \alpha \mu_{\delta,t}\Big]dt + \sigma_D dB_{D,t}, \tag{23}$$

where $\mu_{\delta,t}$ is the endogenous expected consumption growth rate from (16), and $\bar{\mu}_\delta := \mathbb{E}[\mu_{\delta,t}]$ is its unconditional mean. The parameter α governs how strongly dividend growth responds to changes in $\mu_{\delta,t}$—that is, its leverage (see, e.g., Abel, 1999). Dividend risk enters through the term $\sigma_D dB_{D,t}$, where $\sigma_D > 0$ is calibrated to match aggregate dividend volatility, and $B_{D,t}$ is a Brownian motion independent of all other shocks. This setup allows shocks to expected consumption growth to drive fluctuations in dividend growth—amplified by α—while matching long-run growth and volatility levels.

Since expected consumption growth $\mu_{\delta,t}$ drives dividend growth and is itself shaped by the monetary block and investor beliefs, the equilibrium log price-dividend ratio $\Pi(x_t)$ depends on

[8] A full discussion of the signs conjectured here appears in Section 2.5. These signs, along with the partial derivatives I_z in (21)–(22), are computed by solving the PDE for the log wealth-consumption ratio $I(x_t)$ using Chebyshev polynomial methods (Judd, 1998). See Appendix D for details.

Table 1: Market Prices of Risk and Quantities of Risk

Equilibrium expressions for $m_{k,t}$, the market price of each Brownian risk entering the SDF (19), and $s_{k,t}$, the corresponding quantity of risk in the asset's diffusion (24).

Source of Risk	Market Price of Risk $m_{k,t}$	Quantity of Risk $s_{k,t}$
Consumption δ	$\gamma\,\sigma_\delta$	0
Inflation π	$(1-\theta)\Big(\sigma_\pi I_\pi - \frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi}\,\phi_t\, I_{\widehat{a}}\Big)$	$\sigma_\pi\,\Pi_\pi - \frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi}\,\phi_t\,\Pi_{\widehat{a}}$
Output gap y	$(1-\theta)\,\sigma_y\, I_y$	$\sigma_y\,\Pi_y$
Dividend D	0	σ_D

the state vector $x_t := [\pi_t\ y_t\ \widehat{a}_t\ \nu_{a,t}]^\top$. Applying Itô's lemma to the stock price $S_t = D_t e^{\Pi(x_t)}$ yields the following dynamics:[9]

$$\frac{dS_t}{S_t} = \left(r_{R,t} + \mathrm{RP}_t - \frac{D_t}{S_t}\right) dt + s_{\delta,t} dB_{\delta,t} + s_{\pi,t} d\widehat{B}_{\pi,t} + s_{y,t} dB_{y,t} + s_{D_t} dB_{D,t}. \tag{24}$$

Here, $\mathrm{RP}_t = \sum_k m_{k,t} s_{k,t}$ is the equity risk premium, and $s_{k,t}$ is the diffusion loading of the stock price on the k-th shock in $[B_{\delta,t}\ \widehat{B}_{\pi,t}\ B_{y,t}\ B_{D,t}]^\top$. The loadings are:

$$s_{\delta,t} = 0, \qquad s_{\pi,t} = \sigma_\pi \Pi_\pi - \frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi} \phi_t \Pi_{\widehat{a}}, \qquad s_{y,t} = \sigma_y \Pi_y, \qquad s_{D,t} = \sigma_D, \tag{25}$$

where $\Pi_z := \partial \Pi / \partial z$ is the partial derivative of the log price-dividend ratio with respect to state variable $z \in \{\pi, y, \widehat{a}\}$.

Table 1 shows the two key inputs for asset pricing under each fundamental Brownian shock: the market price of risk $m_{k,t}$, which enters the SDF diffusion (19), and the quantity of risk $s_{k,t}$, which enters the return diffusion (24). This side-by-side format lays out the core structure: $s_{k,t}$ drives return volatility, while the interaction $m_{k,t} s_{k,t}$ determines the equity risk premium. The analysis that follows rests on the signs of key partial derivatives—specifically I_z and Π_z for $z \in \{\pi, y, \widehat{a}\}$. Table 2 summarizes the signs of these derivatives for quick reference; Appendix E provides their economic interpretation, analytical approximations, and numerical confirmation.

[9] The derivation applies Itô's lemma to $S_t = D_t e^{\Pi(x_t)}$, using the dynamics of D_t (Eq. 23) and the state-dependent function $\Pi(x_t)$. No-arbitrage implies that the expected return $\mathbb{E}_t[dR_{M,t}] = \mathbb{E}_t[dS_t/S_t] + (D_t/S_t)dt$ must satisfy $\mathbb{E}_t[dR_{M,t}] - r_{R,t} dt = \mathrm{RP}_t dt$, where $\mathrm{RP}_t = \sum_k m_{k,t} s_{k,t}$. See Appendix F for the full derivation.

Consider first consumption and output gap risks. Consumption risk is priced positively ($m_{\delta,t} > 0$) due to risk aversion ($\gamma > 0$), but its quantity of risk is zero ($s_{\delta,t} = 0$) since consumption shocks do not affect dividends or expected growth $\mu_{\delta,t}$. Output gap risk is also priced positively ($m_{y,t} > 0$), assuming a preference for early resolution ($1 - \theta > 0$). A higher output gap raises real rates via the Taylor rule and steepens the expected consumption path via (16), implying $I_y > 0$. This same mechanism boosts equity values, leading to $s_{y,t} > 0$ through $\Pi_y > 0$.

Next, consider inflation risk. A rise in inflation tends to lower the real interest rate (assuming $\beta_\pi < 1$, as validated in Section 3), which flattens the expected consumption path and worsens growth prospects. This leads to $I_\pi < 0$ and $\Pi_\pi < 0$. The effect reflects a "bad news" channel (Piazzesi and Schneider, 2006), in which unexpected inflation signals weaker future growth. As a result, both the market price of inflation risk ($\sigma_\pi I_\pi$) and its return exposure ($\sigma_\pi \Pi_\pi$) are negative. Crucially, beliefs about policy transmission—summarized by $\widehat{a}_t$—amplify this channel: when $\widehat{a}_t$ is low, inflation is seen as more persistent, increasing the negative valuation impact and raising $|I_\pi|$ and $|\Pi_\pi|$.

Finally, consider the effects that arise from learning about monetary-policy transmission, captured by the terms involving $\nu_{a,t}$ and ϕ_t in $m_{\pi,t}$ and $s_{\pi,t}$. These reflect how the sensitivity of valuations to beliefs ($I_{\widehat{a}}$ and $\Pi_{\widehat{a}}$) interacts with the policy stance. When policy is restrictive ($\phi_t > 0$), higher perceived effectiveness $\widehat{a}_t$ implies faster inflation control. Since inflation depresses valuations ($I_\pi < 0$, $\Pi_\pi < 0$), stronger control improves asset values, so $I_{\widehat{a}} > 0$ and $\Pi_{\widehat{a}} > 0$. In contrast, when policy is accommodative ($\phi_t < 0$), greater effectiveness means stronger inflation pressure—which is bad news—implying $I_{\widehat{a}} < 0$ and $\Pi_{\widehat{a}} < 0$. In both cases, $I_{\widehat{a}}$ and $\Pi_{\widehat{a}}$ take the same sign as ϕ_t, so the products $\phi_t I_{\widehat{a}}$ and $\phi_t \Pi_{\widehat{a}}$ are generally positive.

Substituting this insight into Table 1, we see that the learning terms, $-\frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi}\phi_t I_{\widehat{a}}$ in $m_{\pi,t}$ and $-\frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi}\phi_t \Pi_{\widehat{a}}$ in $s_{\pi,t}$, are negative. This is because $I_{\widehat{a}}$ and $\Pi_{\widehat{a}}$ share the sign of ϕ_t. Their magnitudes increase with both transmission uncertainty ($\nu_{a,t}$) and the strength of the policy stance ($|\phi_t|$). This yields a key result: learning amplifies both the pricing and return exposure to inflation risk. The effect is stronger when policy deviates from neutral (large $|\phi_t|$) and monetary-transmission uncertainty is elevated (high $\nu_{a,t}$).

Table 2 summarizes the signs of the key derivatives with respect to output gap, inflation, and transmission effectiveness, along with the underlying economic mechanisms that drive them under different policy regimes.

Table 2: Key Derivative Signs

Derivative	$\phi_t > 0$	$\phi_t < 0$	Economic Mechanism
I_y, Π_y	$+$	$+$	Higher output gap boosts growth and valuations.
I_π, Π_π	$-$	$-$	Higher inflation hurts growth and valuations.
$I_{\hat{a}}$, $\Pi_{\hat{a}}$	$+$	$-$	Stronger transmission valuable in tightening, detrimental in easing.

2.5.1 Stock Market Volatility

Summing the squared quantities of risk from Table 1 yields the instantaneous variance of the aggregate market return:

$$\sigma_t^2 = \sigma_D^2 + \sigma_y^2 \Pi_y^2 + \sigma_\pi^2 \Pi_\pi^2 - 2\lambda_\pi \nu_{a,t} \Pi_{\widehat{a}} \Pi_\pi \phi_t + \tfrac{\lambda_\pi^2 \nu_{a,t}^2}{\sigma_\pi^2} \Pi_{\widehat{a}}^2 \phi_t^2. \tag{26}$$

Market volatility reflects the joint effects of economic fundamentals, monetary policy, and learning. It indirectly depends on perceived policy effectiveness ($\widehat{a}_t$), which shapes Π_π, Π_y, and $\Pi_{\widehat{a}}$, and on two other state variables that enter directly: transmission uncertainty ($\nu_{a,t}$) and the policy stance (ϕ_t).

Two key amplification channels emerge. First, lower $\widehat{a}_t$ raises the sensitivity of valuations to inflation surprises—via larger $|\Pi_\pi|$ and $|\Pi_{\widehat{a}}|$—and thus increases volatility. Second, higher $\nu_{a,t}$ amplifies volatility through both the linear and quadratic terms in (26), with the latter scaling ϕ_t^2. These effects are strongest when policy is far from neutral and beliefs about transmission are highly uncertain. Learning thus magnifies the impact of inflation shocks, especially in uncertain and non-neutral policy environments.

2.5.2 Equity Risk Premium

We now turn to the equity risk premium, RP_t, which reflects how the prices of risk ($m_{k,t}$) interact with the quantities of risk ($s_{k,t}$). Formally, the premium is given by

$$\mathrm{RP}_t = \sum_k m_{k,t} s_{k,t}. \tag{27}$$

Because consumption risk does not affect dividends ($s_{\delta,t} = 0$) and dividend shocks are unpriced ($m_{D,t} = 0$), the premium reflects only output-gap and inflation risk:

$$\mathrm{RP}_t = m_{y,t} s_{y,t} + m_{\pi,t} s_{\pi,t}. \tag{28}$$

Substituting the expressions from Table 1 yields:

$$\begin{aligned} \mathrm{RP}_t = {} & \underbrace{(1-\theta)\Big(\sigma_y^2 \Pi_y I_y + \sigma_\pi^2 \Pi_\pi I_\pi\Big)}_{\text{Monetary Block}} \\ & + \underbrace{(1-\theta)\lambda_\pi \nu_{a,t}\Big(-\Pi_\pi I_{\widehat{a}} - \Pi_{\widehat{a}} I_\pi\Big)\phi_t}_{\text{Learning (Linear)}} \; + \; \underbrace{(1-\theta)\frac{\lambda_\pi^2 \nu_{a,t}^2}{\sigma_\pi^2} \Pi_{\widehat{a}} I_{\widehat{a}} \phi_t^2}_{\text{Learning (Quadratic)}}. \end{aligned} \tag{29}$$

The premium has three components. The first, labeled *Monetary Block*, reflects compensation for exposure to output gap and inflation shocks. With $1 - \theta > 0$ and signs $I_y, \Pi_y > 0$ and $I_\pi, \Pi_\pi < 0$ (Table 2), this term is strictly positive: equity is risky because it falls on inflation surprises and rises on output gap expansions. The size of this component grows when perceived transmission is weak: a low $\widehat{a}_t$ raises $|I_\pi|$ and $|\Pi_\pi|$, amplifying the premium.

The second term, labeled *Learning (Linear)*, captures how learning affects the price of risk through its interaction with the policy stance (ϕ_t) and uncertainty ($\nu_{a,t}$). It scales linearly with both. The key object, $-\Pi_\pi I_{\widehat{a}} - \Pi_{\widehat{a}} I_\pi$, changes sign with ϕ_t because both $I_{\widehat{a}}$ and $\Pi_{\widehat{a}}$ do (Table 2). As a result, the full product $\Big(-\Pi_\pi I_{\widehat{a}} - \Pi_{\widehat{a}} I_\pi\Big)\phi_t$ is mostly positive. The linear term therefore raises the premium, especially when policy deviates from neutral.

The third term, labeled *Learning (Quadratic)*, also raises the premium. Since $I_{\widehat{a}}$ and $\Pi_{\widehat{a}}$ share the sign of ϕ_t, their product $\Pi_{\widehat{a}} I_{\widehat{a}}$ is positive. The full term grows with ϕ_t^2 and $\nu_{a,t}^2$, creating a U-shaped effect: the premium rises as policy becomes more restrictive or accommodative. This reflects added risk from belief updating: when policy is far from neutral and uncertainty is high, inflation surprises trigger larger belief revisions, increasing risk and the compensation investors demand.

In sum, the equity risk premium rises in three cases: (i) when perceived transmission effectiveness $\widehat{a}_t$ is low, (ii) when policy deviates from neutral ($\phi_t \neq 0$), and (iii) when transmission uncertainty $\nu_{a,t}$ is elevated. Markets thus price not just fundamentals, but also beliefs and uncertainty about monetary-policy transmission.

3 Parameter Estimation and Model Fit

We estimate the model's parameters via maximum likelihood, using monthly U.S. data from July 1954 to December 2023 on real GDP, the Federal funds rate, CPI, and the output gap.[10] Real GDP comes from the NIPA tables; the remaining series are from the Federal Reserve Bank of St. Louis (FRED).

Figure 1 plots annualized U.S. real GDP growth, the output gap, CPI inflation, and the Federal funds rate from 1954 to 2023. The plots are segmented by vertical bands that delineate the tenures of successive Federal Reserve Chairs. Inflation shows substantial volatility, peaking in the mid-1970s (around 11%) and early 1980s (above 14%), followed by relative stability until the 2021–2022 surge. The Federal funds rate mirrors these patterns, reaching 18% in the early 1980s before trending downward, with near-zero levels in 2008-2015 and 2020-2021. High inflation—especially during the 1970s and early 1980s stagflation—often coincided with slower GDP growth. The GDP and output gap panels show cyclical behavior, with dips during the 1981-1982 recession, the 2008 financial crisis, and the 2020 COVID pandemic.

Table 3 reports the maximum likelihood estimates. We proxy $\log(\delta_{t+\Delta}/\delta_t)$, π_t, y_t, and $r_{N,t}$ using the log growth rate of real GDP, log CPI growth, the continuously compounded output gap, and the continuously compounded Federal funds rate, respectively. The estimated policy rule coefficients, β_π and β_y, indicate a stronger response of nominal rates to inflation than to output gap, consistent with prior findings (Clarida, Gali, and Gertler, 2000; Ang, Boivin, Dong, and Loo-Kung, 2011). Both inflation and the output gap display low mean-reversion speeds, suggesting persistence, whereas the perceived transmission coefficient $\widehat{a}_t$ reverts more quickly to its mean. Average inflation over the sample is 3.38%, while the average nominal interest rate is 4.44%, implying a mean real rate close to 1%. This average inflation rate notably exceeds the 2% target adopted by many central banks.

Figure 2 shows investors' estimate of the monetary-policy transmission coefficient, $\widehat{a}_t$ (solid line, left axis), alongside the historical inflation rate (dashed line, right axis). The estimate $\widehat{a}_t$

[10] The maximum likelihood procedure is detailed in Appendix G.

Figure 1: GDP Growth, Output Gap, Inflation, and Federal Funds Rate

Observed annualized U.S. real GDP growth rate (first panel), output gap (second panel), CPI inflation rate (third panel), and Federal funds rate (fourth panel). The alternating shaded bands mark the tenures of different Fed Chairs. The data is at the monthly frequency from July 1954 to December 2023.

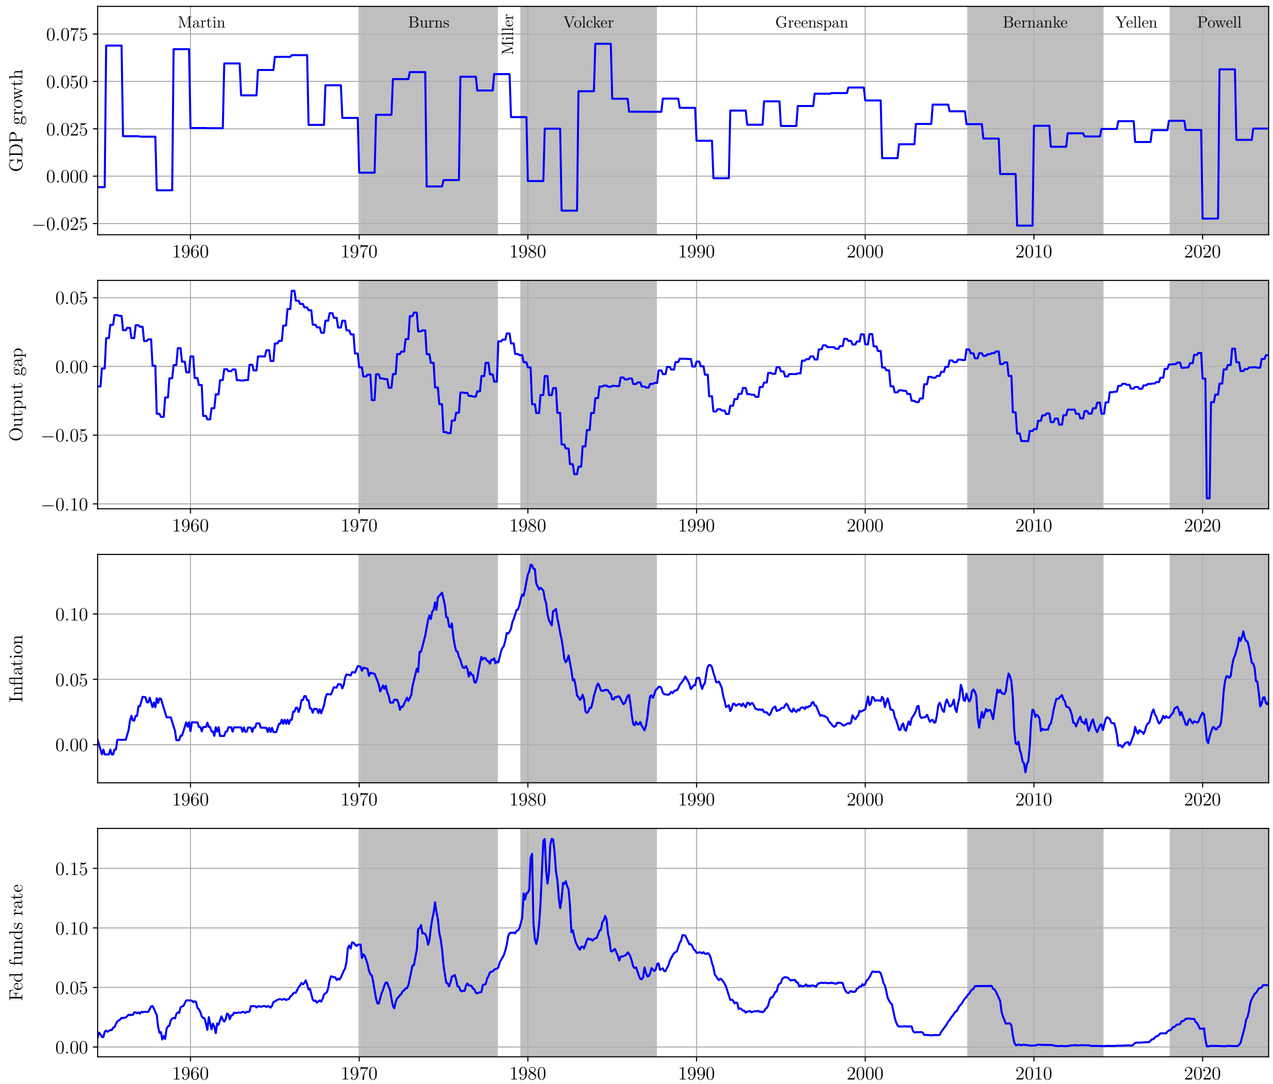

varies substantially over the sample, reflecting changes in investor beliefs about transmission effectiveness. When inflation is above its historical mean, our separate calculations show that changes in inflation and $\widehat{a}_t$ are negatively correlated (-0.48). In contrast, this correlation turns positive (0.64) when inflation is below its mean. This state-dependent correlation aligns with the learning mechanism in (7)—during high-inflation periods often met with restrictive policy ($\phi_t > 0$), positive inflation surprises lead investors to revise $\widehat{a}_t$ downward (perceiving weaker effectiveness). Conversely, during low-inflation periods potentially seeing accommodative policy

Table 3: Parameter Values Estimated by Maximum Likelihood

Parameter values estimated by maximum likelihood using monthly U.S. data from July 1954 to December 2023. The estimation procedure is detailed in Appendix G. Output data is in real terms. Standard errors are in brackets; statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

Parameter	Symbol	Estimate	Std. Error
Output growth volatility	σ_δ	0.0212***	(0.0005)
Output gap volatility	σ_y	0.0215***	(0.0007)
Output gap mean-reversion speed	λ_y	0.4219***	(0.0507)
Mean nominal interest rate	$\bar{r}_N$	0.0444***	(0.0011)
Interest rate sensitivity to inflation	β_π	0.4666***	(0.0277)
Interest rate sensitivity to output gap	β_y	0.1332***	(0.0216)
Inflation volatility	σ_π	0.0131***	(0.0011)
Mean inflation	$\bar{\pi}$	0.0338***	(0.0032)
Inflation mean-reversion speed	λ_π	0.2436***	(0.0262)
Volatility of transmission coefficient	σ_a	0.7434***	(0.1709)
Mean-reversion of transmission coefficient	λ_a	0.9365***	(0.2000)

Figure 2: Investors' Estimate of the Transmission Coefficient

Investors' estimate of the monetary-policy transmission coefficient, $\widehat{a}_t$ (solid line, left axis), and the historical inflation rate (dashed line, right axis). The alternating shaded bands mark the tenures of different Fed Chairs. The time series of $\widehat{a}_t$ is extracted from the maximum likelihood estimation. The data is at the monthly frequency from July 1954 to December 2023.

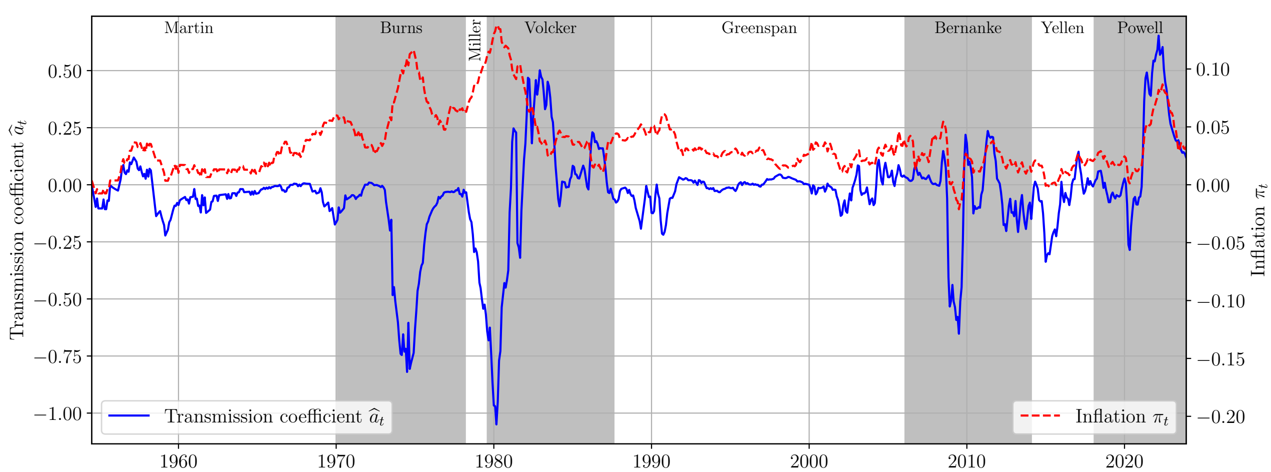

($\phi_t < 0$), negative inflation surprises (falling inflation) also induce downward revisions in $\widehat{a}_t$, generating the positive correlation.

The time series of $\widehat{a}_t$ tracks key policy episodes. For example, $\widehat{a}_t$ fell sharply during the inflationary peaks of the mid-1970s and early 1980s, reflecting investors' belief that transmission effectiveness was weakening. In contrast, $\widehat{a}_t$ rose during the latter part of Paul Volcker's

Figure 3: Monetary Policy Stance and Long-Term Inflation Drift

Monetary policy stance $\phi_t = r_{N,t} - \bar{r}_N$ (top panel) and the perceived long-term inflation drift from equation (11), $\widehat{\pi}_t^{\mathcal{A}} = \bar{\pi} - \widehat{a}_t(r_{N,t} - \bar{r}_N)$ (bottom panel). The alternating shaded bands mark the tenures of different Fed Chairs. The time series are extracted from the maximum likelihood estimation. The data is at the monthly frequency from July 1954 to December 2023.

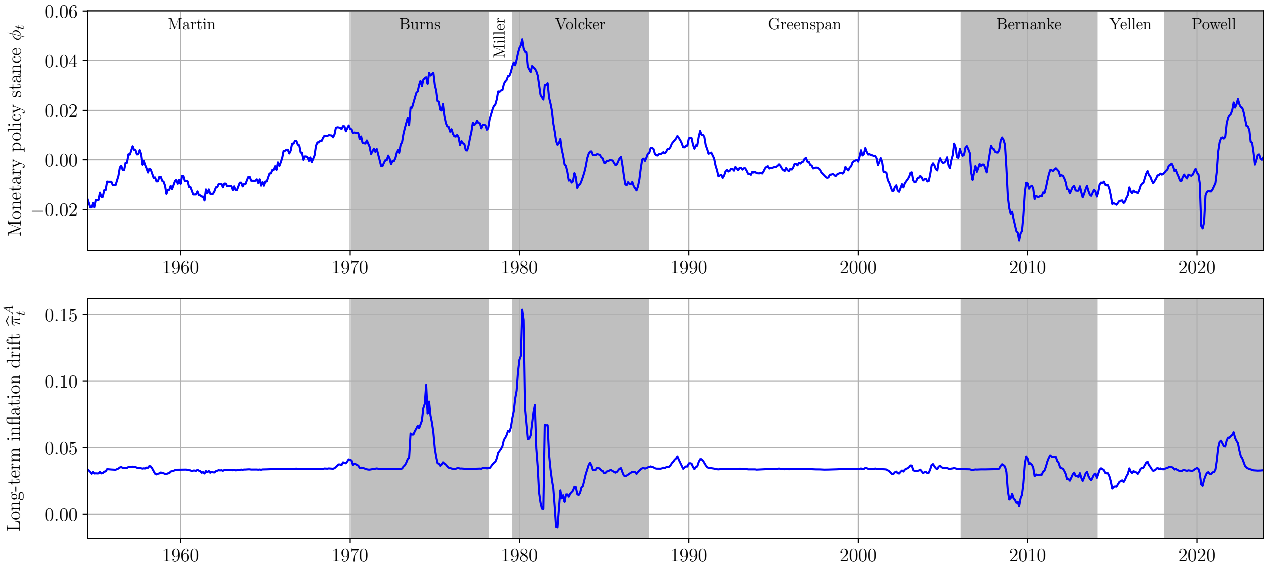

tenure (1979–87), consistent with his success in taming inflation, and remained mostly positive during Ben Bernanke's tenure (2006–14) through the global financial crisis. During the initial COVID-19 response (2020–2021), the Fed maintained accommodative policy, and positive inflation surprises in that setting lifted $\widehat{a}_t$. But as policy turned restrictive in 2022–23 to combat persistent inflation amid heightened uncertainty, $\widehat{a}_t$ declined even as tightening continued. This reflects the learning process in (7), where $\widehat{a}_t$ adjusts based on the sign of ϕ_t and inflation outcomes relative to policy actions.

Figure 3 displays the monetary policy stance, $\phi_t = r_{N,t} - \bar{r}_N$ (top panel), and the perceived long-term inflation drift, $\widehat{\pi}_t^{\mathcal{A}} = \bar{\pi} - \widehat{a}_t\phi_t$ (bottom panel), as defined in (11). A positive ϕ_t means policy is restrictive; a negative ϕ_t means it is accommodative. The Fed broadly maintained a restrictive stance from late 1965 to early 1992 and again in 2022-23, while policy was mostly accommodative from mid-1992 to late 2021. The stance is highly persistent, with autocorrelation 0.987 and annualized volatility 1.28%.

The long-term inflation drift, shaped by updates to $\widehat{a}_t$, shows persistent fluctuations (autocorrelation 0.97, annualized volatility 1.25%) and correlates positively with survey-based

Table 4: Asset-pricing Moments

Comparison of annualized empirical and model-implied asset-pricing moments. Empirical moments are calculated from monthly data on the Fed funds rate (nominal interest rate), Fed funds minus CPI inflation (real rate), and S&P 500 returns for July 1954–December 2023. We obtain model-implied moments by setting the state variables to their long-run means, $y_t = \widehat{a}_t = 0$ and $\pi_t = \bar{\pi}$.

Moment	**Data**	**Model**
Real interest rate	0.0095	0.0106
Nominal interest rate	0.0444	0.0444
Market risk premium	0.0596	0.0569
Market return volatility	0.1490	0.1178

inflation expectations.[11] Notable lows in $\widehat{\pi}_t^{\mathcal{A}}$ occurred in the early 1980s and during the Great Recession (2008-2009); highs appeared during the 1973 oil crisis, the late 1970s, and mid-2021, reflecting periods of elevated inflation pressure.

With the model's structural parameters and state variables estimated, we now calibrate the preference and dividend parameters. Following Bansal and Yaron (2004), we set the relative risk aversion to $\gamma = 10$, the elasticity of intertemporal substitution to $\psi = 2$, the subjective discount rate to $\rho = 0.0095$, and the dividend leverage to output growth to $\alpha = 2.5$. These values are standard in the long-run risk literature. Dividend growth volatility is set to $\sigma_D = 0.0575$, matching historical S&P 500 dividend volatility.

Table 4 compares empirical asset-pricing moments (first column) with model-implied moments (second column). The empirical moments use the Fed funds rate for the nominal rate, the difference between Fed funds and CPI inflation for the real rate, and the S&P 500 as the market. Model-implied moments are computed by setting the state variables to their long-run means, $y_t = \widehat{a}_t = 0$ and $\pi_t = \bar{\pi}$.

Table 4 shows close alignment between model-implied and empirical asset-pricing moments. The model produces a real interest rate of 1.06% (vs. 0.95% in the data) and a nominal rate of 4.44% (vs. 4.44%). The model-implied risk premium is 5.69%, just 27 basis points below the empirical value of 5.96%. Market return volatility shows the largest gap: 11.78% in the

[11] The long-term inflation drift, $\widehat{\pi}_t^{\mathcal{A}}$, has a 0.50 correlation with the median 5-year inflation forecast from the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters (SPF), available since 2005Q3. Regressing the SPF forecast on $\widehat{\pi}_t^{\mathcal{A}}$ yields a coefficient of 0.19, significant at the 1% level (t-stat = 3.43).

model versus 14.90% in the data—a 21% miss, which we discuss in Section 4.3. The close fit is notable, as all parameters are estimated from macro and policy data alone.

4 Asset Pricing Implications and Empirical Evidence

This section presents and tests the model's key predictions. Using the parameter estimates from Section 3, we solve the model numerically (see Appendix D), analyze how the state variables drive valuations and risk, and then test the resulting predictions empirically.

4.1 Model Predictions for Asset Prices

Figure 4 plots the expected consumption growth ($\mu_{\delta,t}$), the real interest rate ($r_{R,t}$), and the log price-dividend ratio (Π_t) against inflation (π_t) and the output gap (y_t). The figure focuses on these two state variables because they account for nearly all the variation in the model-implied paths of $\mu_{\delta,t}$, $r_{R,t}$, and Π_t.

As shown in the left panels of Figure 4, expected consumption growth ($\mu_{\delta,t}$), the real interest rate ($r_{R,t}$), and the log price-dividend ratio (Π_t) all decline as inflation (π_t) rises. Equation (16) shows that higher inflation lowers the expected growth of real consumption, reducing $\mu_{\delta,t}$. The same inflation shock also depresses the real rate through the Fisher identity, $r_{R,t} = r_{N,t} - \pi_t$. When the Taylor-rule coefficient on inflation satisfies $\beta_\pi < 1$ (as estimated in Table 3), the nominal rate $r_{N,t}$ rises by less than one-for-one with π_t, so $r_{R,t}$ falls. Finally, the log price-dividend ratio tracks $\mu_{\delta,t}$: lower expected dividend growth pushes Π_t down, so $\Pi_\pi < 0$, as shown in Section 2.5. Inflation's persistence makes it a source of long-run risk; investors who prefer early resolution of uncertainty therefore discount cash flows more heavily when π_t rises, so Π_t falls. This decline is steeper when perceived inflation persistence is higher—equivalently, when the perceived transmission effectiveness $\widehat{a}_t$ is lower—consistent with the analysis in Section 2.5.

The right panels of Figure 4 show that a higher output gap (y_t) raises expected consumption growth ($\mu_{\delta,t}$), the real interest rate ($r_{R,t}$), and the log price-dividend ratio (Π_t). Under the Taylor rule (2), a larger y_t lifts the nominal rate $r_{N,t}$. With inflation essentially unchanged on impact, the Fisher identity $r_{R,t} = r_{N,t} - \pi_t$ then implies a higher real rate. The Euler equation (16) requires a steeper expected consumption path, so $\mu_{\delta,t}$ increases. Faster expected

Figure 4: Response of Key Variables to Inflation and Output Gap

Model-implied expected consumption growth ($\mu_{\delta,t}$), real interest rate ($r_{R,t}$), and log price-dividend ratio (Π_t) plotted against inflation (π_t, left panels) and the output gap (y_t, right panels).

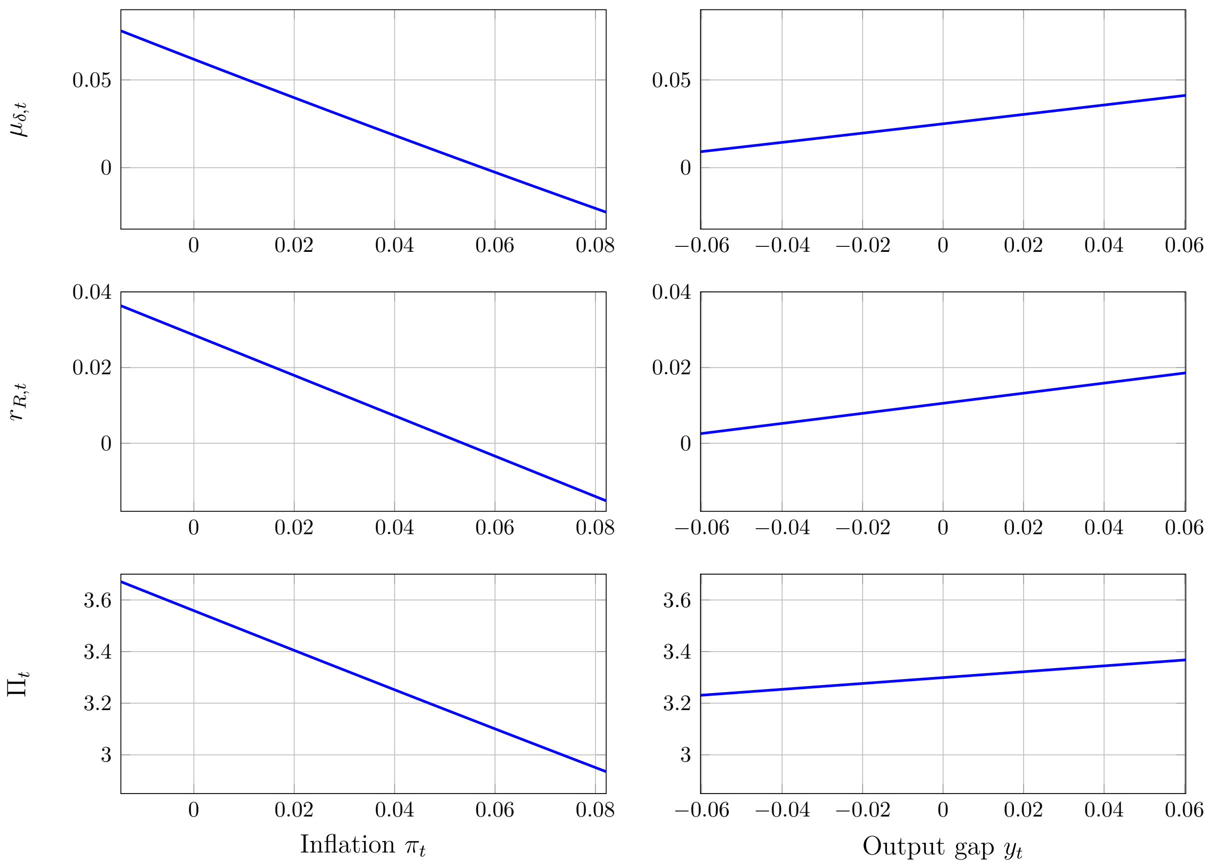

dividend growth pushes the price-dividend ratio up, so $\Pi_y > 0$.

Figure 5 shows how the market risk premium and return volatility respond to their main drivers: perceived policy transmission effectiveness ($\widehat{a}_t$), monetary policy stance (ϕ_t), and transmission uncertainty ($\nu_{a,t}$). These three drivers explain most of the variation in the model-implied risk premium and return volatility. The stance ϕ_t varies from $-\bar{r}_N$ to $+\bar{r}_N$, covering the range from the zero lower bound to aggressive tightening.

The left panels of Figure 5 show a clear downward slope: both the equity risk premium and market return volatility fall as the perceived transmission effectiveness $\widehat{a}_t$ rises. The mechanism follows Section 2.5. When investors view monetary policy as more effective, they expect inflation to revert more quickly, so the valuation elasticities $|I_\pi|$ and $|\Pi_\pi|$ shrink. Inflation becomes a weaker source of long-run risk, reducing the priced exposure $m_{\pi,t}s_{\pi,t}$ and the squared quantity

Figure 5: Market Risk Premium and Return Volatility

Model-implied risk premium and return volatility plotted against the perceived transmission effectiveness ($\widehat{a}_t$, left panels) and the monetary policy stance (ϕ_t, right panels). In the right panels, the solid lines show values with baseline calibration, while the dashed lines show values when transmission uncertainty $\nu_{a,t}$ is doubled.

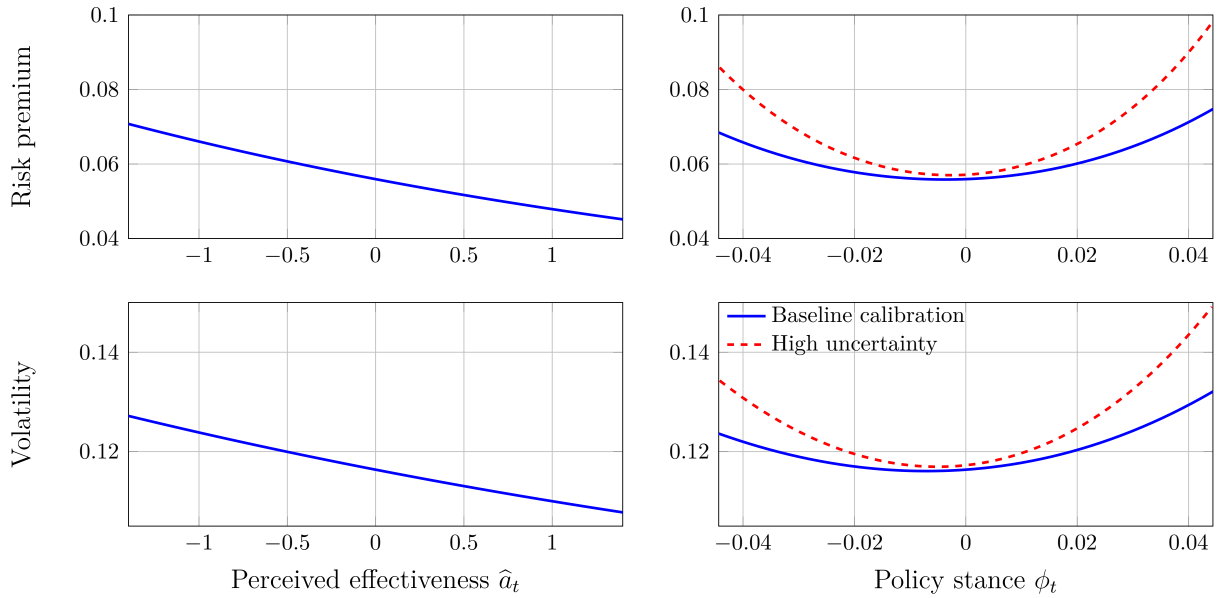

of risk Π_π^2. With less compensation demanded for inflation shocks, both the premium and volatility decline.

The right panels of Figure 5 show a U-shaped pattern: both the equity risk premium and market volatility rise as policy moves away from neutral—whether more restrictive ($\phi_t > 0$) or more accommodative ($\phi_t < 0$). In equations (26) and (29), the learning terms scale with ϕ_t^2 and with transmission uncertainty $\nu_{a,t}$. The dashed lines isolate the role of uncertainty by showing the same relationships under a higher value of $\nu_{a,t}$.[12] A non-neutral stance amplifies the impact of inflation surprises on beliefs, raising both the risk premium and market volatility.

Together, Figures 4 and 5 yield several testable predictions. (i) Higher inflation lowers expected consumption growth, real interest rates, and the price-dividend ratio; a higher output gap raises all three. (ii) When investors see policy transmission as weak (low $\widehat{a}_t$), they treat inflation as more persistent, which raises the equity risk premium and market volatility. (iii) Departures from neutrality—tighter or looser—amplify these effects through learning terms

[12]For this illustration, we double $\nu_{a,t}$ to mirror the behavior of the empirical proxy from Baker, Bloom, and Davis (2016), whose 95th percentile is 2.1 times its mean (see Section 4.2). The left panels fix $\phi_t = 0$, so $\nu_{a,t}$ plays no role; dashed lines are omitted.

that scale with $\nu_{a,t}\phi_t$.

4.2 Empirical Evidence

We test the model's predictions by regressing macro-financial outcomes on inflation, the output gap, and three core monetary state variables: perceived transmission effectiveness $\widehat{a}_t$, policy stance ϕ_t, and monetary-policy uncertainty $\nu_{a,t}$. We proceed in four steps. First, we show how expected output growth, real interest rates, and the price-dividend ratio respond to inflation and the output gap (Figure 4, Table 5). Second, we examine whether the equity risk premium and return volatility co-move with the state variables (Figure 5, Table 6). Third, we replicate these patterns using model-free proxies, allowing for interactions with monetary-policy uncertainty, to confirm that the results do not hinge on state-variable estimation (Tables 7–8). Finally, we test whether those proxies also forecast future excess returns at one-, five-, and ten-year horizons (Table 9).

We measure expected output growth as the median real GDP growth forecast from the Survey of Professional Forecasters. The real interest rate is the Fed funds rate minus CPI inflation. The price-dividend ratio is the S&P 500 index divided by trailing 12-month dividends. We use the 1-month preference-free lower bound ('LBR30') from Chabi-Yo and Loudis (2020) as our measure of the market risk premium, annualized by multiplying by twelve. This series closely tracks the premium in Martin (2016) (correlation ≈ 0.99) but spans a longer sample. Return volatility is proxied by the VIX. Model-implied counterparts are generated by feeding the estimated state variables from Section 3 into the model.

To test the relationships in Figure 4, we regress both empirical outcomes and model-implied counterparts on inflation and the output gap. Table 5 confirms the figure's patterns: expected output growth, the real interest rate, and the log price-dividend ratio all decline with inflation and rise with the output gap. Most coefficients are statistically significant, and in the model, the two variables together explain nearly all the variation.

To assess magnitudes, we rescale all effects into standard deviation units. In the model, a one-standard-deviation increase in inflation lowers expected output growth by 0.97 standard deviations, the real rate by 0.97, and the price-dividend ratio by 0.99. A similar increase in the

Table 5: Expected Output Growth, Real Interest Rate, and Log Price-Dividend Ratio vs. Inflation and the Output Gap.

Relationships between expected output growth, the real interest rate, the log price-dividend ratio, and inflation (π_t) and the output gap (y_t), in both the model and the data. Standard errors in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels. Monthly data: December 1968–December 2023 in column 2, July 1954–December 2023 otherwise.

	Expected output growth		Real interest rate		Log price-dividend ratio	
	Model	Data	Model	Data	Model	Data
π_t	-1.0395^{***}	-0.2251^{***}	-0.5334^{***}	-0.1035^{***}	-7.7655^{***}	-6.4404^{***}
	(0.0082)	(0.0800)	(0.0000)	(0.0375)	(0.0342)	(1.2024)
y_t	0.2618^{***}	0.5962^{***}	0.1332^{***}	0.1619^{*}	1.0987^{***}	1.7354
	(0.0072)	(0.2307)	(0.0000)	(0.0892)	(0.0074)	(1.5164)
R^2_{adj}	0.999	0.138	1.000	0.036	0.999	0.194
Obs.	834	661	834	834	834	834

output gap raises expected growth and the real rate by 0.21, and the price-dividend ratio by 0.12. Empirical responses have the same signs but smaller effects: inflation lowers expected growth by 0.17, the real rate by 0.11, and the price-dividend ratio by 0.42; the output gap raises them by 0.35, 0.16, and 0.10, respectively. While the model somewhat overstates the impact of inflation, its responses to the output gap align broadly with the data. All effects are economically meaningful. These patterns reflect the model's logic: inflation depresses the real rate and expected growth, pulling down equity values, while a stronger output gap lifts all three through its effect on the nominal policy rate.

To test the patterns in Figure 5, we regress the market risk premium and return volatility (in both the model and the data) on $\widehat{a}_t$, ϕ_t, and ϕ_t^2. Table 6 reports the results. Both outcomes decline with $\widehat{a}_t$ and rise with ϕ_t^2, as the model predicts. These effects are statistically significant in the model. In the data, the squared stance ϕ_t^2 is consistently significant, while $\widehat{a}_t$ becomes significant in the full specification. In the model, these variables together explain nearly all the variation in the risk premium and volatility.

The negative link with $\widehat{a}_t$ reflects the model's logic: when policy is seen as less effective, investors view inflation as more persistent, which raises long-run risk—and thus lifts risk premia and volatility. The U-shaped response to ϕ_t stems from learning. As shown in (7), inflation

Table 6: Market Risk Premium and Return Volatility vs. State Variables

Relationships between the market risk premium (Panel A), return volatility (Panel B), and the model's key state variables: perceived transmission effectiveness ($\widehat{a}_t$), policy stance (ϕ_t), and its square (ϕ_t^2), in both the model and the data. Standard errors in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels. Monthly data: July 1954–December 2023 in "Model"; January 1996–February 2023 in Panel A "Data"; January 1990–December 2023 in Panel B "Data".

	Model	Data	Model	Data	Model	Data	Model	Data
Panel A: Market risk premium								
$\widehat{a}_t$	-0.0194^{***}	-0.0374			-0.0096^{***}	-0.0201	-0.0097^{***}	-0.0520^{***}
	(0.0030)	(0.0413)			(0.0001)	(0.0216)	(0.0000)	(0.0125)
ϕ_t							0.0052	0.7680^{*}
							(0.0047)	(0.4036)
ϕ_t^2			13.8913^{***}	69.6739^{***}	9.5004^{***}	62.9804^{***}	9.3448^{***}	67.0555^{***}
			(0.2387)	(8.2166)	(0.0535)	(10.0615)	(0.1959)	(7.7403)
R^2_{adj}	0.775	0.019	0.897	0.061	0.998	0.063	0.998	0.067
Obs.	834	326	834	326	834	326	834	326
Panel B: Market return volatility								
$\widehat{a}_t$	-0.0143^{***}	-0.0627			-0.0074^{***}	-0.0199	-0.0075^{***}	-0.0785^{***}
	(0.0022)	(0.0506)			(0.0002)	(0.0199)	(0.0001)	(0.0019)
ϕ_t							0.0083^{**}	1.5599^{***}
							(0.0034)	(0.2583)
ϕ_t^2			10.0796^{***}	165.2499^{***}	6.7222^{***}	158.8168^{***}	6.4752^{***}	171.1562^{***}
			(0.1738)	(10.1481)	(0.0321)	(14.5063)	(0.0779)	(13.1232)
R^2_{adj}	0.787	0.017	0.885	0.107	0.996	0.107	0.997	0.117
Obs.	834	408	834	408	834	408	834	408

surprises reveal more about policy transmission when the stance is far from neutral, triggering stronger belief updates. These updates feed into the market risk premium and return volatility via (29) and (26). Together, these effects confirm the model's core mechanism: investor learning turns central-bank credibility into a priced risk factor, raising volatility and premia when perceived transmission is weak or policy is far from neutral.

Table 6 also reveals a small but significant positive coefficient on the linear policy stance ϕ_t. While the model emphasizes the U-shaped response from ϕ_t^2, it also generates a modest linear effect through the interaction between belief sensitivity and the stance itself. This interaction introduces an asymmetry: under tight policy ($\phi_t > 0$), a positive inflation surprise is *doubly bad news*: it signals not just higher inflation but also weaker confidence in policy effectiveness, prompting a sharper asset-price reaction. Under accommodative policy ($\phi_t < 0$), the same surprise may suggest that easing is working, leading to a milder update. This asymmetric learning amplifies volatility and premia more in restrictive cycles, helping explain the positive coefficient on ϕ_t. The mechanism aligns with the asymmetric belief updating documented by Bauer et al. (2024).

To gauge magnitudes, we rescale all effects into standard deviation units. In the model, a one-standard-deviation increase in $\widehat{a}_t$ lowers the risk premium by 0.44 standard deviations and return volatility by 0.46. A similar increase in ϕ_t raises the risk premium by 0.01 and volatility by 0.03, while a one-standard-deviation rise in ϕ_t^2 lifts them by 0.64 and 0.60, respectively. The empirical estimates show the same pattern with smaller effects: the risk premium falls by 0.21 when $\widehat{a}_t$ rises, and increases by 0.15 with ϕ_t and 0.24 with ϕ_t^2; return volatility falls by 0.17 with $\widehat{a}_t$, and rises by 0.18 with ϕ_t and 0.34 with ϕ_t^2. While the model overstates the impact of ϕ_t^2, the empirical effects remain economically meaningful.

Table 7 repeats the regressions from the "Data" columns of Table 6, replacing the estimated ϕ_t and $\widehat{a}_t$ with model-free proxies. For ϕ_t, we use demeaned inflation—the main driver of the estimated stance in the model. This proxy closely tracks ϕ_t, with correlation 0.97. To proxy $\widehat{a}_t$, we follow the logic of equation (11): we take long-term inflation expectations from the Cleveland Fed (averaged across 5- to 30-year horizons) and divide them by observed inflation. To avoid extreme values, we replace inflation with its fifth percentile whenever it falls below that threshold. The resulting proxy tracks the model-implied $\widehat{a}_t$ with correlation 0.56.

Table 7 confirms the core patterns from Table 6, using proxies instead of the estimated state variables. Both the risk premium and return volatility fall with the proxy for $\widehat{a}_t$ and rise with the squared proxy for ϕ_t. The $\widehat{a}_t$ proxy explains 3.5% of the variation in the risk premium and 4.1% in volatility; the ϕ_t^2 proxy explains 4.4% and 9.3%. These results show that the key empirical relationships hold even without relying on model-based inputs.

Table 8 tests how monetary-policy uncertainty $\nu_{a,t}$ affects the risk premium and return volatility. We estimate $\nu_{a,t}$ via maximum likelihood (Section 3) and proxy it with the monetary-policy uncertainty index of Baker et al. (2016). The dummy variables $D_{\text{high } \nu_{a,t}}$ and $D_{\text{high } \nu_{a,t}^{proxy}}$ equal one when uncertainty exceeds its median. The results largely align with the model's

Table 7: Market Risk Premium and Return Volatility vs. Model-Free Proxies

Relationships between the market risk premium (Panel A), return volatility (Panel B), and model-free proxies for perceived transmission effectiveness ($\widehat{a}_t^{proxy}$) and policy stance (ϕ_t^{proxy} and its square), based on observed macro data. Standard errors in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels. Monthly data: January 1996–February 2023 in Panel A; January 1990–December 2023 in Panel B.

Panel A: Market risk premium				
$\widehat{a}_t^{proxy}$	-0.0077^{***}		-0.0040	-0.0086
	(0.0027)		(0.0041)	(0.0083)
ϕ_t^{proxy}				0.3127
				(0.4106)
$(\phi_t^{proxy})^2$		19.4076^{***}	13.9956^{***}	8.4033
		(3.3658)	(4.6109)	(9.4781)
R^2_{adj}	0.035	0.044	0.048	0.049
Obs.	326	326	326	326
Panel B: Market return volatility				
$\widehat{a}_t^{proxy}$	-0.0150^{***}		-0.0028	-0.0085
	(0.0047)		(0.0022)	(0.0058)
ϕ_t^{proxy}				0.3795
				(0.3229)
$(\phi_t^{proxy})^2$		49.4756^{***}	45.8699^{***}	39.3831^{***}
		(8.2054)	(8.3655)	(13.5225)
R^2_{adj}	0.041	0.093	0.092	0.092
Obs.	408	408	408	408

predictions from the right panels of Figure 5: when uncertainty is high, the risk premium rises, and the U-shaped response to ϕ_t becomes steeper. For volatility, this steepening is clearly visible with both state variables and proxies. By contrast, the direct coefficient on $\nu_{a,t}$ is mixed when using state variables, but unambiguously positive with proxies. In both the model and the data, the logic is the same and follows from (7): higher uncertainty makes inflation surprises more revealing about $\widehat{a}_t$, prompting sharper belief updates that amplify volatility and risk premia—especially when policy is far from neutral.

We now turn to the model's final implication: return predictability. We test whether the same monetary-policy variables that shape risk premia and volatility also help forecast future excess returns, using proxies for perceived transmission effectiveness and the policy stance in

Table 8: Impact of Monetary Uncertainty on Market Risk Premium and Volatility

Relationships between the market risk premium, return volatility, and monetary-policy uncertainty, captured by either state variables (Panel A) or model-free proxies (Panel B). Uncertainty enters directly and via interactions with the squared policy stance. Standard errors in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels. Monthly data: January 1996–February 2023 for the risk premium; January 1990–December 2023 for volatility.

	Risk Premium		Volatility	
Panel A: State Variables				
$\widehat{a}_t$	-0.0197	-0.0337^{**}	-0.0746^{***}	-0.0804^{***}
	(0.0162)	(0.0148)	(0.0115)	(0.0088)
ϕ_t	0.0994	0.4876	1.4830^{***}	1.7805^{***}
	(0.5106)	(0.4285)	(0.2957)	(0.1872)
$(\phi_t)^2$	82.0765^{***}	74.9999^{***}	177.5218^{***}	166.6612^{***}
	(3.6722)	(4.6297)	(5.3356)	(6.4964)
$D_{\text{high } \nu_{a,t}}$	0.0139^{***}	0.0029	0.0046	-0.0108^{**}
	(0.0019)	(0.0049)	(0.0079)	(0.0054)
$(\phi_t)^2 \times D_{\text{high } \nu_{a,t}}$		298.0847^{***}		541.6875^{***}
		(104.8407)		(75.8286)
R^2_{adj}	0.077	0.097	0.116	0.137
Obs.	326	326	408	408
Panel B: Model-Free Proxies				
$\widehat{a}_t^{proxy}$	-0.0119^{*}	-0.0139^{**}	-0.0107^{*}	-0.0161^{***}
	(0.0067)	(0.0070)	(0.0063)	(0.0049)
ϕ_t^{proxy}	0.2704	0.2154	0.1616	0.0904
	(0.2826)	(0.2755)	(0.2756)	(0.2623)
$(\phi_t^{proxy})^2$	0.5696	-11.2560	33.4682^{**}	4.3499
	(6.6255)	(8.2509)	(15.3279)	(7.7628)
$D_{\text{high } \nu_{a,t}^{proxy}}$	0.0338^{***}	0.0282^{***}	0.0461^{***}	0.0353^{***}
	(0.0025)	(0.0024)	(0.0108)	(0.0113)
$(\phi_t^{proxy})^2 \times D_{\text{high } \nu_{a,t}^{proxy}}$		17.4690^{**}		40.5607^{***}
		(7.4305)		(11.5140)
R^2_{adj}	0.176	0.180	0.179	0.188
Obs.	326	326	408	408

predictive regressions. Table 9 shows that perceived transmission effectiveness helps forecast future excess returns: its proxy enters with a negative and significant coefficient at the 1- and

Table 9: Predicting Future Market Excess Returns with Model-Free Proxies

Relationships between future market excess returns at the 1-year (Panel A), 5-year (Panel B), and 10-year (Panel C) horizons and model-free proxies for the monetary-policy stance (ϕ_t^{proxy}) and the perceived transmission effectiveness ($\widehat{a}_t^{proxy}$). Standard errors are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels. Monthly data: January 1982–December 2023.

Panel A: 1-Year Future Market Excess Return				
$\widehat{a}_t^{proxy}$	-0.0222^{***}		-0.0162^{**}	0.0199
	(0.0028)		(0.0072)	(0.0219)
ϕ_t^{proxy}				-2.3460
				(2.0020)
$(\phi_t^{proxy})^2$		42.3470^{**}	23.5450	60.9133^{***}
		(21.0638)	(27.3260)	(8.5570)
R^2_{adj}	0.016	0.013	0.018	0.039
Obs.	503	503	503	503
Panel B: 5-Year Future Market Excess Return				
$\widehat{a}_t^{proxy}$	-0.0083^{**}		-0.0016	0.0023
	(0.0041)		(0.0037)	(0.0120)
ϕ_t^{proxy}				-0.2527
				(1.0288)
$(\phi_t^{proxy})^2$		28.1021^{**}	26.2369^{*}	30.2613^{**}
		(13.9292)	(14.0202)	(15.2043)
R^2_{adj}	0.009	0.027	0.025	0.024
Obs.	503	503	503	503
Panel C: 10-Year Future Market Excess Return				
$\widehat{a}_t^{proxy}$	-0.0026		0.0029	0.0037
	(0.0028)		(0.0050)	(0.0079)
ϕ_t^{proxy}				-0.0571
				(0.8277)
$(\phi_t^{proxy})^2$		18.2070^{*}	21.5489	22.4589^{**}
		(9.7081)	(15.7022)	(10.0635)
R^2_{adj}	0.000	0.018	0.018	0.016
Obs.	503	503	503	503

5-year horizons, explaining 1.6% and 0.9% of the variation, respectively. Predictive power fades at 10 years. The squared policy stance $(\phi_t^{proxy})^2$ also predicts returns, with positive and significant coefficients at all horizons. Its explanatory power peaks at five years, where it accounts for 2.7% of the variation. For both predictors, statistical strength weakens with horizon, but the patterns remain consistent.

Together, the empirical results confirm the model's core mechanism: investor learning about monetary-policy transmission shapes risk and return in asset markets. Perceived transmission effectiveness $\widehat{a}_t$, along with the level and square of the policy stance ϕ_t, predict the equity risk premium and return volatility, as the model implies. Volatility and risk premia rise when perceived transmission is weak, policy is far from neutral, and monetary-policy uncertainty is high. These effects are economically sizable and hold across model-implied outcomes, empirical data, and model-free proxies. The data also reveal asymmetric belief updating, consistent with the model and with evidence from Bauer et al. (2024): under tight policy, inflation surprises are *doubly bad* or *doubly good*, producing stronger asset-price effects than in loose regimes. These dynamics turn central-bank credibility into a priced risk factor, driving the U-shaped and asymmetric patterns in asset prices.

4.3 Model Magnitudes and Real-World Frictions

This section shows how real-world frictions may amplify the model's learning mechanism. While the baseline setup omits several key complexities, it matches key asset-pricing moments well—closely fitting the real interest rate and equity risk premium (Table 4) but understating return volatility (11.78% vs. 14.90%). To understand this gap, we turn to real-world frictions that likely hinder information flow and belief updating.

Several features of real-world markets likely slow the learning process, raising uncertainty and amplifying risk premia. We identify four such frictions: (i) inflation measurement noise, as real-time inflation data are imprecise;[13] (ii) monetary policy shocks, since not all rate changes reflect systematic responses to macro conditions;[14] (iii) policy rate observation noise, as small discrepancies between announced and effective rates complicate filtering; and (iv) structural

[13] Aruoba (2008) document quarterly revision standard deviations ranging from 0.37 to 0.85 percentage points, with noise-to-signal ratios between 0.15 and 0.33. This implies that real-time inflation data may lose 15–30% of their effective information content.

[14] Romer and Romer (2004) isolate monetary policy shocks by removing the Fed's systematic responses to its own forecasts. They find that only about 28% of funds rate changes around FOMC meetings reflect such forecast-driven responses, with the remaining 72% representing discretionary shifts averaging 39 basis points.

shocks to transmission volatility, which is constant in our model (σ_a) but likely increases during periods of political or macro disruption (e.g., tariffs, or political pressure on central banks), heightening uncertainty and slowing learning.

The impact of these frictions can be understood through the uncertainty dynamics in (8). Frictions (i)–(iii) weaken the negative quadratic term, $(\phi_t \lambda_\pi \nu_{a,t}/\sigma_\pi)^2$, slowing the decline of $\nu_{a,t}$. Friction (iv) raises σ_a^2, directly increasing the drift of $\nu_{a,t}$. Both effects lead to higher and more persistent uncertainty and, by extension, larger volatility and risk premia. As shown in the right panels of Figure 5, the dashed lines—which reflect a doubling of $\nu_{a,t}$, consistent with the empirical proxy from Baker et al. (2016), whose 95th percentile is roughly twice its mean—raise volatility and help narrow the model-data gap from Table 4.

These frictions strengthen the model's predictions without altering their qualitative nature. We omit them from the baseline to avoid introducing extra state variables, which would complicate the filtering problem without affecting the core mechanism. They offer a plausible explanation for the gap in magnitudes between model and data.

5 Conclusion

As Cochrane (2022) puts it, *"Nobody knows how interest rates affect inflation."* Our model begins with this uncertainty. We build an asset-pricing framework in which investors learn how effectively monetary policy affects inflation. Both the perceived effectiveness of this transmission and the uncertainty around it are key drivers of risk premia and return volatility. When policy is seen as effective, inflation risk fades and markets stabilize. When credibility is in doubt, learning amplifies volatility—especially when policy is far from neutral. U.S. data from 1954 to 2023 support these predictions.

The results underscore a central point: in financial markets, beliefs about monetary-policy effectiveness matter as much as policy itself. Central bank credibility shapes risk. A natural implication is that when credibility falters, restoring it may require stronger action. While we do not model this feedback explicitly—namely, the central bank adjusting policy in response to credibility concerns—it warrants further study. During the early Volcker years, for example, the Fed's credibility was in question, and aggressive rate hikes may have been needed to change beliefs. As Volcker (1979, p. 888) put it:

> *"Inflation feeds in part on itself, so part of the job of returning to a more stable and more productive economy must be to break the grip of inflationary expectations."*

That statement captures the core idea—restoring stability requires bold, credible action to re-anchor inflationary expectations. Future work could explore this mechanism further, along with extensions involving heterogeneous beliefs or time-varying investor attention.

References

Abel, A. B. (1999). Risk premia and term premia in general equilibrium. *Journal of Monetary Economics 43*(1), 3–33.

Ang, A., J. Boivin, S. Dong, and R. Loo-Kung (2011). Monetary policy shifts and the term structure. *The Review of Economic Studies 78*(2), 429–457.

Aruoba, S. B. (2008). Data revisions are not well behaved. *Journal of money, credit and banking 40*(2-3), 319–340.

Baker, S. R., N. Bloom, and S. J. Davis (2016, 07). Measuring economic policy uncertainty*. *The Quarterly Journal of Economics 131*(4), 1593–1636.

Bansal, R. and I. Shaliastovich (2013). A long-run risks explanation of predictability puzzles in bond and currency markets. *The Review of Financial Studies 26*(1), 1–33.

Bansal, R. and A. Yaron (2004). Risks for the long run: a potential resolution of asset pricing puzzles. *Journal of Finance 59*(4), 1481–1509.

Bauer, M. D., C. E. Pflueger, and A. Sunderam (2024). Perceptions about monetary policy. *The Quarterly Journal of Economics 139*(4), 2227–2278.

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein (2011). Explaining asset pricing puzzles associated with the 1987 market crash. *Journal of Financial Economics 101*(3), 552–573.

Bernanke, B. S. and K. N. Kuttner (2005). What explains the stock market's reaction to federal reserve policy? *The Journal of Finance 60*(3), 1221–1257.

Bonelli, D., B. Palazzo, and R. Yamarthy (2024). Good inflation, bad inflation:implications for risky asset prices. *Working paper*.

Calvo, G. A. (1983, September). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*(3), 383–398.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*, 205–251.

Chabi-Yo, F. and J. Loudis (2020). The conditional expected market return. *Journal of Financial Economics 137*(3), 752–786.

Cieslak, A. and P. Povala (2016). Information in the term structure of yield curve volatility. *The Journal of Finance 71*(3), 1393–1436.

Clarida, R., J. Gali, and M. Gertler (2000). Monetary policy rules and macroeconomic stability: evidence and some theory. *The Quarterly journal of economics 115*(1), 147–180.

Cochrane, J. H. (2022). Nobody knows how interest rates affect inflation. *Wall Street Journal*, A.15. Eastern edition.

Cochrane, J. H. (2024). Expectations and the neutrality of interest rates. *Review of Economic Dynamics 53*, 194–223.

Cogley, T., C. Matthes, and A. M. Sbordone (2015). Optimized taylor rules for disinflation when agents are learning. *Journal of Monetary Economics 72*, 131–147.

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

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

Cooper, I. and R. Priestley (2009). Time-varying risk premiums and the output gap. *The Review of Financial Studies 22*(7), 2801–2833.

Duffie, D. and L. G. Epstein (1992). Asset pricing with stochastic differential utility. *Review of Financial Studies 5*(3), 411–436.

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica 57*(4), 937–969.

Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Gallmeyer, M., B. Hollifield, F. Palomino, and S. E. Zin (2007). Arbitrage-free bond pricing with dynamic macroeconomic models.

Ghaderi, M., S. B. Seo, and I. Shaliastovich (2024). Learning, subjective beliefs, and time-varying preferences for different inflation ranges. *Available at SSRN*.

Gil de Rubio Cruz, A., E. Osambela, B. Palazzo, F. Palomino, and G. Suarez (2022). Inflation surprises and equity returns. Available at SSRN 4280699.

Judd, K. L. (1998). *Numerical Methods in Economics*. MIT Press.

Kroner, N. (2025). How markets process macro news: The importance of investor attention.

Liptser, R. S. and A. N. Shiryaev (2001). *Statistics of Random Processes II*. Springer Verlag, New York.

Martin, I. (2016, 10). What is the expected return on the market?*. *The Quarterly Journal of Economics 132*(1), 367–433.

Orphanides, A. and J. C. Williams (2004). Imperfect knowledge, inflation expectations, and monetary policy. In B. S. Bernanke and M. Woodford (Eds.), *The Inflation-Targeting Debate*, Chapter 5, pp. 201–246. Chicago: University of Chicago Press.

Orphanides, A. and J. C. Williams (2007). Robust monetary policy with imperfect knowledge. *Journal of monetary Economics 54*(5), 1406–1435.

Pflueger, C. and G. Rinaldi (2022). Why does the fed move markets so much? a model of monetary policy and time-varying risk aversion. *Journal of Financial Economics 146*(1), 71–89.

Piazzesi, M. and M. Schneider (2006). Equilibrium yield curves. In D. Acemoglu, K. Rogoff, and M. Woodford (Eds.), *NBER Macroeconomics Annual 2006*, Volume 21, pp. 389–472. Cambridge, MA: MIT Press.

Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American economic review 94*(4), 1055–1084.

Rotemberg, J. J. (1982, December). Sticky Prices in the United States. *Journal of Political Economy 90*(6), 1187–1211.

Sargent, T. J. (1982). The ends of four big inflations. In R. E. Hall (Ed.), *Inflation: Causes and Effects*, Chapter 2, pp. 41–98. Chicago: University of Chicago Press.

Sargent, T. J. (1999). *The Conquest of American Inflation.* Princeton, NJ: Princeton University Press.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy 39*, 195–214.

Veronesi, P. (1999). Stock market overreaction to bad news in good times: A rational expectations equilibrium model. *Review of Financial Studies 12*(5), 975–1007.

Volcker, P. A. (1979, November). Statement before the joint economic committee of the u.s. congress, october 17, 1979. *Federal Reserve Bulletin 65*, 888.

Weil, P. (1990). Nonexpected utility in macroeconomics. *The Quarterly Journal of Economics 105*(1), 29–42.

Xiong, W. and H. Yan (2010, April). Heterogeneous expectations and bond markets. *Review of Financial Studies 23*(4), 1433–1466.

A New-Keynesian Foundations for the Inflation Process

This appendix derives the reduced-form inflation dynamics (4) from standard New-Keynesian macroeconomic relationships. The derivation draws on insights from foundational New-Keynesian literature (Rotemberg, 1982; Calvo, 1983; Galí, 2015) to obtain the continuous-time inflation process with its policy-dependent anchor.

Forward-Looking Phillips Curve. We begin with a standard forward-looking Phillips curve, which relates current inflation to expected future inflation and the output gap. In deviation-from-mean form, inflation dynamics can be expressed as:

$$\pi_t^d = \beta \mathbb{E}_t[\pi_{t+1}^d] + \kappa y_t, \tag{A1}$$

where y_t is the (zero-mean) output gap, $\pi_t^d := \pi_t - \bar{\pi}$ is inflation's deviation from its mean $\bar{\pi}$, and β and κ are structural parameters reflecting time discounting and the sensitivity of inflation to economic slack, respectively.

Following Cochrane (2024), we can rewrite this discrete-time relationship in continuous time and rearrange to obtain:

$$\mathbb{E}_t[d\pi_t^d] = -\rho_\pi \pi_t^d dt - \rho_y y_t dt, \tag{A2}$$

where $\rho_\pi := (1 - \frac{1}{\beta})$ and $\rho_y := \frac{\kappa}{\beta}$. Reintroducing the mean level $\bar{\pi}$ yields an affine drift for the inflation process:

$$\mathbb{E}_t[d\pi_t] = \rho_\pi(\bar{\pi} - \pi_t)dt - \rho_y y_t dt. \tag{A3}$$

Monetary Policy and Inflation Dynamics. As in the main text, monetary policy follows the Taylor rule, linking the nominal interest rate to inflation and output deviations:

$$r_{N,t} = \bar{r}_N + \beta_\pi(\pi_t - \bar{\pi}) + \beta_y y_t, \tag{A4}$$

where $\bar{r}_N$ is the long-run nominal interest rate, while β_π and β_y capture policy responsiveness to inflation and output gap fluctuations.

We can solve for y_t from the Taylor rule:

$$y_t = \frac{1}{\beta_y}(r_{N,t} - \bar{r}_N - \beta_\pi(\pi_t - \bar{\pi})). \tag{A5}$$

Substituting this expression into the inflation drift equation:

$$\begin{aligned}
\mathbb{E}_t[d\pi_t] &= \rho_\pi(\bar{\pi} - \pi_t)dt - \rho_y \frac{1}{\beta_y}(r_{N,t} - \bar{r}_N - \beta_\pi(\pi_t - \bar{\pi}))dt && \text{(A6)}\\
&= \rho_\pi(\bar{\pi} - \pi_t)dt - \frac{\rho_y}{\beta_y}(r_{N,t} - \bar{r}_N)dt + \frac{\rho_y \beta_\pi}{\beta_y}(\pi_t - \bar{\pi})dt && \text{(A7)}\\
&= \left(\rho_\pi - \frac{\rho_y \beta_\pi}{\beta_y}\right)(\bar{\pi} - \pi_t)dt - \frac{\rho_y}{\beta_y}(r_{N,t} - \bar{r}_N)dt. && \text{(A8)}
\end{aligned}$$

Defining $\lambda_\pi := \rho_\pi - \frac{\rho_y \beta_\pi}{\beta_y}$ and introducing stochastic shocks, we obtain the inflation process used in the main text:

$$d\pi_t = \lambda_\pi(\pi_t^{\mathcal{A}} - \pi_t)dt + \sigma_\pi dB_{\pi,t}, \tag{A9}$$

where $\pi_t^{\mathcal{A}} := \bar{\pi} - a_t(r_{N,t} - \bar{r}_N)$ with $a_t := \frac{\rho_y}{\beta_y \lambda_\pi}$ representing the monetary-policy transmission effectiveness. Stability of inflation requires $\lambda_\pi > 0$. This derivation shows how our reduced-form inflation dynamics can be grounded in standard New-Keynesian relationships, while focusing attention on the key coefficient—the transmission effectiveness a_t—that investors must learn about.

B Real Risk-Free Rate Derivation

This appendix provides a step-by-step derivation of the real risk-free rate (Eq. 15) under Epstein-Zin preferences in our monetary-learning economy. We also show how this connects to the consumption Euler equation and leads to the endogenous determination of expected consumption growth.

As stated in (12)–(13), the representative investor maximizes:

$$J_t = \mathbb{E}_t \left[\int_t^\infty h(C_s, J_s) ds \right], \tag{B1}$$

where the Epstein-Zin aggregator (Duffie and Epstein, 1992) is:

$$h(C, J) = \frac{\rho}{1 - 1/\psi} \left(\frac{C^{1-1/\psi}}{[(1-\gamma)J]^{1/\theta - 1}} - (1-\gamma)J \right), \quad \theta = \frac{1-\gamma}{1-1/\psi}. \tag{B2}$$

Following Benzoni et al. (2011), we adopt the value function:

$$J(C, x_t) = \frac{C^{1-\gamma}}{1-\gamma} \left[\rho e^{I(x_t)} \right]^\theta, \tag{B3}$$

where $I(x_t)$ is the log wealth-consumption ratio, depending on the state vector $x_t = [\pi_t \ y_t \ \widehat{a}_t \ \nu_{a,t}]^\top$.

For this value function, the state-price density (SDF) is:

$$\xi_t = \exp\left[\int_0^t h_J(C_s, J_s) ds \right] h_C(C_t, J_t) \tag{B4}$$

$$= \exp\left[\int_0^t \left(\frac{\theta - 1}{e^{I(x_s)}} - \rho\theta \right) ds \right] \rho^\theta C_t^{-\gamma} e^{(\theta-1)I(x_t)}. \tag{B5}$$

Consumption Euler Equation and Real Risk-Free Rate. For any asset with return process $dR_{i,t}$, the first-order condition for optimal portfolio choice yields:

$$\mathbb{E}_t[dR_{i,t}] - r_{R,t}dt = -\mathrm{Cov}_t \left(\frac{d\xi_t}{\xi_t}, dR_{i,t} \right). \tag{B6}$$

This fundamental asset pricing equation states that the expected instantaneous real excess return on any asset i (above the real risk-free rate), $\mathbb{E}_t[dR_{i,t}] - r_{R,t}dt$, is equal to the negative of the covariance between that asset's return $dR_{i,t}$ and the growth rate of the stochastic discount factor (SDF), $d\xi_t/\xi_t$. The derivation follows from the principle of no-arbitrage, which implies that the "gain process" $\xi_t P_t$ for any traded asset price P_t must be a martingale, meaning its expected infinitesimal change is zero: $\mathbb{E}_t[d(\xi_t P_t)] = 0$. Applying Itô's product rule to the process $\xi_t P_t$, taking expectations to isolate the drift term, and setting this drift to zero yields a relationship between the expected return $\mathbb{E}_t[dR_{i,t}]/dt$, the expected SDF growth (which equals $-r_{R,t}$), and the covariance term $\mathrm{Cov}_t(d\xi_t/\xi_t, dR_{i,t})$. It is then a matter of algebra to rearrange this condition and reach (B6).

Similarly, applying the underlying martingale pricing condition $\mathbb{E}_t[d(\xi_t P_t)] = 0$ directly to the risk-free asset establishes the link between its return $r_{R,t}$ and the expected growth rate of the SDF:

$$\mathbb{E}_t\left[\frac{d\xi_t}{\xi_t}\right] = -r_{R,t}dt. \tag{B7}$$

This is the key consumption Euler equation that links expected consumption growth to the real risk-free rate. It states that the expected change in marginal utility (represented by the SDF ξ_t) equals the negative of the real rate.

Applying Itô's lemma to our SDF specification and taking expectations, we note that:

$$\frac{d\xi_t}{\xi_t} = d(\ln \xi_t) + \frac{1}{2}d\langle \ln \xi_t, \ln \xi_t\rangle_t. \tag{B8}$$

Given that the consumption shock $dB_{\delta,t}$ is independent of the monetary block shocks $(B_{y,t}, B_{\pi,t}, B_{a,t})$ in our model, we have $\mathrm{Cov}_t\left(\frac{dC_t}{C_t}, dI(x_t)\right) = 0$ (because the only Brownian driver of dC_t/C_t is $B_{\delta,t}$, which does not enter $I(x_t)$ in our specification). Applying Itô's lemma to equation (B5) yields

$$\mathbb{E}_t\left[\frac{d\xi_t}{\xi_t}\right] = \Big(\frac{\theta-1}{e^{I(x_t)}} - \rho\theta\Big)dt - \gamma\mu_{\delta,t}dt + (\theta-1)\mu_I(x_t)dt + \frac{1}{2}(\theta-1)^2\sigma_I^2(x_t)dt + \frac{1}{2}\gamma(\gamma+1)\sigma_\delta^2 dt, \tag{B9}$$

where $\mu_I(x_t)dt = \mathbb{E}_t[dI(x_t)]$ is the expected instantaneous change in the log wealth-consumption ratio, and $\sigma_I^2(x_t)dt = \mathbb{E}_t[(dI(x_t))^2]$ is the instantaneous variance of changes in the log wealth-consumption ratio. Setting this equal to $-r_{R,t}dt$ and solving for $r_{R,t}$:

$$r_{R,t} = \rho\theta - \frac{\theta-1}{e^{I(x_t)}} + \gamma\mu_{\delta,t} - (\theta-1)\mu_I(x_t) - \frac{1}{2}(\theta-1)^2\sigma_I^2(x_t) - \frac{1}{2}\gamma(\gamma+1)\sigma_\delta^2. \tag{B10}$$

This expression for the real risk-free rate depends on the unknown drift, $\mu_I(x_t)$, and variance, $\sigma_I^2(x_t)$, of the log wealth-consumption ratio. To solve for these terms and arrive at a simplified expression for $r_{R,t}$, we must characterize the equilibrium behavior of the function $I(x_t)$. The standard method in continuous-time dynamic programming is to use the Hamilton-Jacobi-Bellman (HJB) equation.

The HJB equation is the fundamental condition that the value function $J(C, x_t)$ must satisfy in equilibrium. It states that the utility flow from the aggregator, $h(C, J)$, must exactly offset the expected change in the value function, i.e., $0 = h(C_t, J_t) + \mathbb{E}_t[dJ_t]/dt$. Applying Itô's lemma to our value function $J(C, x_t) = \frac{C^{1-\gamma}}{1-\gamma}\left[\rho e^{I(x_t)}\right]^\theta$ to find its drift and substituting it into this equilibrium condition yields the following nonlinear partial differential equation (PDE) for the log wealth-consumption ratio

$I(x_t)$:

$$
\begin{aligned}
0 = e^{-I} - \rho + \frac{\gamma - 1}{\theta}\left(\frac{\gamma\sigma_\delta^2}{2} - \mu_{\delta,t}\right) + \lambda_\pi\left[\widehat{a}_t(\bar{r}_N - r_{N,t}) + \bar{\pi} - \pi_t\right] I_\pi - \lambda_a \widehat{a}_t I_{\widehat{a}} - \lambda_y y_t I_y \\
+ \frac{\sigma_\pi^2}{2} I_{\pi\pi} + \frac{(\bar{r}_N - r_{N,t})^2 \lambda_\pi^2 \nu_{a,t}^2}{2\sigma_\pi^2} I_{\widehat{a}\widehat{a}} + \frac{\sigma_y^2}{2} I_{yy} + (\bar{r}_N - r_{N,t})\lambda_\pi \nu_{a,t} I_{\pi\widehat{a}} \\
+ \frac{\theta\sigma_\pi^2}{2} I_\pi^2 + \frac{\theta(\bar{r}_N - r_{N,t})^2 \lambda_\pi^2 \nu_{a,t}^2}{2\sigma_\pi^2} I_{\widehat{a}}^2 + \frac{\theta\sigma_y^2}{2} I_y^2 + \theta(\bar{r}_N - r_{N,t})\lambda_\pi \nu_{a,t} I_\pi I_{\widehat{a}}.
\end{aligned}
\tag{B11}
$$

We simplify this HJB equation by noting that: (i) The terms involving first and second derivatives of I (on lines 1-2) correspond to $\mu_I(x_t)$, the drift of $I(x_t)$; (ii) The terms quadratic in first derivatives (line 3) equal $\frac{\theta}{2}\sigma_I^2(x_t)$, where the variance $\sigma_I^2(x_t)$ is found by applying Itô's lemma to $I(x_t)$ to determine its quadratic variation:

$$
\sigma_I^2(x_t) = \left(\sigma_\pi I_\pi - \frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi} \phi_t I_{\widehat{a}}\right)^2 + \sigma_y^2 I_y^2, \quad \text{with } \phi_t = r_{N,t} - \bar{r}_N. \tag{B12}
$$

Thus, the HJB equation (B11) becomes:

$$
0 = e^{-I} - \rho + \frac{\gamma - 1}{\theta}\left(\frac{\gamma\sigma_\delta^2}{2} - \mu_{\delta,t}\right) + \mu_I(x_t) + \frac{\theta}{2}\sigma_I^2(x_t). \tag{B13}
$$

Solving for $\mu_I(x_t)$ and substituting into our risk-free rate expression (B10), then simplifying (combining like terms and using the definition of θ), we obtain:

$$
r_{R,t} = \rho + \frac{\mu_{\delta,t}}{\psi} - \frac{\gamma(1+\psi)}{2\psi}\sigma_\delta^2 - \frac{1-\theta}{2}\sigma_I^2(x_t). \tag{B14}
$$

Noting that $\sigma_I^2(x_t) = \sigma_{W,t}^2 - \sigma_\delta^2$, where $\sigma_{W,t}^2$ is the instantaneous variance of total wealth returns:

$$
\sigma_{W,t}^2 = \sigma_\delta^2 + \left(\sigma_\pi I_\pi - \frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi} \phi_t I_{\widehat{a}}\right)^2 + \sigma_y^2 I_y^2. \tag{B15}
$$

This variance includes three components: direct consumption risk (σ_δ^2), inflation risk (including the effect of learning about policy transmission), and output gap risk.

We arrive at the standard form of the real risk-free rate in our monetary-learning economy:

$$
r_{R,t} = \rho + \frac{\mu_{\delta,t}}{\psi} - \frac{\gamma(1+\psi)}{2\psi}\sigma_\delta^2 - \frac{1-\theta}{2}(\sigma_{W,t}^2 - \sigma_\delta^2). \tag{B16}
$$

This is the standard form of the real risk-free rate in long-run risk models, which appears as (15) in the main text. Note that (15) can be readily solved for $\mu_{\delta,t}$ to obtain (16).

The nominal interest rate $r_{N,t}$ follows from the Fisher equation:

$$
r_{N,t} = r_{R,t} + \pi_t, \tag{B17}
$$

The neutral rate $\bar{r}_N$, a central component of the Taylor rule (2), must be consistent with the economy's long-run equilibrium. While a true unconditional expectation $\mathbb{E}[r_{N,t}]$ is difficult to compute analytically, we can obtain a tractable, closed-form expression for $\bar{r}_N$ by evaluating the Fisher equation at the economy's steady state. This common approach provides a consistent anchor for monetary policy.

Specifically, guided by the Fisher equation (B17), we compute the neutral rate by evaluating the expression for the real rate, $r_{R,t}$, at the long-run means of the state variables and adding the long-run mean inflation, $\bar{\pi}$. This yields:

$$\bar{r}_N = \rho + \bar{\pi} + \frac{\bar{\mu}_\delta}{\psi} - \frac{\gamma(1+\psi)}{2\psi}\sigma_\delta^2 - \frac{1-\theta}{2}\Big(\sigma_\pi^2 \bar{I}_\pi^2 + \sigma_y^2 \bar{I}_y^2\Big), \tag{B18}$$

where $\bar{\mu}_\delta$ is the long-run mean of consumption growth, and the partial derivatives $\bar{I}_\pi$ and $\bar{I}_y$ are components of the variance of the log wealth-consumption ratio, evaluated at the steady state: $\pi_t = \bar{\pi}$, $\widehat{a}_t = 0$, $y_t = 0$ (which implies $\phi_t = 0$), and $\nu_{a,t} = \bar{\nu}_a$.

C Derivation of Market Prices of Risk

This appendix derives the market prices of risk in our monetary-learning economy and explains why the Brownian motion $B_{a,t}$ does not enter the investor's pricing kernel.

From Appendix B, the state-price density (SDF) dynamics can be written as:

$$\frac{d\xi_t}{\xi_t} = -r_{R,t}dt - m_t^\top d\widehat{B}_t, \tag{C1}$$

where m_t is the vector of market prices of risk and $d\widehat{B}_t$ is the vector of Brownian innovations.

Applying Itô's lemma to (B5) and focusing on diffusion terms (since drift terms determine $r_{R,t}$ as derived in Appendix B), we get:

$$\frac{d\xi_t}{\xi_t}|_{\text{diffusion}} = -\gamma\frac{dC_t}{C_t}|_{\text{diffusion}} + (\theta-1)dI(x_t)|_{\text{diffusion}} \tag{C2}$$

Since $\frac{dC_t}{C_t}|_{\text{diffusion}} = \sigma_\delta dB_{\delta,t}$, the market price of consumption risk (Eq. (20) from the main text) is:

$$m_{\delta,t} = \gamma\sigma_\delta. \tag{C3}$$

We compute $dI(x_t)|_{\text{diffusion}}$ using Itô's lemma along with the dynamics for π_t, $\widehat{a}_t$, and y_t given in equations (10), (7), and (1), respectively. Substituting and collecting terms yields:

$$dI(x_t)|_{\text{diffusion}} = \left(\sigma_\pi I_\pi - \frac{\lambda_\pi\nu_{a,t}}{\sigma_\pi}\phi_t I_{\widehat{a}}\right) d\widehat{B}_{\pi,t} + \sigma_y I_y dB_{y,t}. \tag{C4}$$

The SDF diffusion can now be written as:

$$\frac{d\xi_t}{\xi_t}|_{\text{diffusion}} = -\gamma\sigma_\delta dB_{\delta,t} - (1-\theta)\left[\left(\sigma_\pi I_\pi - \frac{\lambda_\pi\nu_{a,t}}{\sigma_\pi}\phi_t I_{\widehat{a}}\right) d\widehat{B}_{\pi,t} + \sigma_y I_y dB_{y,t}\right]. \tag{C5}$$

Matching coefficients with $-m_t^\top d\widehat{B}_t$ yields the remaining market prices of risk:

$$m_{\pi,t} = (1-\theta)\left(\sigma_\pi I_\pi - \frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi}\phi_t I_{\widehat{a}}\right), \tag{C6}$$

$$m_{y,t} = (1-\theta)\sigma_y I_y. \tag{C7}$$

These are precisely (21)–(22) from the main text.

D Numerical Solution Method

This appendix provides details on the numerical method used to solve for the log wealth-consumption ratio $I(x_t)$.

The Hamilton-Jacobi-Bellman (HJB) equation for the log wealth-consumption ratio $I(x_t)$, where $x_t = [\pi_t \ y_t \ \widehat{a}_t \ \nu_{a,t}]^\top$, is given in (B11) in Appendix B. This equation is a second-order nonlinear partial differential equation (PDE) with non-constant coefficients.

For the numerical implementation, we make a simplifying assumption by setting $\nu_{a,t} = \breve{\nu}_a$, where $\breve{\nu}_a$ represents the empirical average of $\nu_{a,t}$. This reduces the dimensionality of the state space from four to three variables, significantly streamlining the numerical solution process.

We solve the PDE for $I(\pi_t, \widehat{a}_t, y_t)$ using the Chebyshev collocation method (Judd, 1998). The solution approach approximates the function $I(\pi_t, \widehat{a}_t, y_t)$ as:

$$I(\pi_t, \widehat{a}_t, y_t) \approx \mathcal{P}(\pi_t, \widehat{a}_t, y_t) = \sum_{i=0}^{I}\sum_{j=0}^{J}\sum_{k=0}^{K} a_{i,j,k} \times T_i[\pi] \times T_j[\widehat{a}] \times T_k[y], \tag{D1}$$

where $T_m[\cdot]$ is the Chebyshev polynomial of order m, and the coefficients $a_{i,j,k}$ are determined by the solution algorithm.

The interpolation nodes are constructed from the scaled roots of Chebyshev polynomials of orders $I+1$, $J+1$, and $K+1$. We scale these roots to cover approximately 99% of the unconditional distributions of the three mean-reverting state variables. This ensures that our approximation is accurate over the most relevant region of the state space.

The solution procedure consists of:

1. Substituting the polynomial approximation $\mathcal{P}(\pi_t, \widehat{a}_t, y_t)$ and its partial derivatives into the PDE
2. Evaluating the resulting expression at each interpolation node
3. Solving the resulting system of $(I+1)\times(J+1)\times(K+1)$ equations for the $(I+1)\times(J+1)\times(K+1)$ unknown coefficients $a_{i,j,k}$

After obtaining the coefficients, we can evaluate both the function $I(x_t)$ and its partial derivatives (such as I_π, $I_{\widehat{a}}$, and I_y) at any point within the solution domain. These partial derivatives are central to determining the market prices of risk in equations (21)–(22) and the corresponding quantities of risk that drive asset returns.

We adopt polynomial orders $I = J = K = 4$, yielding a highly accurate approximation, and confirm it through cross-platform replication in Mathematica (Wolfram Research) and Python.

The same numerical procedure is applied to solve for the market log price-dividend ratio $\Pi(x_t)$ and its partial derivatives. The PDE for the market log price-dividend ratio Π_t of the asset that is a claim to the dividend process (23) is given by:

$$\begin{aligned}
0 &= e^{-\Pi} - r_{R,t} + \bar{\mu}_\delta + \alpha\left(\mu_{\delta,t} - \bar{\mu}_\delta\right) \\
&+ \left(\lambda_\pi[\widehat{a}_t(\bar{r}_N - r_{N,t}) + \bar{\pi} - \pi_t] - m_{\pi,t}\sigma_t\right)\Pi_\pi - \left(\lambda_a \widehat{a}_t + \frac{m_{\pi,t}(\bar{r}_N - r_{N,t})\lambda_\pi \breve{\nu}_a}{\sigma_\pi}\right)\Pi_{\widehat{a}} \\
&- \left(\lambda_y y_t + m_{y,t}\sigma_y\right)\Pi_y + \frac{\sigma_\pi^2}{2}\Pi_{\pi\pi} + \frac{(\bar{r}_N - r_{N,t})^2\lambda_\pi^2\breve{\nu}_a^2}{2\sigma_\pi^2}\Pi_{\widehat{a}\widehat{a}} + \frac{\sigma_y^2}{2}\Pi_{yy} + (\bar{r}_N - r_{N,t})\lambda_\pi\breve{\nu}_a\Pi_{\pi\widehat{a}} \\
&+ \frac{\sigma_\pi^2}{2}\Pi_\pi^2 + \frac{(\bar{r}_N - r_{N,t})^2\lambda_\pi^2\breve{\nu}_a^2}{2\sigma_\pi^2}\Pi_{\widehat{a}}^2 + \frac{\sigma_y^2}{2}\Pi_y^2 + (\bar{r}_N - r_{N,t})\lambda_\pi\breve{\nu}_a\Pi_\pi\Pi_{\widehat{a}}.
\end{aligned} \tag{D2}$$

We derive this PDE by enforcing the fundamental no-arbitrage condition for a dividend-paying stock. For a stock with price $S_t = D_t e^{\Pi(x_t)}$ and dividend process D_t, this condition implies that the expected change in its discounted value, $\xi_t S_t$, over an infinitesimal interval dt, plus the discounted dividend $\xi_t D_t dt$ paid in that interval, must sum to zero:

$$\mathbb{E}_t[d(\xi_t S_t)] + \xi_t D_t dt = 0. \tag{D3}$$

Applying Itô's lemma to the product $\xi_t S_t$ to determine its expected change (its drift component multiplied by dt), and then incorporating the dividend term $\xi_t D_t dt$, we arrange all resulting terms to form the PDE (which is set to zero). The first line of this PDE contains the dividend yield ($e^{-\Pi} \equiv D_t/S_t$), the risk-free rate ($r_{R,t}$), and expected dividend growth terms ($\bar{\mu}_\delta + \alpha(\mu_{\delta,t} - \bar{\mu}_\delta)$). The remaining terms capture how the log price-dividend ratio responds to changes in state variables, with terms reflecting both direct effects of state variable dynamics and quadratic variation terms arising from Itô's lemma. The market prices of risk ($m_{\delta,t}$, $m_{\pi,t}$, and $m_{y,t}$) appear because they determine how the SDF covaries with asset returns.

We first replace the solution for the log-wealth consumption ratio I in the real interest rate $r_{R,t}$ (Eq. 15) and in the market prices of risk $m_{\delta,t}$, $m_{\pi,t}$, and $m_{y,t}$ (Eqs. 20–22), which are then replaced in the above PDE. We then solve for the market log price-dividend ratio Π using the same numerical procedure as above.

E Economic Interpretation of Partial Derivatives

This appendix gives approximations and intuition for key partial derivatives of the log wealth-consumption ratio $I(x_t)$ and log price-dividend ratio $\Pi(x_t)$.

We study the partial derivatives using two approaches. First, we differentiate the PDE for $I(x_t)$ (Eq. B11) and keep only first-order terms to derive analytical approximations that reveal the economic drivers of their signs. Second, these sign patterns are confirmed numerically in our full computational analysis.

Sensitivity to Inflation. We infer that $I_\pi < 0$ because, when $\beta_\pi < 1$ (as in our estimated model), higher inflation lowers the real interest rate, flattening expected consumption growth and worsening long-run consumption prospects. This effect is stronger when perceived transmission effectiveness $\widehat{a}_t$ is low, since inflation deviations are then expected to persist longer.

Sensitivity to Perceived Transmission Strength $\widehat{a}_t$. Differentiating the PDE (B11) with respect to $\widehat{a}_t$ yields a complex expression for $I_{\widehat{a}}$, which we simplify by dropping terms involving second or higher-order derivatives with respect to $\widehat{a}_t$ (e.g., $I_{\widehat{a}\widehat{a}}$, $I_{\pi\widehat{a}}$, $I_{y\widehat{a}}$) and other complex dynamic interactions. This leads to the approximation:

$$I_{\widehat{a}}^{\text{approx}} \approx \frac{1}{e^{-I} + \lambda_a} \lambda_\pi \phi_t (-I_\pi), \tag{E1}$$

which shows that $I_{\widehat{a}}$ shares the sign of the policy stance ϕ_t. When policy is restrictive ($\phi_t > 0$), greater transmission effectiveness raises the wealth-consumption ratio ($I_{\widehat{a}} > 0$), as policy more effectively reduces inflation, and $I_\pi < 0$. When policy is accommodative ($\phi_t < 0$), stronger transmission implies faster inflation, lowering valuations and making $I_{\widehat{a}} < 0$.

Sensitivity to Output Gap. Similarly, differentiating with respect to y_t yields:

$$I_y^{\text{approx}} \approx \frac{\beta_y \left[(\psi - 1) + \lambda_\pi \widehat{a}_t (-I_\pi)\right]}{e^{-I} + \lambda_y} \tag{E2}$$

showing that I_y is generally positive under our parameters. A higher output gap raises valuations by: (1) increasing the real interest rate via the Taylor rule when $\psi > 1$; and (2) reducing expected future inflation, which helps since $I_\pi < 0$. The second effect is stronger when transmission is effective ($\widehat{a}_t$ high) and the gap is persistent (low λ_y).

These approximations clarify the signs and state dependence of the market prices of risk in (21)–(22). The price of inflation risk is negative because $I_\pi < 0$, and becomes more negative when the learning term $-\frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi} \phi_t I_{\widehat{a}}$ is negative—which occurs when $\phi_t I_{\widehat{a}} > 0$, as confirmed by our approximation. The price of output gap risk is positive since $I_y > 0$.

Similar relationships hold for the price-dividend ratio sensitivities Π_π, $\Pi_{\widehat{a}}$, and Π_y, which shape the risk quantities in (25).

F Stock Price Dynamics and Equity Risk Premium

This appendix provides a detailed derivation of the stock price dynamics (Eq. 24) and the equity risk premium in our monetary-learning economy.

The stock price S_t represents a claim to the dividend process D_t, which follows (23) in the main text. In equilibrium, the stock price can be expressed as $S_t = D_t e^{\Pi(x_t)}$, where $\Pi(x_t)$ is the log price-dividend ratio, determined by the state vector $x_t = [\pi_t \; y_t \; \widehat{a}_t \; \nu_{a,t}]^\top$.

Applying Itô's lemma to S_t:

$$\frac{dS_t}{S_t} = \frac{dD_t}{D_t} + d\Pi(x_t) + \frac{1}{2} d\langle \Pi(x_t), \Pi(x_t)\rangle + d\left\langle \frac{D_t}{D_t}, \Pi(x_t) \right\rangle. \tag{F1}$$

The first term in this expression is the dividend process, which is given by (23). For the dynamics of $\Pi(x_t)$, we apply Itô's lemma:

$$\begin{aligned}
d\Pi(x_t) = {} & \mu_\Pi(x_t)dt + \Pi_\pi d\pi_t + \Pi_{\widehat{a}} d\widehat{a}_t + \Pi_y dy_t + \Pi_\nu d\nu_{a,t} \\
& + \frac{1}{2}\left[\Pi_{\pi\pi}(d\pi_t)^2 + \Pi_{\widehat{a}\widehat{a}}(d\widehat{a}_t)^2 + \Pi_{yy}(dy_t)^2 + \Pi_{\nu\nu}(d\nu_{a,t})^2\right] \\
& + \Pi_{\pi\widehat{a}} d\pi_t d\widehat{a}_t + \Pi_{\pi y} d\pi_t dy_t + \Pi_{\pi\nu} d\pi_t d\nu_{a,t} + \Pi_{\widehat{a}y} d\widehat{a}_t dy_t + \Pi_{\widehat{a}\nu} d\widehat{a}_t d\nu_{a,t} + \Pi_{y\nu} dy_t d\nu_{a,t}.
\end{aligned} \tag{F2}$$

Using the dynamics of the state variables from (10), (7), (1), and (8), and collecting terms:

$$d\Pi(x_t) = \mu_\Pi(x_t)dt + \Pi_\pi \sigma_\pi d\widehat{B}_{\pi,t} + \Pi_{\widehat{a}}\left(-\frac{\phi_t \lambda_\pi \nu_{a,t}}{\sigma_\pi}\right) d\widehat{B}_{\pi,t} + \Pi_y \sigma_y dB_{y,t} \tag{F3}$$

$$= \mu_\Pi(x_t)dt + \left(\sigma_\pi \Pi_\pi - \frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi}\phi_t \Pi_{\widehat{a}}\right) d\widehat{B}_{\pi,t} + \sigma_y \Pi_y dB_{y,t}. \tag{F4}$$

Noting that the Brownian motions are independent, the stock price dynamics are:

$$\begin{aligned}
\frac{dS_t}{S_t} = {} & \underbrace{\left[\mu_D(x_t) + \mu_\Pi(x_t) + \frac{1}{2}\left(\sigma_\pi^2 \Pi_\pi^2 + \sigma_y^2 \Pi_y^2 + \sigma_D^2 + \text{other quadratic terms}\right)\right]}_{\text{drift term}} dt \\
& + \underbrace{0 \cdot dB_{\delta,t}}_{\text{consumption shock}} + \underbrace{\left(\sigma_\pi \Pi_\pi - \frac{\lambda_\pi \nu_{a,t}}{\sigma_\pi}\phi_t \Pi_{\widehat{a}}\right) d\widehat{B}_{\pi,t}}_{\text{inflation shock}} \\
& + \underbrace{\sigma_y \Pi_y dB_{y,t}}_{\text{output gap shock}} + \underbrace{\sigma_D dB_{D,t}.}_{\text{dividend shock}}
\end{aligned} \tag{F5}$$

where $\mu_D(x_t) = (1-\alpha)\bar{\mu}_\delta + \alpha\mu_{\delta,t}$. The diffusion coefficients on each Brownian motion are the quantities of risk, $s_{k,t}$ (Eq. 25), that are the essential inputs for the equity risk premium derivation.

By the fundamental asset pricing equation (B6), the expected excess real return on the stock must equal the negative covariance of its return with the stochastic discount factor (SDF):

$$\mathbb{E}_t[dR_{M,t}] - r_{R,t}dt = -\mathrm{Cov}_t\left(\frac{d\xi_t}{\xi_t}, dR_{M,t}\right), \tag{F6}$$

where $dR_{M,t} = \frac{dS_t}{S_t} + \frac{D_t}{S_t}dt$ is the total real return on the stock. The equity risk premium is therefore $\mathrm{RP}_t = -\mathrm{Cov}_t\left(\frac{d\xi_t}{\xi_t}, dR_{M,t}\right)/dt$. Since the dividend yield term, $\frac{D_t}{S_t}dt$, is locally deterministic, the covariance is driven entirely by the capital gain term, $\frac{dS_t}{S_t}$:

$$\mathrm{RP}_t dt = -\mathrm{Cov}_t\left(\frac{d\xi_t}{\xi_t}, \frac{dS_t}{S_t}\right). \tag{F7}$$

The SDF dynamics are given by (19) in the main text:

$$\frac{d\xi_t}{\xi_t} = -r_{R,t}dt - m_t^\top d\widehat{B}_t, \tag{F8}$$

where $m_t = [m_{\delta,t}\ m_{\pi,t}\ m_{y,t}]^\top$ is the vector of market prices of risk. Computing the covariance between the diffusion of the real SDF and the real stock return:

$$\mathrm{Cov}_t\left(\frac{d\xi_t}{\xi_t}, \frac{dS_t}{S_t}\right) = \mathbb{E}_t\left[\left(-m_t^\top d\widehat{B}_t\right)\left(s_{\delta,t}dB_{\delta,t} + s_{\pi,t}d\widehat{B}_{\pi,t} + s_{y,t}dB_{y,t} + s_{D,t}dB_{D,t}\right)\right] \tag{F9}$$

$$= -m_{\delta,t}s_{\delta,t}dt - m_{\pi,t}s_{\pi,t}dt - m_{y,t}s_{y,t}dt - 0\cdot s_{D,t}dt \tag{F10}$$

$$= -\left(m_{\delta,t}s_{\delta,t} + m_{\pi,t}s_{\pi,t} + m_{y,t}s_{y,t}\right)dt. \tag{F11}$$

Therefore, the equity risk premium is $\mathrm{RP}_t = m_{\delta,t}s_{\delta,t} + m_{\pi,t}s_{\pi,t} + m_{y,t}s_{y,t}$. Due to dividend-shock independence, its price of risk $m_{D,t}$ is zero. Also, since $s_{\delta,t} = 0$ (as consumption shocks do not directly affect dividends in our specification), the risk premium simplifies to:

$$\mathrm{RP}_t = m_{\pi,t}s_{\pi,t} + m_{y,t}s_{y,t} \tag{F12}$$

$$= m_{\pi,t}\left(\sigma_\pi\Pi_\pi - \frac{\lambda_\pi\nu_{a,t}}{\sigma_\pi}\phi_t\Pi_{\widehat{a}}\right) + m_{y,t}\sigma_y\Pi_y. \tag{F13}$$

Substituting the expressions for $m_{\pi,t}$ and $m_{y,t}$ from (21)–(22) yields the full expression for the equity risk premium given in (29) in the main text.

Having derived all the necessary components, we can now construct the full expression for the stock price dynamics. The drift of the real stock price process must equal its expected real capital gain. The expected real capital gain is the expected total real return minus the real dividend yield:

$$\mathrm{Drift}\left(\frac{dS_t}{S_t}\right) = \mathbb{E}_t[dR_{M,t}] - \frac{D_t}{S_t} = (r_{R,t} + \mathrm{RP}_t) - \frac{D_t}{S_t}.$$

The diffusion of the stock price process is determined by its exposure to the underlying Brownian motions. These diffusion loadings, $s_{k,t}$, were derived earlier in this appendix and are summarized in (25) in the main text:

$$\mathrm{Diffusion}\left(\frac{dS_t}{S_t}\right) = s_{\delta,t}dB_{\delta,t} + s_{\pi,t}d\widehat{B}_{\pi,t} + s_{y,t}dB_{y,t} + s_{D,t}dB_{D,t}.$$

Combining the derived drift and diffusion components, we arrive at the complete expression for the stock price dynamics as presented in (24) in the main text:

$$\frac{dS_t}{S_t} = \left(r_{R,t} + \mathrm{RP}_t - \frac{D_t}{S_t}\right)dt + s_{\delta,t}dB_{\delta,t} + s_{\pi,t}d\widehat{B}_{\pi,t} + s_{y,t}dB_{y,t} + s_{D,t}dB_{D,t}. \tag{F14}$$

G Maximum Likelihood Estimation

This appendix details our estimation methodology for inferring the model's structural parameters from macroeconomic time series. Our empirical implementation uses monthly U.S. data spanning July 1954 to December 2023. Output data (δ_t) are drawn from real GDP values in the NIPA tables. We obtain the Federal funds rate and output gap from the Federal Reserve Bank of St. Louis (FRED), using their continuously compounded values as empirical counterparts to the nominal interest rate $r_{N,t}$ and output gap y_t, respectively. For inflation π_t, we employ the year-over-year log growth rate of the Consumer Price Index (CPI).

We begin by estimating the consumption growth volatility parameter by maximizing the log-likelihood function:

$$l_\delta(\Theta_\delta; u_{\delta,\Delta}, \ldots, u_{\delta,J\Delta}) = \sum_{j=1}^{J} \log\left(\frac{1}{(2\pi)^{1/2}\sqrt{\sigma_\delta^2\Delta}}\right) - \frac{1}{2}\left(\sigma_\delta^2\Delta\right)^{-1} u_{\delta,j\Delta}^2, \tag{G1}$$

where $\Delta = 1/12$ represents the monthly time increment, $\Theta_\delta \equiv (\sigma_\delta)^\top$ is the parameter vector, J denotes the number of observations, and $\top$ is the transpose operator. The innovations $u_{\delta,t+\Delta}$ are specified as:

$$u_{\delta,t+\Delta} = \log\left(\delta_{t+\Delta}/\delta_t\right) - \left(\text{avg(GDP growth)} - \frac{1}{2}\sigma_\delta^2\right)\Delta. \tag{G2}$$

Here, avg(GDP growth) is the annualized average of observed GDP growth over the sample period.

For the Taylor rule parameters, we construct a separate log-likelihood function:

$$l_r(\Theta_r; u_{r,\Delta}, \ldots, u_{r,J\Delta}) = \sum_{j=1}^{J} \log\left(\frac{1}{(2\pi)^{1/2}\sqrt{\sigma_r^2\Delta}}\right) - \frac{1}{2}\left(\sigma_r^2\Delta\right)^{-1} u_{r,j\Delta}^2, \tag{G3}$$

where $\Theta_r \equiv (\bar{r}_N, \beta_\pi, \beta_y, \sigma_r)^\top$ comprises the neutral rate, policy response coefficients, and residual volatility. The monetary policy innovations are defined as:

$$u_{r,t} = r_{Nt} - \left[\bar{r}_N + \beta_\pi\left(\pi_t - \text{avg(Inflation)}\right) + \beta_y y_t\right]. \tag{G4}$$

The term avg(Inflation) is the annualized average of the inflation rate over the sample period.

Estimating the inflation and transmission coefficient parameters presents a more complex challenge due to the unobservable nature of a_t. We approach this by discretizing the continuous-time dynamics from equations (7) and (4) as follows:

$$\pi_{t+\Delta} = \pi_t e^{-\lambda_\pi\Delta} + \widehat{\pi}_t^{\mathcal{A}}\left(1 - e^{-\lambda_\pi\Delta}\right) + \sqrt{var_\pi}\,\epsilon_{\pi,t+\Delta}, \tag{G5}$$

$$\widehat{\pi}_t^{\mathcal{A}} = \bar{\pi} - \widehat{a}_t(r_{Nt} - \bar{r}_N) \tag{G6}$$

$$\widehat{a}_{t+\Delta} = \widehat{a}_t e^{-\lambda_a\Delta} - \frac{(r_{Nt} - \bar{r}_N)\lambda_\pi\nu_{a,t}}{\sigma_\pi}\sqrt{\frac{1 - e^{-2\lambda_a\Delta}}{2\lambda_a}}\,\epsilon_{\pi,t+\Delta}, \tag{G7}$$

$$\nu_{a,t+\Delta} = \nu_{a,t} + \left[\sigma_a^2 - 2\lambda_a\nu_{a,t} - \left(\frac{(r_{Nt} - \bar{r}_N)\lambda_\pi\nu_{a,t}}{\sigma_\pi}\right)^2\right]\Delta, \tag{G8}$$

where $var_\pi = \frac{\sigma_\pi^2}{2\lambda_\pi}\left(1 - e^{-2\lambda_\pi\Delta}\right)$ and $\epsilon_{\pi,t+\Delta}$ represents a standard normal random variable. This discretization preserves the essential features of the continuous-time filter while enabling estimation through conventional likelihood methods.

The inflation component parameters are recovered by maximizing:

$$l_\pi(\Theta_\pi; u_{\pi,\Delta}, \ldots, u_{\pi,J\Delta}) = \sum_{j=1}^{J} \log\left(\frac{1}{(2\pi)^{1/2}\sqrt{var_\pi}}\right) - \frac{1}{2}\left(var_\pi\right)^{-1} u_{\pi,j\Delta}^2, \tag{G9}$$

where $\Theta_\pi \equiv (\sigma_\pi, \bar{\pi}, \lambda_\pi, \sigma_a, \lambda_a)^\top$ collects the inflation volatility, target, mean-reversion speed, and transmission coefficient parameters. The inflation innovations are specified as:

$$u_{\pi,t+\Delta} = \pi_{t+\Delta} - \left[\pi_t e^{-\lambda_\pi \Delta} + \widehat{\pi}_t^{\mathcal{A}} \left(1 - e^{-\lambda_\pi \Delta}\right)\right]. \tag{G10}$$

The dynamics of the filtered transmission coefficient $\widehat{a}_t$ and its uncertainty $\nu_{a,t}$ follow the recursive updating rules given in equations (G7) and (G8), respectively. This approach allows us to jointly estimate the parameters governing the inflation process and the learning dynamics without requiring direct observation of the transmission coefficient itself.