Loud Markets, Quiet Markets: Asset Price Informativeness and Monetary Policy

Wei Cui and Xiao Yin *

Preliminary. Please do not cite or circulate

Abstract

We build a model of monetary policy under disagreement in which the Fed imperfectly infers the market's belief from asset prices before setting the interest rate. Asset prices are nonlinear in the market's belief, so they become more informative when the market holds extreme views. The unconditional response of the interest rate to market optimism is near zero: in normal times, the asset price signal is too noisy for the Fed to act on. However, the response becomes large and significant when either (i) the difference of opinions between the Fed and the market is large, or (ii) the Fed faces high uncertainty about fundamentals. We also show that imperfect information can be welfare-improving when market sentiment contains noise: the nonlinear asset price acts as an endogenous filter that dampens the Fed's response to sentiment fluctuations while preserving its response to genuine disagreement.

JEL codes:

Keywords:

*Cui: University College London, w.cui@ucl.ac.uk; Yin: University College London, xiao.yin@ucl.ac.uk.

I Introduction

The Fed and financial markets frequently disagree about the outlook for economic activity and the appropriate path of interest rates. Empirical evidence shows that the Fed's predicted rate path and the market's forward rates often diverge for extended periods (Caballero and Simsek, 2022). Survey evidence further confirms that these disagreements reflect genuinely different views about future aggregate demand, not merely differences in risk premia (Sastry, 2024). In this paper, we study monetary policy under an imperfect information environment in which the Fed must infer the market's belief from noisy financial signals.

A natural benchmark assumes *perfect information*: the Fed and the market observe each other's beliefs and can agree to disagree. In this case, the Fed sets the policy rate as a weighted average of its own view and the market's view, with the weight on the market governed by the persistence of disagreements. The Fed partially accommodates the market because expectations of future policy mistakes feed through forward rates and affect current output. In practice, the Fed does not observe the market's belief directly. Instead, it must *infer* it from public signals, such as equity prices, bond yields, forward rates, survey data, and other financial indicators, that are noisy and nonlinear functions of underlying beliefs. This inference problem, combined with the nonlinearity of asset prices, generates predictions absent from the perfect information benchmark.

We use equity price as the leading example because the model's asset pricing equation provides a natural microfoundation: the price of the market portfolio is proportional to output in levels, so the expected asset price is an *exponential* function of the market's perceived output gap. However, the mechanism is not specific to equity prices. Any public signal whose sensitivity to the market's belief increases with the level of that belief, a property shared by many financial indicators in levels, would generate the same qualitative predictions. For instance, corporate bond spreads, housing prices, and consumer confidence indices all exhibit greater responsiveness to sentiment shifts in booms than in busts, and could serve as alternative channels for the Fed's inference.

The paper delivers four sets of findings, organized around the comparison between the perfect and imperfect information environments.

First, we characterize the Fed's policy response to market optimism. Under perfect information, the response is constant: the Fed always adjusts the rate by a fraction of disagreement. Under imperfect information, the local response coefficient is utilizes the Fed's effective Kalman gain on the asset price signal, increasing in the Fed's prior estimate of market optimism. When the market is expected to be mildly optimistic or pessimistic, the asset price is on a flat part of the exponential and provides little signal, so Kalman gain is small and the Fed barely responds. When the market is expected to hold extreme views, the exponential is steep and the Fed responds nearly as strongly as under perfect information.

This means that the *unconditional* response of the interest rate to market optimism, the coefficient one would obtain from a simple regression of the rate on the market's perceived output gap, is near zero. In normal times, the asset price is too noisy for the Fed to act on, and the Fed effectively sets rates based on its own view alone. However, the response "turns on" when the difference of opinions is large enough for the asset price signal to cut through the noise.

Second, we show that the Fed's response intensifies when it is more uncertain about its own view. Higher prior uncertainty raises the Kalman gain, since the Fed puts more weight on the market price signal when its own prior is less informative. This comparative static is absent under perfect information, where the response is independent of uncertainty.

Third, we show that imperfect information generates an additional source of gradualism in monetary policy. Under perfect information, the only source of gradualism is the Fed's expectation that the market will learn from data and converge toward the Fed's view. Under imperfect information, there is a second layer: the Fed itself is learning the market's view from sequential asset price observations. We call this "double gradualism," showing that the Fed's expected rate path is flatter under imperfect information, with the gap closing gradually as the Fed accumulates information.

Fourth, we show that imperfect information *can* be welfare-improving when market sentiment contains noise. Under perfect information, the Fed is forced to respond to every fluctuation in the market's belief, including sentiment-driven noise that does not reflect fundamentals. Under imperfect information, the nonlinear asset price acts as an

endogenous filter: small sentiment fluctuations (which are more likely to be noise) are invisible to the Fed, while large belief shifts (which are more likely to reflect genuine disagreement) cut through the noise and elicit a policy response. This filtering can reduce the welfare cost of sentiment shocks relative to the perfect information benchmark.

Our findings are consistent with a specific empirical pattern (to be added later): a linear regression of the interest rate on the market's perceived output gap yields an insignificant coefficient, but the coefficient becomes significant when the sample is split by the level of disagreement or by the degree of Fed uncertainty. The model provides a structural explanation for this pattern: the unconditional insignificance reflects the noise in asset prices at normal times, while the conditional significance reflects the state-dependent informativeness of the exponential price function.

Related literature. Our paper contributes to several strands of the literature.

Disagreement and monetary policy. Our model builds directly on Caballero and Simsek (2022), who show that opinionated disagreements between the Fed and the market generate gradualism, Fed belief surprises, and tantrum shocks. We share their belief structure (heterogeneous interpretations, agree-to-disagree, Kalman filtering) and their policy framework. Our contribution is to relax the assumption that the Fed perfectly observes the market's belief and to show that the resulting inference problem—combined with the nonlinearity of asset prices—generates novel predictions about the state dependence of monetary policy. Stein and Sunderam (2018) study a related setting in which the Fed and the bond market have different models, generating gradualism through a different channel. Sastry (2019) provides empirical evidence that disagreements about the Fed's reaction to news drive monetary policy shocks, consistent with the heterogeneous-interpretation structure we adopt.

The Fed and asset prices. A long-standing debate concerns whether and how the Fed should respond to asset prices (Bernanke and Gertler 1999, 2001; Cecchetti et al. 2000; Gilchrist and Leahy 2002). This literature focuses on whether asset prices contain information about future inflation or output that should enter the policy rule. Our question is different: we ask whether asset prices contain information about the *market's*

belief, which the Fed needs to know in order to anticipate the market's reaction to its policy decisions. Bond, Edmans, and Goldstein (2012) survey the broader literature on "feedback effects"—the idea that prices aggregate information that decision-makers can use. Our model provides a specific instance of this feedback in a monetary policy context. Cieslak and Vissing-Jorgensen (2021) document that the Fed responds to the stock market between meetings, consistent with the idea that equity prices serve as a signal of the market's view.

The Fed information effect. Romer and Romer (2000) show that the Fed's policy actions convey information about the state of the economy. Nakamura and Steinsson (2018) formalize this "Fed information effect": when the Fed raises rates, the market partly interprets the hike as a signal of Fed optimism about fundamentals. Our model features a *reverse* information flow: the market's asset prices convey information about the market's belief *to the Fed*. The two channels are complementary—the Fed learns from the market before its announcement, and the market learns from the Fed after the announcement. Bauer and Swanson (2023) argue that the apparent Fed information effect is better explained by the Fed responding to public data more than the market expected, which is closely related to the heterogeneous-interpretation structure in our model and in CS.

Heterogeneous beliefs and financial markets. Our belief structure draws on the literature on agree-to-disagree models with heterogeneous priors or heterogeneous interpretation of signals (Harrison and Kreps 1978; Scheinkman and Xiong 2003; Simsek 2013, 2021). Sethi and Yildiz (2016) provide the formal foundation for heterogeneous interpretations of public signals that CS and we adopt. Kandel and Pearson (1995) provide early empirical evidence that trading volume around earnings announcements is driven by heterogeneous interpretation of the same information, consistent with the signal structure in our model. Andrade et al. (2016) and Bordalo et al. (2020) document systematic patterns in forecast disagreement among professional forecasters.

Nonlinear policy rules. The empirical literature on nonlinear Taylor rules (Dolado, MarÃa-Dolores, and Naveira 2005; Surico 2007; Cukierman and Muscatelli 2008) typically finds that the Fed responds more strongly to positive output gaps than to negative ones.

The standard explanation involves asymmetric preferences or a convex Phillips curve (Debelle and Laxton 1997; Filardo 1998). We provide an alternative, information-based explanation: the Fed's inference from asset prices is better in booms than in busts because asset prices are more sensitive to the market's belief when the market is optimistic. The asymmetry in our model arises endogenously from the asset pricing equation, not from asymmetric preferences.

Uncertainty and monetary policy. Brainard (1967) shows that parameter uncertainty induces the central bank to respond less aggressively to shocks—an effect that predicts a *negative* relationship between uncertainty and the policy response. Our model generates the opposite prediction: higher uncertainty about the market's belief makes the Fed respond *more* to market signals, because an uncertain Fed places greater weight on external information. The distinction is that Brainard uncertainty concerns the transmission mechanism (how does the rate affect output?), while our uncertainty concerns the state (what does the market believe?). Orphanides and Williams (2005) study monetary policy under imperfect knowledge of the natural rate, which shares the feature that the policymaker must infer a key state variable from noisy data.

Learning and monetary policy. Evans and Honkapohja (2001) and Eusepi and Preston (2011) study environments where agents learn about the economy and policy must account for this learning. In their models, the agents' beliefs are boundedly rational but the policymaker understands the true model. In our setting, neither the Fed nor the market has the "true" model—both are Bayesian given their own interpretations. The policy problem is not to correct the market's learning but to infer the market's belief and respond appropriately.

Central bank communication. Blinder et al. (2008) survey the literature on central bank transparency, documenting that greater transparency improves the predictability of policy. Woodford (2005) argues that communication is valuable because it helps align private-sector expectations with the central bank's intentions. In our model, communication operates in the reverse direction: the Fed benefits from transparency not in order to influence the market, but in order to *read* the market. When the market's belief is partially revealed through asset prices, the Fed can set more appropriate policy—a

form of "listening" rather than "speaking."

Below, Section 2 describes the model environment. Section 3 derives the equilibrium under discretion, including the Fed's inference problem, the state-dependent policy response, and the empirical predictions. Section 4 analyzes the welfare implications of imperfect information, showing that it can act as a beneficial filter against sentiment noise. Section 5 concludes.

II The Basic Environment

The model setup has a standard New Keynesian framework with different opinions by a representative household and the Fed. Policy actions affect the economy with a lag. The transmission lags imply that belief disagreements about future aggregate demand/productivity matter since the market expects the Fed to make "mistakes." Importantly, we allow the Fed to have imperfect information about the market's belief, and we allow asset prices to partially reveal this information. Given that asset prices are non-linear in the market's belief, their informativeness becomes state-dependent, which is the key mechanism of the paper. The perfect information benchmark—in which the Fed directly observes the market's belief—is a special case obtained by setting the asset price noise to zero.

Time is discrete and continues forever. The Fed sets the policy rate in each period before the realization of an aggregate shock. We divide each period into three stages. In the first stage, both the Fed and the market observe a public signal, but they have different interpretations of the signal, and the interpretations are private at this stage. In the second stage, the market's pre-shock belief is partially revealed through an asset price; then the Fed sets the nominal risk-free interest rate. In the third stage, the aggregate shock (another public signal without heterogeneous interpretation) is realized, the market chooses optimal allocations, and the equilibrium output is determined.

We will use $j \in \{F, M\}$ to denote the agents as the Fed and the market. We will use $\mathbb{E}_t^j$ as agent j's expectation with all information at time t, and $\hat{\mathbb{E}}_t^j$ as his expectation *before* the aggregate shock at time t. So without the hat notation, we mean that all public information is received.

II.A Preferences and technology

The market features a representative household with utility

$$\mathbb{E}_t^M \sum_{h=0}^{\infty} \beta^h \left[\log C_t - \frac{N_t^{1+\eta}}{1+\eta} \right] \tag{1}$$

where C_t is consumption, N_t is labor supply, and η is the inverse labor supply elasticity. After observing the aggregate shock at time t, the household maximizes the above utility subject to

$$P_t C_t + \frac{B_{t+1}}{R_t} = B_t + W_t N_t + \int_0^1 D_t(i) di \tag{2}$$

where P_t is the nominal price level, B_t is nominal bond, R_t is the safe gross nominal interest rate, W_t is nominal wage, and $D_t(i)$ is the profit/dividend paid by intermediate-good firm i. The first-order conditions (for labor and for bonds) are

$$\frac{W_t}{P_t} = C_t N_t^{\eta} \tag{3}$$

$$\frac{1}{C_t} = \beta R_t \mathbb{E}_t^M \left[\frac{P_t}{P_{t+1}} \frac{1}{C_{t+1}} \right] \tag{4}$$

The production side has a competitive final goods sector that produces according to

$$Y_t = \left(\int_0^1 Y_t(i)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{5}$$

where $\sigma > 1$ measures the substitution among goods and $i \in [0,1]$ is an index for monopolistically competitive intermediate goods firms. Each firm i produces according to

$$Y_t(i) = A_t N_t(i)^{1-\alpha} \tag{6}$$

where A_t is the level of aggregate productivity and $N_t(i)$ is the level of hours hired by each firm i. If (nominal) prices are flexible, the equilibrium output will reach its potential level Y_t^*, so the output is $Y_t^* = A_t \left(N^*\right)^{1-\alpha}$ for some potential level of hours N^*.

II.B Price rigidities

In each period, a randomly selected fraction, $1-\theta$, of firms reset their nominal prices. The firms that do not adjust their price in period t set their labor input to meet the demand for their goods (since firms operate with a markup).

As a convention, we log-linearize the equilibrium conditions around the allocation with real potential outcomes and zero inflation. Throughout, we use the notation of

$x \equiv \log X$ as the natural log of variable X and $\tilde{x} \equiv \log\left(X_t / X_t^*\right)$ as the log deviation of X from the potential outcome.

The New-Keynesian Phillips curve has the usual form of

$$\pi_t = \kappa \tilde{y}_t + \beta \mathbb{E}_t^M \left[\pi_{t+1}\right], \text{ where } \kappa \equiv \frac{1-\theta}{\theta}\left(1-\theta\beta\right)\frac{1+\eta}{1-\alpha+\alpha\sigma} \tag{7}$$

where $\pi_t = \log(P_t/P_{t-1})$, $\tilde{y}_t$ is the output gap relative to the potential output y_t^*, and κ is a price flexibility parameter.

II.C Shocks and signals

We discuss the demand-stage shock first. We capture aggregate shocks with news about growth as in CS. Let production technology/demand $a_t = \log A_t$ be determined one period in advance. The process for potential aggregate productivity growth $\Delta a_{t+1} = a_{t+1} - a_t \equiv g_t$ is given by

$$\Delta a_{t+1} = g_t = z_t + \nu_t, \text{ where } \nu_t \sim N(0, \sigma_\nu^2) \tag{8}$$

where z_t is the fundamental part of the growth rate and ν_t is a transitory noise. Therefore, Δa_{t+1} is realized in the demand stage of period t because of the realization of z_t and ν_t.

We log-linearize the consumption-savings choice of the household (the market) to obtain the dynamic IS curve with the z_t component:

$$\tilde{y}_t = -\left(\iota_t - \mathbb{E}_t^M\left[\pi_{t+1}\right] - \rho\right) + g_t + \mathbb{E}_t^M\left[\tilde{y}_{t+1}\right] \tag{9}$$

where $\rho \equiv -\log\beta$, and naturally $\iota_t - \mathbb{E}_t^M\left[\pi_{t+1}\right]$ is the (market-expected) real interest rate. For a given real rate and expected future output gap, the potential growth rate, g_t, raises the current output gap one-for-one. Since g_t directly enters the IS curve, we follow CS in labeling it as the *demand shock*. Notice that at the signal stage of period t, there is no uncertainty about the potential output of the economy at t—the demand shock reveals the potential output at time $t+1$. The fundamental part z_t follows a unit root process

$$z_t = z_{t-1} + \varepsilon_t, \text{ where } \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \tag{10}$$

The unit root assumption is a tractable way to capture the idea that the "underlying state of the economy"—the trend growth rate—can shift permanently. This is what makes disagreements persistent: agents who form different views about the level of z_t will continue to disagree until new data persuades one of them to revise.

At the signal stage, before seeing g_t, agents receive a public signal s_t about z_t: $s_t = z_t + e_t$, which is a noisy observation of z_t with $e_t \sim N(0, \sigma_e^2)$. One can think of this public signal as arising from earnings announcements of major firms (e.g., S&P 500 constituents), macroeconomic data releases, or other high-frequency indicators that both the Fed and the market observe simultaneously.

For understanding how disagreements arise from a shared signal, we use *heterogeneous uncertain beliefs* (HUB). That is, agents see the same signal s_t, but they interpret it differently. The idea is that a single data release—say, a jobs report—can be read as strong or weak depending on one's prior model of the economy, one's weighting of seasonal factors, or one's assessment of structural change. Agents understand that their counterparts interpret the signal s_t in their own way; they simply disagree about whose interpretation is correct.

Precisely, agent j interprets the signal of the permanent component with a different bias c_t^j:

$$s_t =^j z_t - c_t^j + e_t \tag{11}$$

Effectively, agent j treats the signal as $s_t^j = s_t + c_t^j$. The interpretation biases are drawn from a joint normal distribution that is i.i.d. across periods:

$$c_t^F, c_t^M \sim N(0, \sigma_\delta^2) \text{ and } corr(c_t^F, c_t^M) = \rho_c \tag{12}$$

For the main discussion, we set $\rho_c = 0$, so the Fed's and the market's interpretations are independent. This is the most conservative assumption: any positive correlation would reduce disagreement and make the Fed's inference easier.

Importantly, at the signal stage, the Fed and the market cannot see each other's interpretation (though they know the distributional structure). The Fed needs to form a belief about the market's interpretation in order to set the interest rate. We allow the asset price Q_t^{pre}, determined by the market's belief, to be observed by the Fed before the rate decision. More broadly, any public signal that reflects the market's belief—whether an equity index, a forward interest rate, or a survey of market participants—could serve this role. We focus on the asset price because the model provides a natural microfoundation for its functional form. After the demand shock g_t is realized, the market makes optimal labor,

consumption, and savings choices, and the Fed can fully learn the market's interpretation from the realized equilibrium.

II.D Monetary policy

The Fed sets the nominal interest rate $\iota_t = R_t - 1$ at the asset price stage, before the realization of g_t, while private agents choose spending and working after the shock in the demand stage. The Fed minimizes a conventional objective function

$$\hat{\mathbb{E}}_t^F \left[V_t^F \right] \equiv \hat{\mathbb{E}}_t^F \left[\sum_{h=0}^{\infty} \beta^h \left(\gamma \tilde{y}_{t+h}^2 + \pi_{t+h}^2 \right) \right], \tag{13}$$

where γ denotes the weight on the output gap relative to inflation. We assume the Fed's policy does not have commitment, so the Fed's problem is

$$\min_{\iota_t} \hat{\mathbb{E}}_t^F \left[\gamma \tilde{y}_t^2 + \pi_t^2 \right] + \hat{\mathbb{E}}_t^F \left[V_{t+1}^F \right] \tag{14}$$

subject to (7) and (9). The expected continuation value will be treated as exogenous to its decision ι_t—this is because the model has no endogenous state variables that link the current rate to future outcomes. Therefore, effectively the Fed solves a sequence of static quadratic problems.

The absence of commitment is important for our results. Under discretion, the Fed cannot promise to follow a particular rule in the future, so the market's expectation of future policy is determined by the market's belief about future fundamentals, not by the Fed's current action. This means the Fed's interest rate decision affects the current output gap but not the market's expected future output gap. As we show below, this implies certainty equivalence: the Fed's optimal rate is linear in its estimate of the market's belief, with a constant coefficient. The state dependence in our model arises entirely from the quality of the Fed's estimate, not from the Fed's decision rule.

II.E Beliefs and disagreement

In summary, there are three stages within each period t:

1. **Signal stage.** The permanent component of demand follows a random walk:

$$z_t = z_{t-1} + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2). \tag{15}$$

   but agents do not observe z_t directly. Both agents observe a public signal and draw private interpretations. Each agent updates to $\hat{\mu}_t^i$. Each agent $i \in \{F, M\}$ observes

a private interpreted signal $s_t^i = z_t + e_t + c_t^i$, where $e_t \sim N(0, \sigma_e^2)$ is common noise and $c_t^i \sim N(0, \sigma_\delta^2)$ are independent private interpretations.

2. **Asset price stage.** An asset price Q_t^{pre} is determined by the market's pre-shock belief. The Fed observes Q_t^{pre} and sets ι_t.
3. **Demand stage.** After the rate is set, both agents observe:

$$g_t = z_t + \nu_t, \qquad \nu_t \sim N(0, \sigma_\nu^2). \tag{16}$$

The market acts, and $\hat{\mu}_t^M$ is fully revealed.

The updating works as follows: (1). For agent $i \in \{F, M\}$, given some prior $z_t \sim N(\mu_{t-1}^i, \Sigma_{t-1}^i)$ and after observing new information s_t (with different interpretation about the noise), agent i updates his belief about the mean and variance of z_t, with a distribution $N(\hat{\mu}_t^i, \hat{\Sigma}_t^i)$; (2). After seeing the Δa_{t+1}, agent i further updates his belief to $N(\mu_t^i, \Sigma_t^i)$, the prior for the beginning of next period. Standard Kalman filter results imply the learning outcome. After observing s_t, the recursive relationships can be expressed as

$$\hat{\mu}_t^i = \underbrace{\frac{\Sigma_{t-1}^i + \sigma_\varepsilon^2}{\Sigma_{t-1}^i + \sigma_\varepsilon^2 + \sigma_i^2}}_{\omega_s^i} \left(s_t + c_t^i\right) + (1 - \omega_s^i)\mu_{t-1}^i \tag{17}$$

$$\hat{\Sigma}_t^i = \left[\frac{1}{\Sigma_{t-1}^i + \sigma_\varepsilon^2} + \frac{1}{\sigma_i^2}\right]^{-1} \tag{18}$$

where ω_s^i is the Kalman gain (i specific). Notice that $\hat{\mu}_t^i$ is simply the conditional expectation $\hat{\mathbb{E}}_t^i[z_t]$. After observing the demand shock, we have

$$\mu_t^i = \underbrace{\frac{\hat{\Sigma}_t^i}{\hat{\Sigma}_t^i + \sigma_\nu^2}}_{\omega_g^i} g_t + (1 - \omega_g^i)\hat{\mu}_t^i = \omega_g^i g_t + (1 - \omega_g^i)\omega_s^i \left(s_t + c_t^i\right) + \left(1 - \omega_g^i\right)\left(1 - \omega_s^i\right)\mu_{t-1}^i$$

$$\Sigma_t^i = \left[\frac{1}{\Sigma_{t-1}^i + \sigma_\varepsilon^2} + \frac{1}{\sigma_i^2} + \frac{1}{\sigma_\nu^2}\right]^{-1} \tag{19}$$

where ω_g^i is the Kalman gain for agent i after observing g_t. Importantly, μ_t^i is conditional expectation $\mathbb{E}_t^i[z_t]$ (different from $\hat{\mathbb{E}}_t^i[z_t]$). Below we characterize the recursion for $\hat{\mu}_t^i$ as it is more important for interest rate policy (before the shock).

In the learning steady state, the Kalman gains are:

$$\omega_s = \frac{\Sigma + \sigma_\varepsilon^2}{\Sigma + \sigma_\varepsilon^2 + \sigma^2}, \qquad \omega_g = \frac{\hat{\Sigma}}{\hat{\Sigma} + \sigma_\nu^2}, \qquad \phi = (1 - \omega_s)(1 - \omega_g), \tag{20}$$

where $\sigma^2 = \sigma_e^2 + \sigma_\delta^2$. The parameter $\phi \in (0, 1)$ governs the persistence of beliefs: it is

the fraction of the prior that survives after incorporating the two signals (the interpreted signal and the demand shock). High ϕ means agents are stubborn—signals are noisy relative to the prior, so beliefs move slowly and disagreements persist. Low ϕ means signals are precise and agents converge quickly.

The pre-shock belief satisfies:

$$\hat{\mu}_t^i = \phi\, \hat{\mu}_{t-1}^i + \omega_s\, s_t^i + (1 - \phi - \omega_s)\, g_{t-1}. \tag{21}$$

This recursion has a simple interpretation: the agent's current belief is a weighted average of three objects—the previous belief (weight ϕ), the new interpreted signal (weight ω_s), and the most recent demand shock (weight $1 - \phi - \omega_s$). The weights sum to one.

The disagreement $\Delta_t = \hat{\mu}_t^M - \hat{\mu}_t^F$ evolves as:

$$\Delta_t = \phi\Delta_{t-1} + \omega_s(c_t^M - c_t^F). \tag{22}$$

Disagreements are persistent (they decay at rate ϕ per period) but are regenerated each period by new interpretation differences $c_t^M - c_t^F$. On average, each agent expects disagreements to shrink: the other agent will learn from data and come toward one's own view. But the other agent thinks the same thing in reverse—this is the "agree to disagree" structure.

At the demand stage, $\hat{\mu}_t^M$ is fully revealed, so Δ_{t-1} is known to the Fed at the start of period t. This is a crucial feature of the information structure: the Fed always enters a new period knowing *last period's* disagreement perfectly. The uncertainty is only about the *current* innovation to disagreement, driven by the fresh interpretation difference $c_t^M - c_t^F$.

II.F The market portfolio

Following Caballero and Simsek, the log price of the market portfolio satisfies $q_t = q^* + a_t + \tilde{y}_t$, where $a_t = \log A_t$ is log productivity. In levels:

$$Q_t = Q^* \cdot A_t \cdot \exp(\tilde{y}_t). \tag{23}$$

This equation is the standard result under log utility: the price of a claim on future output is proportional to current output. The exponential arises because the log price is linear in the output gap, so the level price is exponential.

At the asset price stage (before g_t is realized), the market's expectation of Q_t is:

$$\hat{E}_t^M[Q_t] = Q^* A_t \cdot \hat{E}_t^M[\exp(\tilde{y}_t)]. \tag{24}$$

Under the market's belief, the output gap has a conditional distribution (from the IS curve) with mean $\hat{E}_t^M[\tilde{y}_t]$ and variance $V = \hat{\Sigma} + \sigma_\nu^2$. Using the log-normal formula:

$$\hat{E}_t^M[Q_t] = Q^* A_t \cdot \exp\Big(\hat{E}_t^M[\tilde{y}_t] + \tfrac{1}{2}V\Big). \tag{25}$$

The key observation is that the expected asset price is an exponential function of the market's expected output gap. This nonlinearity is not an assumption we impose—it is an unavoidable consequence of working with prices in levels rather than in logs. Log-linearized models replace $\exp(\tilde{y}_t)$ with $1 + \tilde{y}_t$ and lose the nonlinearity entirely. But the Fed observes the actual price, not its log-linearized approximation.

While we use the equity market portfolio as our leading example, the same nonlinearity arises in any asset whose price is convex in the market's expected output gap. For instance, the price of a long-term nominal bond is convex in expected future rates, which in turn depend on the market's belief about demand. Corporate bond spreads, real estate values, and option-implied volatilities all share the property that their sensitivity to underlying beliefs increases with the level of those beliefs. The key requirement is not that the signal comes from the equity market specifically, but that it comes from an object whose level—not log—the Fed can observe, and whose level is a convex function of the market's belief.

III Equilibrium under Policy Discretion

For simplicity, we analyze an economy with fully rigid prices (so that inflation does not move) and with discretionary monetary policy. These restrictions are not essential for the qualitative results, but they simplify the analysis considerably: with fully rigid prices, the NKPC drops out and the Fed focuses exclusively on stabilizing the output gap.

III.A The observed asset price and the Fed's inference

With discretionary policy, the standard result gives $\hat{E}_t^F[\tilde{y}_t] = 0$ (see also CS). The Fed targets a zero output gap in expectation under its own belief. Since the market may hold

a different belief, the market's expected output gap is generally nonzero:

$$\hat{E}_t^M[\tilde{y}_t] = \Delta_t. \tag{26}$$

When $\Delta_t > 0$, the market is more optimistic than the Fed about aggregate demand and expects the output gap to be positive—it thinks the Fed is setting rates "too high" and failing to accommodate the strength of the economy. When $\Delta_t < 0$, the market thinks the Fed is "too easy" and expects a negative output gap. Substituting into the asset price:

$$\hat{E}_t^M[Q_t] = Q^* A_t \, e^{V/2} \cdot \exp(\Delta_t). \tag{27}$$

This equation makes the economic content of the nonlinearity transparent. The expected asset price is exponential in disagreement Δ_t because the market portfolio is a claim on output, and output is exponential in the output gap. When the market is very optimistic (Δ_t large and positive), it expects high output and values the market portfolio highly; the exponential means that an additional unit of optimism raises the price by more when the market is already optimistic. Conversely, when the market is pessimistic (Δ_t negative), prices are depressed and an additional unit of pessimism has a smaller effect—the price is already near its lower bound.

The Fed observes a noisy version of the expected asset price:

$$Q_t^{\text{pre}} = \hat{E}_t^M[Q_t] + \zeta_t = Q^* A_t \, e^{V/2} \cdot \exp(\Delta_t) + \zeta_t, \tag{28}$$

where $\zeta_t \sim N(0, \sigma_\zeta^2)$ captures noise from liquidity shocks, risk premium fluctuations, or non-fundamental trading. The noise ζ_t is the reason the Fed cannot perfectly infer Δ_t from the price—in a frictionless world with $\sigma_\zeta^2 = 0$, the asset price would perfectly reveal the market's belief and our model would reduce to the CS benchmark.

The Fed knows all the constants and A_t (productivity is public), so it can compute:

$$R_t \equiv \frac{Q_t^{\text{pre}}}{Q^* A_t \, e^{V/2}} = \exp(\Delta_t) + \tilde{\zeta}_t, \tag{29}$$

where $\tilde{\zeta}_t = \zeta_t / (Q^* A_t e^{V/2})$ is rescaled noise with variance $\tilde{\sigma}_\zeta^2 = \sigma_\zeta^2 / (Q^* A_t e^{V/2})^2$. The ratio R_t strips away all the components the Fed already knows, leaving only the signal about Δ_t plus noise.

The nonlinearity of $R_t = \exp(\Delta_t) + \tilde{\zeta}_t$ in Δ_t is the source of everything that follows. To see why the nonlinearity matters, consider two scenarios. In the first, the market

is mildly optimistic ($\Delta_t = 0.01$): $\exp(0.01) \approx 1.01$, and a small change in Δ_t shifts R_t by approximately 1.01—a gentle slope. In the second, the market is very optimistic ($\Delta_t = 0.10$): $\exp(0.10) \approx 1.105$, and a small change shifts R_t by approximately 1.105—a steeper slope. The noise $\tilde{\zeta}_t$ is the same in both cases, but the signal is stronger in the second. The Fed can extract Δ_t more precisely when the market is optimistic, because the "signal" (the derivative of $\exp(\Delta_t)$) is larger relative to the "noise" ($\tilde{\sigma}_\zeta^2$). Economically, this reflects a basic property of financial markets: bull markets are "loud"—prices move a lot and trading volumes are high, making sentiment easy to read—while bear markets can be "quiet," with prices compressed near their lower bound and less variation to extract information from.

The Fed's prior. The Fed knows Δ_{t-1} and the dynamics (22). The innovation $\omega_s(c_t^M - c_t^F)$ has mean zero and variance $\omega_s^2\sigma_\xi^2$ where $\sigma_\xi^2 = 2\sigma_\delta^2$. So the Fed's prior is:

$$\Delta_t \sim N(\phi\,\Delta_{t-1},\ P_0)\,, \qquad P_0 = \omega_s^2\sigma_\xi^2. \tag{30}$$

The prior mean $\bar{\Delta}_t \equiv \phi\Delta_{t-1}$ reflects the Fed's expectation that disagreements shrink toward zero at rate ϕ. The prior variance P_0 reflects the Fed's uncertainty about the current period's interpretation difference—this is the irreducible "opinion risk" that the Fed faces each period.

Linearized signal extraction. The observed ratio $R_t = \exp(\Delta_t) + \tilde{\zeta}_t$ is a nonlinear function of Δ_t. To extract information, the Fed linearizes around its prior mean $\bar{\Delta}_t \equiv \phi\Delta_{t-1}$:

$$R_t \approx \exp(\bar{\Delta}_t) + \exp(\bar{\Delta}_t) \cdot (\Delta_t - \bar{\Delta}_t) + \tilde{\zeta}_t. \tag{31}$$

Subtracting the known constant $\exp(\bar{\Delta}_t)$:

$$R_t - \exp(\bar{\Delta}_t) \approx \exp(\bar{\Delta}_t) \cdot (\Delta_t - \bar{\Delta}_t) + \tilde{\zeta}_t. \tag{32}$$

This is a linear signal of $\Delta_t - \bar{\Delta}_t$ with slope $\exp(\bar{\Delta}_t)$ and noise $\tilde{\zeta}_t$. The effective signal-to-noise ratio is:

$$\mathrm{SNR}(\bar{\Delta}_t) = \frac{\exp(2\bar{\Delta}_t)}{\tilde{\sigma}_\zeta^2}. \tag{33}$$

This ratio is **increasing** in $\bar{\Delta}_t$: when the Fed expects the market to be optimistic, the asset price is more sensitive to Δ_t (the exponential is steeper), so the price is a better signal.

Bayesian update. Using the linearized signal with prior $N(\bar{\Delta}_t, P_0)$ and signal noise $\tilde{\sigma}_\zeta^2/\exp(2\bar{\Delta}_t)$, we obtain

$$P(\bar{\Delta}_t) = \left(\frac{1}{P_0} + \frac{\exp(2\bar{\Delta}_t)}{\tilde{\sigma}_\zeta^2}\right)^{-1} = \frac{P_0\,\tilde{\sigma}_\zeta^2}{P_0\exp(2\bar{\Delta}_t) + \tilde{\sigma}_\zeta^2}, \tag{34}$$

$$K(\bar{\Delta}_t) = \frac{P_0\exp(2\bar{\Delta}_t)}{P_0\exp(2\bar{\Delta}_t) + \tilde{\sigma}_\zeta^2}, \tag{35}$$

so that the posterior mean is:

$$m_t = \bar{\Delta}_t + K \cdot \exp(-\bar{\Delta}_t)(R_t - \exp(\bar{\Delta}_t)). \tag{36}$$

The posterior mean m_t is the Fed's best estimate of Δ_t given all available information. It is a weighted average of the prior mean $\bar{\Delta}_t$ (from persistent memory) and the signal extracted from the asset price (from real-time market data), with the Kalman gain K governing the weight on the signal.

Proposition 1 (State-dependent informativeness)**.** *The effective Kalman gain K and posterior variance P satisfy:*

1. *$K(\bar{\Delta}_t)$ is increasing in $|\bar{\Delta}_t|$ for $\bar{\Delta}_t > 0$.*
2. *$P(\bar{\Delta}_t)$ is decreasing in $|\bar{\Delta}_t|$ for $\bar{\Delta}_t > 0$.*
3. *As $|\bar{\Delta}_t| \to \infty$: $K \to 1$, $P \to 0$ (the price becomes perfectly informative).*
4. *As $\bar{\Delta}_t \to 0$: $K \to K_0 \equiv P_0/(P_0 + \tilde{\sigma}_\zeta^2)$, $P \to P_0\tilde{\sigma}_\zeta^2/(P_0 + \tilde{\sigma}_\zeta^2)$.*

Proof. All statements follow directly from (35) and (34): $\exp(2\bar{\Delta}_t)$ is increasing in $\bar{\Delta}_t$ for $\bar{\Delta}_t > 0$, so the numerator of K grows while the denominator of P grows. □

Remark 1 (Asymmetry)**.** *For $\bar{\Delta}_t < 0$ (expected market pessimism), $\exp(2\bar{\Delta}_t) < 1$ and the asset price is* less *informative. This generates an asymmetry: the Fed's inference is better in booms than in busts.*

III.B The interest rate policy

Under discretion, we know that $\hat{E}_t^F[\tilde{y}_t] = 0$ and therefore (9) implies that

$$\iota_t = \rho + \hat{\mu}_t^F + \phi\, m_t. \tag{37}$$

The interest rate is linear in the Fed's posterior estimate m_t, with constant coefficient ϕ. Certainty equivalence holds: the coefficient does not depend on P.

The economic logic behind (37) is the same as in CS: the Fed sets the rate to offset its own expected demand plus the market's expected future "mistakes." The term $\hat{\mu}_t^F$ captures the Fed's own view of demand. The term ϕm_t captures the Fed's estimate of expected future mistakes: the market expects disagreement to persist (at rate ϕ), generating expected future output gaps that feed back into current output through forward rates. The Fed accommodates the market's view to the extent that these expected mistakes put pressure on current output.

Now, we can obtain the policy response to market optimism. From (37):

$$\frac{d\iota_t}{d\hat{\mu}_t^M} = \phi\, \frac{dm_t}{d\hat{\mu}_t^M}. \tag{38}$$

When $\hat{\mu}_t^M$ rises by 1, Δ_t rises by 1, $\exp(\Delta_t)$ rises by $\exp(\Delta_t)$, and R_t rises by $\exp(\Delta_t)$. From the linearized posterior (36):

$$\frac{dm_t}{d\hat{\mu}_t^M} = K(\bar{\Delta}_t). \tag{39}$$

Therefore, the policy response to the market's belief is

$$\frac{d\iota_t}{d\hat{\mu}_t^M} = \phi\, K(\bar{\Delta}_t) = \phi\, \frac{P_0 \exp(2\bar{\Delta}_t)}{P_0 \exp(2\bar{\Delta}_t) + \tilde{\sigma}_\zeta^2}. \tag{40}$$

This is **positive** (the Fed always leans against the wind) and **increasing in** $\bar{\Delta}_t$ (the Fed leans harder when the market is more optimistic).

The result has an intuitive decomposition. The constant ϕ is the "desire to respond"—how much the Fed would like to react to disagreement if it could observe it perfectly. The term $K(\bar{\Delta}_t)$ is the "ability to respond"—how much of the true disagreement the Fed actually sees through the asset price. The state dependence comes entirely from the second term: the Fed always wants to respond by ϕ, but it can only act on what it observes, and it observes more when the market is optimistic.

Importantly, when $\bar{\Delta}_t \approx 0$, the typical state in normal times, the Kalman gain is

$K_0 = P_0/(P_0 + \tilde{\sigma}_\zeta^2)$. If the asset price noise is large relative to the prior ($\tilde{\sigma}_\zeta^2 \gg P_0$), then $K_0 \approx 0$ and the Fed's unconditional response to market optimism is near zero. The Fed essentially ignores the market price and sets rates based on its own view. The response "turns on" only when $|\bar{\Delta}_t|$ is large enough that $\exp(2\bar{\Delta}_t)$ amplifies the signal sufficiently to overcome the noise.

Figure 1 makes the contrast between the two information environments visually stark. Panel (a) is the central result: under perfect information (blue line), the Fed's response to market optimism is a flat line at $\phi = 0.85$—the Fed responds with the same intensity regardless of the level of disagreement. Under imperfect information (red curve), the response is near zero when prior disagreement is small and rises toward the perfect-information level only when disagreement is large. The blue shaded region between the two lines represents the "unobserved disagreement"—the gap between what the Fed would do if it could see the market's belief and what it actually does. This gap is largest when disagreement is moderate, precisely the regime that dominates in normal times. Panel (b) reinforces this by showing the role of Fed uncertainty: under perfect information, the response is invariant to P_0; under imperfect information, higher uncertainty shifts all three curves upward, as the Fed leans more heavily on the market price when its own prior is weak.

Panels (c) and (d) illustrate the consequences for stabilization. Panel (c) shows output gap variance: under perfect information, it equals V (the irreducible demand shock variance) regardless of the state. Under imperfect information, there is additional variance from the inference error, and this additional variance is *largest when disagreement is small*—the stabilization paradox. The Fed is least effective precisely when the problem appears manageable and the asset price signal is weakest. Panel (d) shows the underlying mechanism: the ratio P/P_0 (how much the asset price reduces inference error relative to the prior) is close to 1 when the market is pessimistic (the price is uninformative) and drops significantly when the market is optimistic (the price is highly informative). This boom-bust asymmetry in the Fed's inference quality drives the asymmetric policy response documented in Proposition 2.

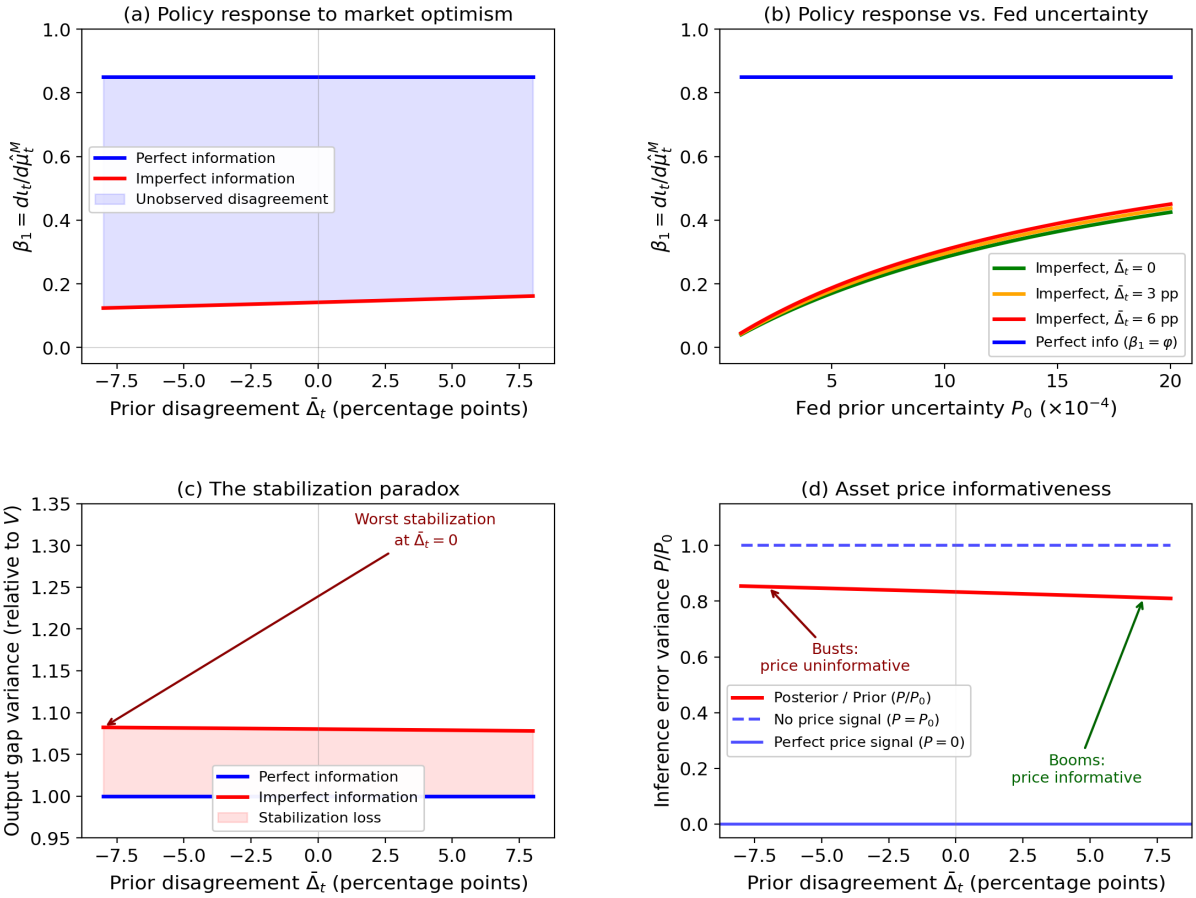


Figure 1. **Perfect vs. imperfect information about the market's belief.** Panel (a): the policy response β_1 is constant under perfect information (blue) and state-dependent under imperfect information (red). Panel (b): β_1 increases with Fed uncertainty under imperfect information but is flat under perfect information. Panel (c): output gap variance is higher under imperfect information, with the worst stabilization occurring when disagreement is small. Panel (d): the asset price reduces inference error more when the market is optimistic (asymmetry between booms and busts).

III.C Asymmetry: Booms vs. Busts

The exponential nonlinearity $\exp(\Delta_t)$ generates a natural asymmetry between optimistic and pessimistic markets. Define the local response coefficient:

$$\beta_1(\bar{\Delta}_t) \equiv \frac{d\iota_t}{d\hat{\mu}_t^M} = \phi\, K(\bar{\Delta}_t). \tag{41}$$

Then, the following proposition illustrates the asymmetric response.

Proposition 2 (Asymmetric response)**.** *For $\bar{\Delta}_t > 0$ (expected market optimism), $K > K_0$ and the Fed responds more than at baseline. For $\bar{\Delta}_t < 0$ (expected market pessimism),*

$K < K_0$ and the Fed responds less. Specifically:

$$\beta_1(\bar{\Delta}_t) - \beta_1(-\bar{\Delta}_t) = \phi P_0 \tilde{\sigma}_\zeta^2 \frac{\exp(2\bar{\Delta}_t) - \exp(-2\bar{\Delta}_t)}{(P_0 \exp(2\bar{\Delta}_t) + \tilde{\sigma}_\zeta^2)(P_0 \exp(-2\bar{\Delta}_t) + \tilde{\sigma}_\zeta^2)} > 0. \tag{42}$$

Proof. Follows from $K(\bar{\Delta}_t) > K(-\bar{\Delta}_t)$ for $\bar{\Delta}_t > 0$, since $\exp(2\bar{\Delta}_t) > \exp(-2\bar{\Delta}_t)$. □

The Fed leans harder against the wind in booms than in busts. This is not because the Fed has asymmetric preferences—the loss function is symmetric. It is because asset prices are more informative about bullish sentiment than bearish sentiment. In a boom, asset prices are elevated and highly sensitive to the market's belief. In a bust, prices are depressed and compressed, providing less signal.

This asymmetry connects to a recurring theme in central banking: the perception that the Fed responds more aggressively to overheating than to slack. Our model provides a novel, information-based explanation for this pattern. The Fed is not behaving asymmetrically by choice—it is behaving asymmetrically because the information environment is asymmetric.

This generates a testable prediction: after controlling for the magnitude of disagreement, the Fed's response should be stronger when the market is optimistic than when it is pessimistic.

III.D Double gradualism

Under perfect information, there is a single source of gradualism: the Fed expects the market to learn from data and converge toward the Fed's view, so it expects to implement its own belief gradually (Caballero and Simsek 2022, Proposition 1). Under imperfect information, a second layer of gradualism emerges: the Fed itself is learning the market's view from sequential asset price observations.

Figure 2 illustrates this "double gradualism" for the case in which the market becomes more optimistic. Under perfect information (solid blue line), the Fed immediately raises the rate to partially accommodate the market's view. Under imperfect information (dashed blue line), the Fed barely responds initially because the asset price at low disagreement is uninformative. The blue shaded area between the two lines represents the additional, information-based gradualism. Over subsequent periods, past disagreements are revealed through realized demand, and the Fed gradually catches up to the

perfect-information rate path.

The two layers of gradualism have different economic sources and different persistence. The first layer (present under both information regimes) arises because the Fed *chooses* to go slowly—it expects the market to learn and converge. The second layer (unique to imperfect information) arises because the Fed *has to* go slowly—it cannot read the market's view clearly enough to respond fully. The total gradualism is the sum of both.

When belief persistence is high (ϕ large, left panel of Figure 2), both layers are large and the gap between perfect and imperfect information is substantial. When belief persistence is low (ϕ small, right panel), disagreements resolve quickly and both layers are small—imperfect information matters less when data is informative enough to resolve disagreements fast.

Figure 2 makes the two layers of gradualism visually distinct. In each panel, the red shading between the market's expected rates (solid red) and the baseline represents the market's expectation of future policy "mistakes"—the standard source of gradualism that is present under both information regimes. The market thinks the Fed will set rates too low (relative to the market's optimism) in every future period, and this expectation depresses forward rates and feeds back into current output. The blue shading between the perfect-information Fed rate (solid blue) and the imperfect-information Fed rate (dashed blue) represents the *additional* gradualism that is unique to our model. This gap reflects the Fed's inability to read the market's belief precisely: the Fed knows the market may have become more optimistic, but the asset price signal at the short end is too noisy to confirm this, so the Fed barely raises rates initially.

Comparing the two panels reveals how belief persistence governs the relative importance of the two layers. In the high-persistence case (left, $\varphi = 0.85$), both layers are large because disagreements decay slowly: the market expects prolonged mistakes, and the Fed's own learning is also slow. The blue gap remains substantial for several periods before the Fed catches up to the perfect-information path. In the low-persistence case (right, $\varphi = 0.50$), disagreements resolve quickly, the market expects mistakes to dissipate fast, and the Fed's learning is also rapid—the blue gap closes within two or

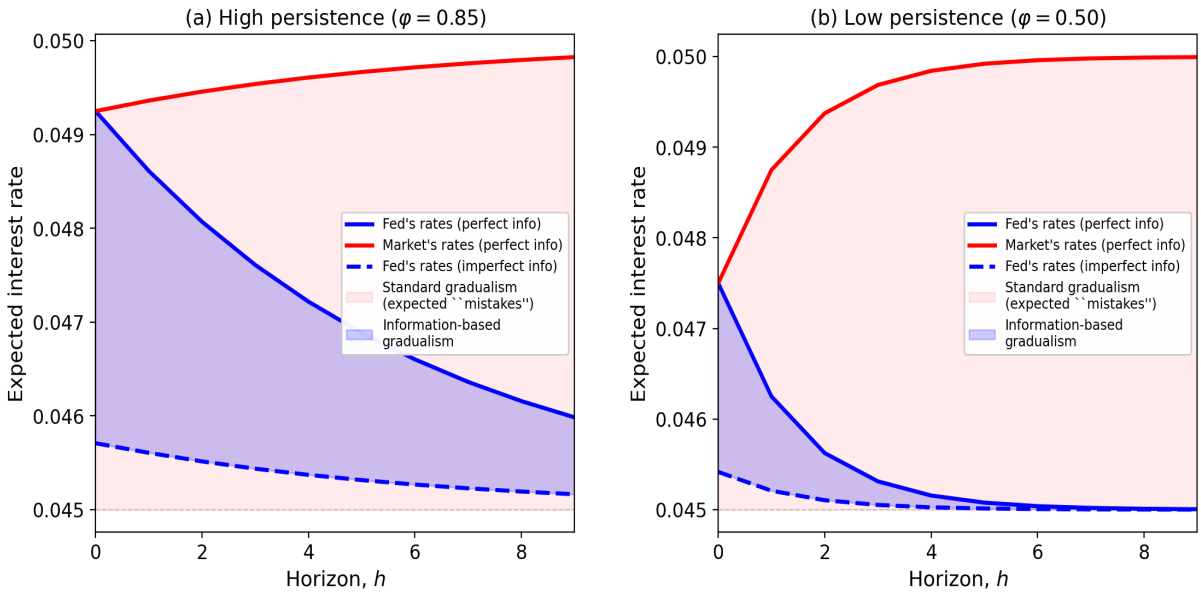


Figure 2. **Double gradualism when the market becomes optimistic.** Solid lines: expected interest rate paths under perfect information. Dashed blue line: the Fed's expected rate path under imperfect information. Red shading: gradualism from expected future "mistakes" (present under both information regimes). Blue shading: additional gradualism from imperfect observation of the market's belief (unique to imperfect information). Left panel: high persistence ($\varphi = 0.85$). Right panel: low persistence ($\varphi = 0.50$).

three periods. This suggests that the information-based gradualism is most empirically relevant in environments where disagreements are persistent, such as episodes of structural change, financial crises, or shifts in the monetary policy regime where both the Fed and the market are forming views about a new economic landscape.

III.E The role of Fed uncertainty and disagreement

The central empirical predictions emerge from the comparative statics of the local response coefficient $\beta_1(\bar{\Delta}_t) = \phi K(\bar{\Delta}_t)$. We show that β_1 increases both with the level of disagreement and with the Fed's uncertainty about its own view, and that these two channels interact.

Prediction 1: β_1 increases with disagreement. From Proposition 1:

$$\frac{\partial \beta_1}{\partial \bar{\Delta}_t} = \phi \frac{2P_0\tilde{\sigma}_\zeta^2 \exp(2\bar{\Delta}_t)}{(P_0 \exp(2\bar{\Delta}_t) + \tilde{\sigma}_\zeta^2)^2} > 0 \quad \text{for all } \bar{\Delta}_t. \tag{43}$$

The Fed responds more to market optimism when the market is already optimistic, because the asset price is more informative. This is the key result: even though

the *unconditional* relationship between the interest rate and market optimism may be statistically insignificant (because $K_0 \approx 0$ in normal times), the relationship becomes strong and significant when restricted to periods of large disagreement.

Empirically, this suggests splitting the sample by the size of the disagreement between the Fed's Greenbook/Tealbook forecast and the market's (Blue Chip) forecast—the same disagreement measure used in CS Figure 2. In periods of large disagreement, the regression of the interest rate on the market's perceived output gap should yield a significantly positive coefficient, while in periods of small disagreement, the coefficient should be near zero.

Prediction 2: β_1 increases with Fed uncertainty. The Fed's prior variance is $P_0 = \omega_s^2 \sigma_\xi^2$. Higher P_0 means the Fed is less confident in its own view. From (40):

$$\frac{\partial \beta_1}{\partial P_0} = \phi \, \frac{\tilde{\sigma}_\zeta^2 \exp(2\bar{\Delta}_t)}{(P_0 \exp(2\bar{\Delta}_t) + \tilde{\sigma}_\zeta^2)^2} > 0 \quad \text{for all } \bar{\Delta}_t. \tag{44}$$

The Fed leans harder against the wind when it is more uncertain about its own view, because it puts more weight on the market price. This is the standard Bayesian logic: an agent with a diffuse prior relies more heavily on incoming signals. In this context, a Fed that is unsure about the state of the economy defers more to the market's view as expressed in asset prices.

In the model, P_0 is driven by $\sigma_\xi^2 = 2\sigma_\delta^2(1 - \rho_c)$—the variance of interpretation differences. Empirically, P_0 can be proxied by several observable measures: disagreement among FOMC participants (as revealed by the dispersion of the dot plot), the magnitude of recent Greenbook forecast revisions, the width of the fan chart in the Summary of Economic Projections, or broader measures of macroeconomic uncertainty. Note that P_0 also increases when the demand shock is very volatile—high σ_ν^2 raises ϕ (beliefs are more stubborn when data is noisy) and raises the prior variance, making the Fed more reliant on the market's signal.

The interaction. The cross-partial:

$$\frac{\partial^2 \beta_1}{\partial \bar{\Delta}_t \, \partial P_0} = \phi \, \frac{2\tilde{\sigma}_\zeta^2 \exp(2\bar{\Delta}_t)[\tilde{\sigma}_\zeta^2 - P_0 \exp(2\bar{\Delta}_t)]}{(P_0 \exp(2\bar{\Delta}_t) + \tilde{\sigma}_\zeta^2)^3}. \tag{45}$$

The sign depends on whether $P_0 \exp(2\bar{\Delta}_t) \lessgtr \tilde{\sigma}_\zeta^2$:

- When the prior dominates the signal ($P_0 \exp(2\bar{\Delta}_t) < \tilde{\sigma}_\zeta^2$): the cross-partial is **positive**. Higher Fed uncertainty amplifies the effect of market optimism. Intuitively, when the Fed is already uncertain and the asset price signal is relatively noisy, any improvement in signal quality (from a more optimistic market) is especially valuable.
- When the signal dominates the prior ($P_0 \exp(2\bar{\Delta}_t) > \tilde{\sigma}_\zeta^2$): the cross-partial is **negative**. The asset price is already so informative that additional Fed uncertainty has diminishing returns.

In practice, the first regime is likely more relevant: the asset price is usually noisier than the prior, so the interaction is positive. This means the increase in β_1 with disagreement should be more pronounced when the Fed is also uncertain ($\bar{\Delta}_t > 0$), reflecting the joint amplification of both channels.

These predictions provide a joint test that is specific to our mechanism: no other model we are aware of predicts that β_1 increases with both market optimism and Fed uncertainty, with a positive interaction term.

IV Welfare: Is Imperfect Information a Blessing or a Curse?

The analysis so far has compared the Fed's policy response under perfect and imperfect information, treating the perfect information benchmark as the ideal. In this section, we challenge this presumption. When market sentiment contains noise, i.e., when the market's belief is sometimes driven by non-fundamental factors, imperfect information can *improve* welfare by preventing the Fed from chasing sentiment-driven fluctuations.

IV.A Sentiment noise in the data generating process

To evaluate welfare, we need to take a stand on the data generating process (DGP). Suppose under the DGP, the correct interpretation of the public signal is

$$c_t^{DGP} = \lambda^F c_t^F + \lambda^M c_t^M + \epsilon_t^{DGP}, \tag{46}$$

where ϵ_t^{DGP} is a zero-mean residual uncorrelated with (c_t^F, c_t^M). When $\lambda^M < 1$, the market's interpretation overshoots the truth—the market sometimes gets too excited or too gloomy relative to fundamentals. We refer to the component of the market's belief that does not reflect the DGP as *sentiment noise*.

Under the DGP, the actual disagreement $\Delta_t = \hat{\mu}_t^M - \hat{\mu}_t^F$ can be decomposed into a *fundamental* component Δ_t^{DGP} (the disagreement that would arise if interpretations were correct) and a *sentiment* component $\Delta_t - \Delta_t^{DGP}$ (the excess disagreement from the market's overreaction):

$$\Delta_t = \underbrace{\Delta_t^{DGP}}_{\text{genuine signal}} + \underbrace{(\Delta_t - \Delta_t^{DGP})}_{\text{sentiment noise}}. \tag{47}$$

Welfare under perfect information Under perfect information, the Fed observes Δ_t directly and sets $\iota_t = \rho + \hat{\mu}_t^F + \phi\Delta_t$. The welfare-relevant expected squared output gap under the DGP is

$$E^{DGP}[\tilde{y}_t^2]\big|_{\text{PI}} = V + \phi^2 \cdot \text{Var}^{DGP}(\Delta_t - \Delta_t^{DGP}). \tag{48}$$

The first term $V = \hat{\Sigma} + \sigma_\nu^2$ is the irreducible demand shock variance. The second term is the cost of accommodation: the Fed responds to Δ_t with weight ϕ, but part of Δ_t is sentiment noise. Every unit of sentiment noise that the Fed chases creates real output gap volatility.

Under perfect information, the Fed has no choice but to respond to the full Δ_t. In other words, it can see the market's belief perfectly and must accommodate it to offset the forward-rate pressure. The Fed is, in a sense, a *prisoner of its own information*: it knows exactly what the market thinks, and this knowledge forces it to partially validate sentiment-driven beliefs.

Welfare under imperfect information Under imperfect information, the Fed responds to its estimate m_t rather than the true Δ_t. The welfare-relevant expected squared output gap is

$$E^{DGP}[\tilde{y}_t^2]\big|_{\text{II}} = V + \phi^2 K^2 \cdot \text{Var}^{DGP}(\Delta_t - \Delta_t^{DGP}) + \phi^2(1-K)^2 \cdot \text{Var}^{DGP}(\Delta_t^{DGP}). \tag{49}$$

The first extra term is the **attenuated** cost of chasing sentiment noise: the Fed only responds to what it sees ($K\Delta_t$), so sentiment noise is dampened by $K^2 < 1$. The second extra term is the cost of **missing genuine information**: the Fed fails to respond to the real component of the market's signal by a factor of $(1-K)^2$.

IV.B When is imperfect information welfare-improving?

Comparing (48) and (49), imperfect information is welfare-improving when

$$\underbrace{(1-K^2)\,\mathrm{Var}^{DGP}(\Delta_t-\Delta_t^{DGP})}_{\text{gain: reduced chasing of noise}} > \underbrace{(1-K)^2\,\mathrm{Var}^{DGP}(\Delta_t^{DGP})}_{\text{loss: missed genuine signal}}, \tag{50}$$

which simplifies to

$$\frac{\mathrm{Var}^{DGP}(\Delta_t-\Delta_t^{DGP})}{\mathrm{Var}^{DGP}(\Delta_t^{DGP})} > \frac{1-K}{1+K}. \tag{51}$$

Since $\frac{1-K}{1+K} < 1$ for all $K > 0$, imperfect information can be welfare-improving even when the market's belief is *somewhat* informative, as long as the noise-to-signal ratio exceeds the threshold $\frac{1-K}{1+K}$.

Proposition 3 (Welfare effect of imperfect information)**.** *Define the noise-to-signal ratio* $\eta = \mathrm{Var}^{DGP}(\Delta_t-\Delta_t^{DGP})/\,\mathrm{Var}^{DGP}(\Delta_t^{DGP})$*. Then:*

1. *If* $\eta > \frac{1-K}{1+K}$*: imperfect information is welfare-improving (**blessing**).*
2. *If* $\eta < \frac{1-K}{1+K}$*: imperfect information is welfare-reducing (**curse**).*
3. *If* $\eta = 0$ *(no sentiment noise): imperfect information is always a curse.*
4. *If* $\eta \to \infty$ *(pure sentiment): imperfect information is always a blessing.*

Proof. Points 1 and 2 follow directly from (51). For point 3, when $\eta = 0$, the left side of (51) is zero while the right side is positive. For point 4, when $\eta \to \infty$, the left side diverges while the right side is bounded. □

The nonlinear asset price as an endogenous filter The welfare result takes on a deeper significance when combined with the state dependence of K. In our model, $K(\bar{\Delta}_t)$ is increasing in $|\bar{\Delta}_t|$: the Fed sees more when disagreement is large. This means the nonlinear asset price automatically acts as a filter with desirable properties:

- **Small disagreements** ($|\bar{\Delta}_t|$ near zero): K is low. The Fed largely ignores the market. If small disagreements are more likely to be noise-driven (as is plausible—it takes a strong fundamental signal to generate large disagreement), then the Fed's inattention is welfare-improving.
- **Large disagreements** ($|\bar{\Delta}_t|$ far from zero): K is high. The Fed responds strongly.

If large disagreements are more likely to reflect genuine fundamentals, then the Fed's responsiveness is welfare-improving.

In other words, the exponential nonlinearity creates a filter that lets small sentiment fluctuations pass through without affecting policy, while transmitting large, conviction-driven belief shifts into the interest rate. This is precisely the filtering property that an optimally designed information structure would have.

Remark 2 (The optimal level of asset price noise)**.** *In the mixed case (both signal and noise in the market's belief), there exists an optimal level of asset price noise* $\tilde{\sigma}_\zeta^{2*}$*—and hence an optimal Kalman gain* K^**—that minimizes welfare losses. The optimum equates the marginal benefit of filtering noise with the marginal cost of missing signal. This implies that policies aimed at improving the transparency of asset prices (reducing* $\tilde{\sigma}_\zeta^2$*) are not unambiguously welfare-improving. If market sentiment is noisy, making prices more transparent can harm welfare by forcing the Fed to chase noise it would otherwise ignore.*

IV.C Policy implications

The welfare analysis overturns the conventional wisdom that the Fed should always seek better information about the market's belief. Instead, our model suggests the following

Distinguish signal from noise. The Fed should invest in understanding the *composition* of market sentiment. That is, how much reflects genuine disagreement about fundamentals versus noise from sentiment shocks, technical trading, or liquidity events. The relevant question is not "what does the market think?" but "how much of what the market thinks is real?"

Over-engineer price transparency may backfire. Policies that reduce asset price noise (e.g., improving market microstructure, reducing dark pool trading) are beneficial only if the market's beliefs are predominantly fundamental. If market sentiment is noisy, some degree of price opacity is welfare-improving because it insulates the Fed from chasing noise.

The state-dependent filter is a feature, not a bug. The fact that the Fed's inference is poor when disagreement is small and good when disagreement is large is often seen as a limitation. Our welfare analysis shows it can be an advantage: the Fed automatically ignores small (likely noisy) sentiment fluctuations while responding to large (likely genuine) belief shifts. Central banks that attempt to "improve" their market-reading capabilities uniformly—making K high for all levels of disagreement—may inadvertently amplify the welfare cost of sentiment noise.

V Data

We focus on forecasts of real GDP growth and CPI inflation to study the Fed's response function. The Fed's forecasts are from the Teal Book on the forecasts of FOMC members. To measure these forecasts, we use the latest values in each month as the Fed's forecasts. Our data on the market's forecasts are the Blue Chip Financial Forecasts (BCFF), a monthly survey of professional forecasters. The survey collects forecasts for economic growth and inflation. The survey includes between 30 and 50 participating institutions in each wave, and individual forecasters are identified by name. Since Teal Book data is not available every month, we only focus on the months with available Teal Book data. The latest data available for the Teal Book is December 2020. Therefore, our final sample runs from 1993 through December 2020, yielding a total of 216 monthly survey waves.

The BCFF survey reports forecasts for quarterly horizons from the current quarter up to five quarters ahead. Let $E_t^{(j)} y_{t+h}$ denote the forecast made by forecaster j at time t for variable y at horizon $t+h$. Since the survey is conducted monthly but forecasts refer to end-of-quarter values, the forecast horizon h (measured in months) varies with both the survey month and the forecasted quarter. For instance, in the January 2000 survey, the one-quarter-ahead forecast corresponds to June 2000, implying $h = 5$ months. Prior to 1997, forecasts extend only four quarters ahead. To construct summary measures of the Fed's and the market's beliefs, we average forecasts across horizons $h = 0, \ldots, 6$. Specifically, for each horizon h, we first compute the cross-sectional median forecast across forecasters. We then take the average of these median forecasts across horizons $h = 0$ to 6. We apply this procedure to each variable of interest (e.g., inflation) and refer to the resulting measure as the market's forecast.

For macroeconomic variables, we retrieve quarterly real GDP growth, CPI inflation, and the unemployment rate from the St. Louis Fed (FRED). We follow Bauer et al. (2024) to calculate output gap forecasts. In particular, we impute forecasts for the level of real GDP from GDP growth forecasts and take projections of potential GDP from the Congressional Budget Office (CBO).

Our preferred measure of the Fed's interest rate is the shadow rate constructed by Wu and Xia (2016). The shadow rate extends the conventional policy rate by allowing it to become negative when the zero lower bound is binding, thereby capturing the effects of unconventional policy tools such as quantitative easing and forward guidance. When the policy rate is above the lower bound, the shadow rate closely tracks the federal funds rate, while during periods when the policy rate is constrained at zero it provides a continuous measure of monetary policy accommodation derived from the yield curve. As a result, the shadow rate offers a more informative summary of the overall stance of monetary policy than the observed federal funds rate alone.

VI Empirical Evidence

VI.A Disagreement between the Market and the Fed

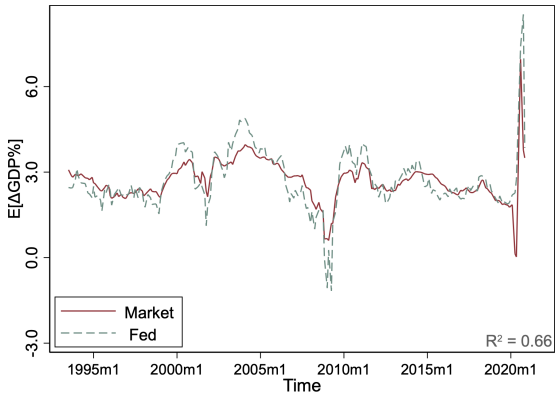
Figure 3 plots the time series of average near-term forecasts held by the market and the Fed for the period January 1993 to December 2020. Panel A shows annualized real GDP growth forecasts; Panel B shows CPI inflation forecasts. In both panels, each series is the horizon average of forecasts from $h = 0$ to 6 quarters ahead, with the red solid line giving the market's belief and the teal dashed line giving the Fed's belief.

Several features of Figure 3 are noteworthy. First, the market's and the Fed's beliefs broadly co-move, with correlations of approximately 0.7 for both GDP growth and inflation. At the same time, there is a fair and persistent degree of disagreement throughout the sample, and the market's forecasts are noticeably more stable than the Fed's: the Teal Book revises more aggressively at cyclical turning points, while the BCFF median moves gradually. As a result, GDP growth and inflation expectations tend to be higher for the market than for the Fed during expansions, when the Fed has already begun incorporating early signs of a slowdown, and lower during recessions, when the Fed has

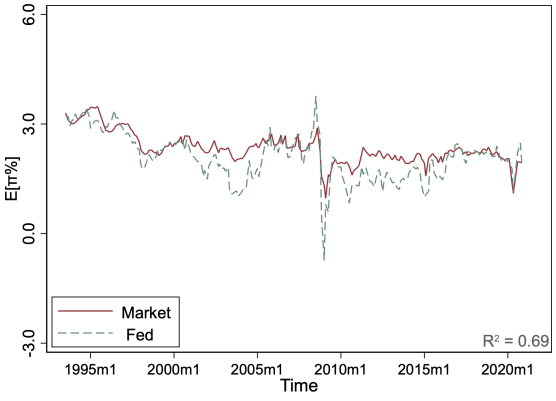
Figure 3. Disagreement between the Market and the Fed

This figure plots the time series of average near-term forecasts for the market and the Fed over the sample period January 1993 to December 2020. For each month t, the market forecast is computed as the cross-sectional median of Blue Chip Financial Forecasts (BCFF) responses, averaged across horizons $h = 0$ to $h = 6$ quarters ahead. The Fed forecast is the corresponding horizon average of Teal Book projections, using the latest available release in each calendar month. Panel A plots the annualized real GDP growth forecasts; Panel B plots the CPI inflation forecasts. In both panels, the red solid line gives the market's forecast and the teal dashed line gives the Fed's forecast.

Panel A: GDP Growth Forecasts



Panel B: CPI Inflation Forecasts



already started revising its outlook upward. Since the Teal Book is not publicly released contemporaneously, these disagreements represent real-time differences in outlook rather than mechanical anchoring to disclosed Fed targets.

VI.B The Fed's Response to Market Optimism

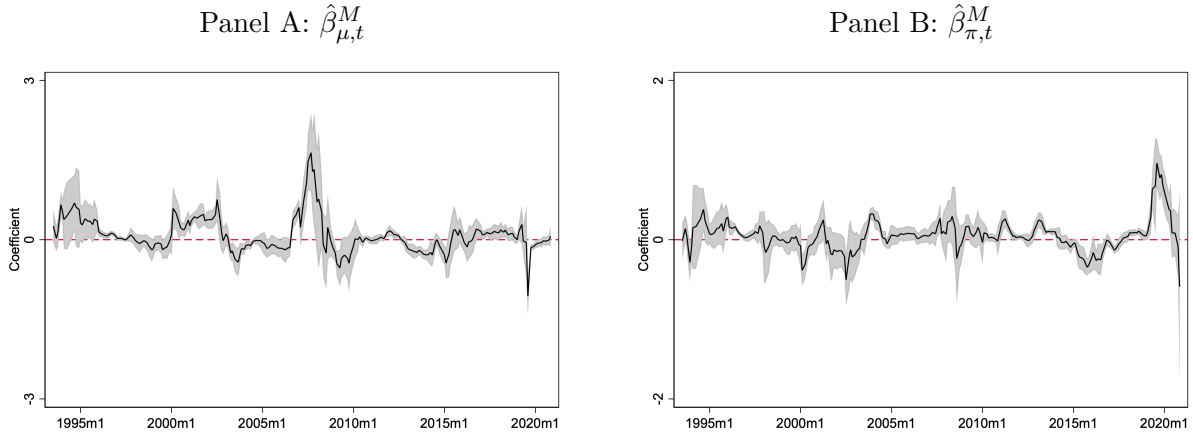
Our empirical analysis proceeds in two stages. In the first stage, we estimate a rolling-window regression to obtain a time-varying measure of how strongly the shadow rate co-moves with the market's GDP growth expectation, holding the Fed's own forecasts fixed. For each center month t, we restrict the sample to all observations within a two-year surrounding window around t and estimate by OLS:

$$\text{rate}_\tau = \alpha_t + \beta^M_{\mu,t}\big(\mu^M_\tau - \mu^F_\tau\big) + \beta^M_{\pi,t}\big(\pi^M_\tau - \pi^F_\tau\big) + \gamma_t' X_\tau + \epsilon_\tau, \tag{52}$$

where $t-12 \le \tau \le t+12$, rate_τ is the shadow rate, μ^M_τ, μ^F_τ, π^M_τ, and π^F_τ are the market's and Fed's horizon-averaged GDP growth and CPI inflation forecasts, respectively, and X_τ is a set of controls including the Fed's own forecasts of GDP growth and CPI inflation,

Figure 4. Rolling Coefficients on Market Forecasts

This figure plots rolling-window estimates of the sensitivity of the Wu–Xia shadow rate to the market's GDP growth expectation ($\hat{\beta}^M_{\mu,t}$, Panels A) and CPI inflation expectation ($\hat{\beta}^M_{\pi,t}$, Panels B). Each point is obtained from a two-year symmetric surrounding window OLS regression of the shadow rate on the market's and Fed's horizon-averaged GDP growth and CPI inflation forecasts, the output gap, realized CPI inflation, and the lagged shadow rate. Forecasts are averaged over horizons $h = 0$ to $h = 6$ quarters ahead. Shaded bands are pointwise 95% confidence intervals based on HAC standard errors with a 3-month bandwidth. The sample runs from January 1993 to December 2020.



the output gap, and realized CPI inflation. The coefficients $\beta^M_{\mu,t}$ and $\beta^M_{\pi,t}$ therefore measure the Fed's response to the market's forecast *over and above* what the Fed's own projection already implies for the rate level. A positive $\beta^M_{\mu,t}$ means that the Fed sets rates higher precisely because the market is more optimistic than itself—a lean-against-the-wind motive operating through the informational wedge. Standard errors are HAC-based with a Bartlett kernel and 4-month bandwidth.

Figure 4 plots the rolling estimates $\beta^M_{\mu,t}$ and $\beta^M_{\pi,t}$ along with their 95% confidence bands. The output growth sensitivity (Panel A) displays substantial time variation over the sample. It is significantly positive and elevated during the economic expansions of the late 1990s and mid-2000s, indicating that in those periods the Fed systematically set higher rates the more bullish the market's near-term growth outlook was relative to the Teal Book projection, even conditioning on the Fed's own forecast. The coefficient turns near zero or modestly negative in the early 2010s, when the zero lower bound was binding and the Fed was operating through unconventional tools whose relationship to the market's growth signal is necessarily attenuated. The inflation sensitivity (Panel B)

exhibits much smaller and largely insignificant variation throughout, consistent with the stability of inflation expectations noted above.

VI.C Determinants of the Fed's Sensitivity to Market Growth Expectations

We then examine what determines the time variation in the Fed's response function to market forecasts. Because inflation is low and stable over the sample period and BCFF headline CPI forecasts reflect transitory fluctuations of limited relevance for monetary policy, following Bauer et al. (2024) we focus exclusively on the output growth disagreement coefficient $\beta^M_{\mu,t}$ and treat the inflation disagreement coefficient as secondary throughout.

We regress the rolling estimates $\beta^M_{\mu,t}$ on contemporaneous measures of disagreement, financial conditions, and market sentiment:

$$\beta^M_{\mu,t} = a + b_1\big(\mu^M_t - \mu^F_t\big) + b_2\,\mathrm{SD}(Ret)_t + b_3\big(\mu^M_t - \mu^F_t\big) \times \mathrm{SD}(Ret)_t + X_t'\Gamma + e_t, \qquad (53)$$

where $\mathrm{SD}(Ret)_t$ is the standard deviation of daily S&P 500 returns over the prior three calendar months and X_t is a vector of controls including the Fed's own forecasts of GDP growth and CPI inflation, the output gap, and realized CPI inflation. Column (1) includes only the two disagreement levels and the Fed's own forecasts μ^F_t and π^F_t. Column (2) adds macroeconomic controls: output gap, realized CPI inflation, and the lagged shadow rate. Column (3) further includes financial and sentiment controls: the VIX, the 10-year minus 3-month yield spread, the Michigan Consumer Sentiment Index, and the FOMC minute-based sentiment score. Column (4) introduces the interaction $(\mu^M_t - \mu^F_t) \times \mathrm{SD}(Ret)_t$. As a robustness check, column (5) repeats column (4) using forecasts averaged over $h = 0$ to $h = 4$ horizons, and column (6) repeats column (4) using a three-year surrounding window instead of the baseline two-year window. All specifications are estimated with HAC standard errors with a 3-month bandwidth.

The coefficient on $\mu^M_t - \mu^F_t$ is positive and statistically significant across all six specifications in Table I, with estimates ranging from 0.234 to 0.407. Recall that the dependent variable is the rolling coefficient $\beta^M_{\mu,t}$ itself, which measures how strongly the shadow rate co-moves with the market's GDP growth forecast in the surrounding two-year window after conditioning on the Fed's own Teal Book projection. A positive coefficient on

TABLE I. Determinants of the Fed's Sensitivity to Market Growth Expectations

The dependent variable is $\hat{\beta}^M_{\mu,t}$, the rolling-window OLS coefficient on the market's GDP growth forecast from a local projection of the shadow rate, as described in Section II.B. Columns (1)–(4) and (6) use forecast horizons $h = 0$–6; column (5) uses $h = 0$–4 as a robustness check. Columns (1)–(5) use a two-year surrounding window; column (6) uses three years. $\mu^M - \mu^F$ is the market's GDP growth forecast minus the Fed's Teal Book forecast (relative market optimism). $\pi^M - \pi^F$ is the analogous inflation disagreement. μ^F and π^F are the Fed's Teal Book forecasts, included to control for the Fed's own outlook. Output Gap is realized GDP minus CBO potential. π is realized CPI inflation. Shadow Rate$_{-3}$ is the 3-month lagged Wu–Xia shadow rate. SD(Ret) is the standard deviation of daily S&P 500 returns over the prior three months. VIX is the CBOE implied volatility index. Spread is the 10-year minus 3-month Treasury yield spread. Michigan Sentiment is the University of Michigan Consumer Sentiment Index. Minute Sentiment is an LLM-based hawkish-dovish score from FOMC meeting minutes (0 = fully dovish, 1 = fully hawkish).

	(1)	(2)	(3)	(4)	(5)	(6)
			$\hat{\beta}^M_{\mu,t}$			
$\mu^M - \mu^F$	0.242*	0.328***	0.382***	0.407***	0.234***	0.234***
	(0.134)	(0.116)	(0.093)	(0.088)	(0.076)	(0.052)
SD(Ret)			−0.072**	−0.071**	−0.072*	−0.019
			(0.033)	(0.032)	(0.039)	(0.018)
SD(Ret) × ($\mu^M - \mu^F$)				−0.088***	−0.074***	−0.040***
				(0.028)	(0.022)	(0.015)
μ^F	0.216***	0.043	0.003	−0.072	−0.024	0.151*
	(0.066)	(0.132)	(0.138)	(0.142)	(0.131)	(0.077)
π^F	0.062	0.100*	0.136***	0.101**	0.013	0.038
	(0.057)	(0.053)	(0.048)	(0.046)	(0.041)	(0.032)
$\pi^M - \pi^F$	0.043	−0.114	−0.206	−0.222	−0.148	0.084
	(0.128)	(0.162)	(0.150)	(0.145)	(0.136)	(0.084)
Output Gap		−0.040**	−0.057**	−0.060***	−0.008	−0.081***
		(0.020)	(0.025)	(0.022)	(0.035)	(0.017)
π		0.062	0.060*	0.076**	0.068*	0.038
		(0.043)	(0.033)	(0.031)	(0.039)	(0.028)
Shadow Rate$_{-3}$		0.040**	0.043**	0.052**	0.033	0.033**
		(0.017)	(0.022)	(0.021)	(0.024)	(0.014)
VIX			0.002	0.002	0.004	−0.000
			(0.004)	(0.004)	(0.004)	(0.002)
Spread			−0.004	0.021	−0.007	−0.011
			(0.029)	(0.029)	(0.036)	(0.025)
Michigan Sentiment			0.003	0.004	0.001	0.005
			(0.004)	(0.004)	(0.005)	(0.003)
Minute Sentiment			−0.378**	−0.253*	−0.294	−0.021
			(0.164)	(0.149)	(0.214)	(0.139)
N Obs	216	216	213	213	213	213

HAC standard errors (Newey–West, 3-month bandwidth) in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

$\mu^M_t - \mu^F_t$ therefore means that in months when the market is more optimistic about growth than the Fed, the Fed's shadow rate is also more responsive to any given upward revision in the market's growth expectation. This is consistent with the lean-against-the-wind mechanism. The key to interpreting this result is the conditioning on μ^F_t in equation (52): because the Fed's own forecast enters the regression, the market's signal is not proxying

for an optimistic economy that the Fed also recognizes. Rather, the result says that when the market sees a brighter outlook than the Fed, the Fed responds by leaning harder against market exuberance. From columns (5) and (6), this finding is robust to the choice of forecast horizon ($h = 0$ to 4) and to extending the rolling window to three years.

Columns (3) through (5) show that $\text{SD}(Ret)_t$, the standard deviation of daily S&P 500 returns over the prior three calendar months, enters with a negative and mostly significant coefficient of approximately -0.07. This means that in periods of elevated realized equity volatility, the rolling sensitivity $\beta^M_{\mu,t}$ is systematically lower: the Fed places less weight on the market's growth signal when asset prices are turbulent. The interpretation is natural from a signal extraction perspective. The market's growth forecast, as reflected in equity valuations and the professional forecasts that respond to them, is a noisy signal whose informational content about the true state of the economy is diluted when equity markets are driven by shifts in risk premiums, liquidity dynamics, or sentiment rather than genuine revisions to expected cash flows. Periods of high realized return volatility are precisely the periods when this noise-to-signal ratio is elevated, making it rational for a Bayesian policymaker to discount the market's optimism.

Column (4) introduces the interaction $(\mu^M_t - \mu^F_t) \times \text{SD}(Ret)_t$, which carries a negative coefficient significant at the 1% level. When equity volatility is low, the marginal effect of $\mu^M_t - \mu^F_t$ is positive and large, meaning that market optimism over and above the Fed's forecast strongly amplifies the lean-against-the-wind sensitivity. But as volatility rises, this amplifying effect is progressively eroded: at sufficiently high levels of equity volatility, the marginal effect of an optimism gap on the Fed's responsiveness approaches zero. In other words, when markets are calm the Fed reads a large market-versus-Fed optimism gap as a meaningful signal that warrants an aggressive lean; when markets are volatile, that same optimism gap is treated with considerably more skepticism, because it is harder to distinguish genuine bullishness about the growth outlook from noise-driven price movements. The finding is robust across the alternative horizon in column (5) and the three-year window in column (6), confirming that the interaction reflects a systematic feature of the Fed's information processing rather than an artifact of the baseline specification. Together, Table I paints a coherent picture of a central bank

that treats the market's growth outlook as a valuable but imperfect signal: it responds more forcefully when the market's optimism exceeds its own considered judgment, but tempers this response when the reliability of that market signal is compromised by elevated financial market turbulence.

References

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

Wu, J. C. and F. D. Xia (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking 48*(2-3), 253–291.