

Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.

ISSN 1936-2854 (Print)

ISSN 2767-3898 (Online)

Inflation Uncertainty and Endogenous Planning Horizons

Christopher Gust, Edward Herbst, David López-Salido

2026-055

Please cite this paper as:

Gust, Christopher, Edward Herbst, and David López-Salido (2026). “Inflation Uncertainty and Endogenous Planning Horizons,” Finance and Economics Discussion Series 2026-055. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2026.055>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Inflation Uncertainty and Endogenous Planning Horizons*

Christopher Gust[†] Edward Herbst[‡] David López-Salido[§]
Federal Reserve Board Federal Reserve Board Banco de España

July 7, 2026

Abstract

We develop a finite-horizon planning model in which firms choose how far ahead to plan when setting prices. Planning further ahead improves a firm's pricing decision but requires cognitive effort. We derive analytical solutions for a firm's chosen planning horizon and show that large and persistent aggregate demand or supply disturbances induce firms to plan further ahead, making inflation more sensitive to shocks and generating endogenous movements in inflation uncertainty. Quantitatively, we show that the model matches the positive relationship between the size of inflation forecast revisions and inflation uncertainty observed in the data.

JEL CLASSIFICATION: E31, D84, E71, E32

KEYWORDS: Inflation uncertainty, finite-horizon planning, inflation expectations.

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System or of the Banco de España or the Eurosystem.

[†]Email: christopher.gust@frb.gov.

[‡]Email: edward.herbst@frb.gov.

[§]Email: david.lopez-salido@bde.es.

1 Introduction

Economic uncertainty varies substantially over time, and a large literature studies how time-varying volatility shapes macroeconomic dynamics (e.g., [Bloom, 2009](#); [Justiniano and Primiceri, 2008](#); [Jurado et al., 2015](#)). While much of this literature has focused on the relationship between uncertainty and real activity, it has also emphasized that inflation uncertainty is not constant and rises and falls over time with notable increases around major macroeconomic episodes.¹ Existing work often models this time variation in inflation uncertainty as stochastic volatility in reduced-form or structural innovations, so that inflation volatility changes because the variances of exogenous shocks change over time.² That approach is useful, but it leaves open an important question: can part of the variation in inflation uncertainty arise endogenously through expectations formation and the way firms set prices?

We answer this question using the finite-horizon planning (FHP) framework of [Woodford \(2019\)](#). Under this framework, decision makers optimize over a finite horizon and, beyond that point, evaluate outcomes using a continuation value function learned from past experience to assess contingencies outside their planning window. This departure from the standard rational expectations assumption under which private-sector agents formulate fully informed, infinitely long plans has proven to be a fruitful one. In particular, it has been shown to help fit macroeconomic time series better than other behavioral models, generate realistic movements in longer-run inflation expectations, and account for survey evidence that inflation forecast errors are predictable.³

However, one limitation of the approach thus far is that the planning horizon over which agents make sophisticated plans is assumed to be fixed and exogenous. While this assumption is a useful simplification, it suppresses the decision of how far ahead an agent should plan and the associated cost or effort involved in planning over future contingencies.⁴ In this

¹See, for example, [De Polis et al. \(2024\)](#) who document substantial time variation and asymmetry in U.S. inflation risks.

²See, for example, [Primiceri \(2005\)](#) and [Stock and Watson \(2007\)](#). Relatedly, [Vavra \(2014\)](#) studies state-dependent pricing in an environment with volatility shocks. In his model, higher volatility raises aggregate price flexibility, so nominal disturbances pass through more to inflation and less to output. Our mechanism is distinct: inflation uncertainty varies even with constant innovation variances because firms' selected planning horizons are sensitive to the size of standard additive shocks.

³See [Gust, Herbst, and López-Salido \(2022\)](#) who estimate the model of [Woodford \(2019\)](#) and compare it to other behavioral models as well as the hybrid NK model that features rational expectations, habit persistence in consumption, and exogenous price indexation. [Gust, Herbst, and López-Salido \(2026\)](#) extend this work and show that the NK-FHP model can account for the initial underreaction and subsequent overreaction of inflation forecasts emphasized in [Angeletos et al. \(2020\)](#).

⁴Closely related to our work is [Boutros \(2026\)](#), who also studies a costly finite-planning margin, focusing on household consumption responses to windfall income shocks. We instead study a costly planning decision for firms' price-setting, derive closed-form solutions for a firm's selected planning horizon, and build explicitly

paper, we abandon the assumption of fixed planning and extend the FHP framework to one in which firms choose their planning horizon. In our application, a firm sets a price over multiple periods à la Calvo-style price contracts and chooses how far to plan ahead when setting their price, balancing the expected gains from optimizing over plans that better reflect future contingencies against the cost of thinking through those contingencies. We model this decision sequentially: starting from a short horizon, the firm asks whether evaluating one more future period ahead is worth its marginal cognitive cost and continues extending its horizon only while the answer is yes. We use our extended FHP framework to quantify the cost or mental effort involved in evaluating future contingencies in firms’ price-setting decisions and the degree of forward-looking, intertemporal optimization that they choose.

In our setting, the cognitive cost associated with planning captures the mental effort required to think systematically about future possibilities. Extending the planning horizon means devoting more attention and analytical effort to projecting how future states of the economy might evolve and how current decisions depend on those outcomes. When this effort is costly, firms may decide not to think too far ahead, even though doing so could in principle improve their decisions. In contrast, assuming a fixed planning horizon removes this margin altogether: firms no longer weigh the benefits of forecasting and planning over future contingencies against its cognitive burden. In our framework, as in [Woodford \(2019\)](#), expectations remain forward looking up to the chosen horizon, but beyond that point firms rely on simpler, experience-based assessments of their future profits reflected in a continuation value function.

We show that the extension of the FHP model to one in which firms choose their planning horizon matters for inflation dynamics and generates endogenous movements in one-quarter-ahead inflation volatility. When demand or supply disturbances are large and persistent, they make future conditions especially relevant for current price setting and firms choose to devote more analytical effort to projecting how future states of the economy might evolve and how current pricing decisions depend on those outcomes. In turn, firms’ longer planning horizons make inflation more sensitive to shocks, which implies an increase in the conditional variance of one-quarter-ahead inflation even if the variance of shocks is constant.

Analytically, we characterize this sequential planning-horizon decision as a stopping problem.⁵ We focus on two economically important cases: one in which demand shocks are

on the finite-horizon planning approach in [Woodford \(2019\)](#), which can be applied to a broad class of intertemporal optimization problems.

⁵Optimal stopping methods have a long tradition in economics. Classic examples include job-search models such as [McCall \(1970\)](#); for a modern treatment of inaction and threshold rules in dynamic economic problems, see [Stokey \(2009\)](#). On one-stage look-ahead rules and the monotonicity conditions under which

permanent, and one in which demand shocks are stationary but firms do not update their continuation value function. In both cases, we show that the marginal benefit of extending the horizon is monotone, implying that the sequential rule is globally optimal over the set of possible planning horizons. We also derive closed-form solutions for these cases and show that firms plan further ahead when persistent shocks move inflation sufficiently far from firm’s continuation value function.

We show that endogenous planning preserves the main empirical strengths of fixed-planning FHP models while adding a new mechanism through which inflation uncertainty varies over time. In particular, the FHP model with endogenous planning continues to generate persistent movements in inflation through firms’ longer-run beliefs and replicates the survey evidence of [Coibion and Gorodnichenko \(2015\)](#) that inflation forecast errors are predictable. Endogenous planning also gives the planning horizon a cost-based microfoundation: in the baseline calibration, firms plan about four quarters ahead on average, and average planning costs paid by resetting firms remain small. Finally, unlike the FHP model under fixed planning, endogenous planning generates endogenous movements in one-quarter-ahead inflation uncertainty. In both the model and the data, larger forecast revisions are associated with higher inflation uncertainty, and we use this relationship to identify the marginal cost associated with a firm’s choice of its planning horizon. After calibrating the planning cost based on this relationship, we show that the mechanism accounts for a modest but meaningful share of the observed variation in inflation uncertainty. Moreover, we show that the model’s endogenous mechanism lines up with elevated inflation uncertainty during the 1970s and early 1980s, the Great Financial Crisis, and the COVID/post-pandemic inflation episode.

The paper is related to a broader literature on cognitive costs, limited information, and boundedly rational expectations. Sticky-information models and inattentive-producer models make information updating infrequent or costly ([Mankiw and Reis, 2002](#); [Reis, 2006](#)), while rational-inattention models treat attention as a scarce resource that agents allocate optimally ([Sims, 2003](#); [Maćkowiak and Wiederholt, 2009](#)). Behavioral New Keynesian models provide another way to discipline departures from fully rational expectations ([Gabaix, 2020](#); [Angeletos et al., 2020](#)). Our mechanism shares the view that forming expectations is costly, but differs in the margin it endogenizes: firms choose how far ahead to solve the dynamic pricing problem, rather than how often to update information or how precise a signal to acquire.

The paper is also related to models of endogenous price review and information acquisition. [Woodford \(2009\)](#) studies state-dependent pricing under information constraints, [Alvarez](#)

they are optimal, see [Chow et al. \(1971\)](#) and [Ferguson \(2008\)](#).

et al. (2011) separate observation costs from menu costs, and Paciello and Wiederholt (2014) analyze endogenous information acquisition and monetary policy. These papers endogenize when firms review information, whether they adjust prices, or what information they acquire. In contrast, we keep the familiar Calvo reset structure and endogenize the degree to which firms look ahead and evaluate future contingencies. This distinction lets us compare our setup directly with fixed-planning FHP models while giving the planning horizon a cost-based microfoundation.

It is important to note that our contribution is not to provide a complete theory of inflation uncertainty. Inflation uncertainty reflects many forces. Instead, we identify and quantify one endogenous contributor: the state-dependent planning horizon of price-setting firms. The analytical results make the mechanism transparent, and the quantitative results measure its importance in a calibrated NK-FHP model.

The rest of the paper proceeds as follows. Section 2 develops the endogenous-planning FHP model. Section 3 derives the implied inflation dynamics and characterizes the firm’s planning decision. Section 4 calibrates the model and presents the quantitative results on planning horizons, planning costs, forecast predictability, and inflation uncertainty. The last section concludes.

2 Model

We extend the finite-horizon planning model of Woodford (2019) by allowing firms to choose how far ahead to plan when setting prices. We focus on the determination of inflation in a partial-equilibrium setting where aggregate demand is exogenously determined. Except for this endogenous planning horizon, the behavior of price-setting firms is the same as in Woodford (2019).

2.1 Price-Setting Firms

Prices in the economy are set by monopolistically competitive firms that face random opportunities to adjust their prices as in Calvo (1983). When a firm resets its price, it chooses how far into the future to plan. In particular, a firm recognizes that its reset price is expected to remain in effect over multiple periods and chooses a planning horizon over which the firm takes future contingencies and their implications for its pricing decision. Beyond that horizon, the firm relies on a continuation value that is easy to evaluate but does not fully reflect the future contingencies that may arise.

We formalize the firm's choice of planning horizon as a sequential one-stage planning rule in discrete time. The firm contemplates planning horizons $k \in \{0, 1, 2, \dots\}$ sequentially. For a given horizon, it chooses the reset price associated with that horizon and then decides whether to plan one stage further into the future by comparing the marginal benefit and marginal cost of doing so. We first characterize the pricing decision of a firm that plans k periods ahead, and then turn to the rule governing whether it extends the horizon from k to $k + 1$ periods.

Each price-setting firm produces a differentiated good that is aggregated into a final good, Y_t , according to:

$$Y_t = \left(\int_0^1 Y_{ft}^{\frac{\varepsilon_t-1}{\varepsilon_t}} df \right)^{\frac{\varepsilon_t}{\varepsilon_t-1}}, \quad (1)$$

where Y_{ft} denotes the quantity of good f , and $\varepsilon_t > 1$ is the time-varying elasticity of substitution across differentiated goods. In our partial equilibrium setup, fluctuations in the final good are exogenous. Specifically, we assume aggregate demand, Y_t evolves according to:

$$\hat{y}_t = \rho_y \hat{y}_{t-1} + \sigma_y e_{yt} \quad (2)$$

where $\hat{y}_t = \log(Y_t)$, $0 < \rho_y \leq 1$ and $e_{yt} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$. The elasticity ε_t evolves according to:

$$\hat{\varepsilon}_t = \rho_\varepsilon \hat{\varepsilon}_{t-1} + \sigma_\varepsilon e_{\varepsilon t}, \quad (3)$$

where $\hat{\varepsilon}_t = \log(\varepsilon_t) - \log(\bar{\varepsilon})$ and $e_{\varepsilon t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$. We treat innovations to ε_t as markup, or aggregate supply, disturbances. We include these shocks in our empirical analysis to generate plausible fluctuations in inflation. In particular, in our empirical analysis, we find that both disturbances to aggregate demand and aggregate supply are needed to match movements in inflation.

Each firm $f \in [0, 1]$ produces its differentiated good using labor N_{ft} :

$$Y_{ft} = N_{ft}^{1-\alpha}, \quad (4)$$

where the parameter $0 < \alpha \leq 1$ determines the return to labor in the production of the intermediate good f . To produce Y_{ft} , a firm hires labor in a competitive labor market, taking the real wage as given. Labor is supplied elastically to the firm so that the real wage, $\bar{w}$, is constant:

$$\bar{w} = (1 - \alpha) \frac{\bar{\varepsilon} - 1}{\bar{\varepsilon}}. \quad (5)$$

As implied by expression (5), the economy's real wage reflects the steady-state share of in-

come going to workers and the (inverse) of the steady-state price over marginal cost markup.

At each date, a fraction $1 - \theta$ of firms have the opportunity to reset their price and choose their planning horizon over which that reset price applies. In periods in which a firm does not have this opportunity, we assume that a firm indexes its price increases to the steady-state (gross) inflation rate: $P_{ft} = \bar{\pi} P_{ft-1}$. The aggregate price level can be rewritten as:

$$P_t = \left[(1 - \theta) (P_t^{k_t^*})^{1-\varepsilon_t} + \theta (\bar{\pi} P_{t-1})^{1-\varepsilon_t} \right]^{\frac{1}{1-\varepsilon_t}}, \quad (6)$$

where P_t denotes the aggregate price level, $P_t^{k_t^*}$ is the common reset price chosen by identical firms that re-optimize at time t , and k_t^* is the common planning horizon that they choose.

2.2 Pricing Decision

Conditional on choosing planning horizon k , the firm chooses its relative price, $p_{ft} \equiv \frac{P_{ft}}{\bar{\pi} P_{t-1}}$, to solve:

$$\max_{p_{ft}} H_t^k(p_{ft}, S_t). \quad (7)$$

A horizon k firm's objective function, $H_t^k(p_{ft}, S_t)$, is given by:

$$H_t^k(p_{ft}, S_t) = \mathbb{E}_t^k \left\{ \sum_{\tau=t}^{t+k} (\beta\theta)^{\tau-t} \left[\left(\frac{\bar{\pi}^{\tau-t+1} p_{ft}}{\prod_{i=0}^{\tau-t} \pi_{t+i}} \right) - \bar{w} Y_{f\tau}^{\frac{\alpha}{1-\alpha}} \right] Y_{f\tau} + (\beta\theta)^{k+1} H_t \left(\frac{\bar{\pi}^{k+1} p_{ft}}{\prod_{i=0}^k \pi_{t+i}} \right) \right\}, \quad (8)$$

where $Y_{f\tau}$, satisfies:

$$Y_{f\tau} = \left(\frac{\bar{\pi}^{\tau-t+1} p_{ft}}{\prod_{i=0}^{\tau-t} \pi_{t+i}} \right)^{-\varepsilon_\tau} Y_\tau.$$

In expression (8), $\pi_t = \frac{P_t}{P_{t-1}}$ denotes the inflation rate, $S_t = (y_t, \varepsilon_t, \pi_t)'$ denotes the aggregate state, and $\beta \in (0, 1)$. The function $H_t(\cdot)$ is a firm's continuation value function, which captures its beliefs about the expected discounted losses associated with price p_{ft} after date $t+k$. Reflecting the cognitive cost of forecasting and planning future events, this continuation value function is not fully state contingent as would be the case under rational expectations. Instead, following [Woodford \(2019\)](#), a firm's continuation value function is intended to reflect a firm's past experience and knowledge about how its pricing decisions affect its profits. As discussed below, a firm updates it as part of its pricing problem, averaging over past information that a firm receives over time. This makes a firm's continuation value function time dependent, and the firm's k -horizon objective function, $H_t^k(p_{ft}, S_t)$, reflects this time dependence and thus evolves over time as well. Importantly, a firm takes its continuation value function as given at each date t in making its pricing and planning decisions.

A firm that plans k periods ahead uses the model's structural relationships and shock processes to evaluate contingencies from date t through date $t + k$. Events beyond that horizon are not evaluated using structural relationships or the implied evolution of future reset prices; instead, they are summarized by the continuation value function, which does not fully reflect the states that may occur in the future. Because of finite planning, a firm's plans at time $t + 1$ will not generally be the same as those when it makes its decisions at time $t + 1$, since this would imply that a firm making decisions at date t evaluates contingencies through period $t + k + 1$ rather than truncating its planning at period $t + k$. Thus, $\mathbb{E}_t^k$ denotes subjective rather than rational expectations: the firm uses the model's stochastic structure only through its k -period planning horizon. For a firm's expectations to be rational, it would have to evaluate contingencies in periods $\tau \geq t + k + 1$ and beyond. Instead, a firm at date t makes plans in period $t + 1$ taking into account structural relationships and contingencies only $k - 1$ periods into the future. Similarly, in planning period $t + 2$, a firm makes plans in period $t + 2$ only taking into account structural relationships and contingencies through $k - 2$ periods into the future, and so on over the course of the k -period planning horizon.

A firm must also evaluate the expectations of future inflation in expression (8). This requires forming beliefs about the planning horizons of other firms. In order for a firm not to have to consider what the model's structural relationships imply for states outside of its k -period planning horizon, it is assumed as in [Woodford \(2019\)](#) that a firm believes that other firms share its k -period planning horizon. This assumption introduces an additional departure from rational expectations besides the truncated planning that a firm does.

Because firms plan only over a finite horizon, it is convenient to define π_t^k and p_t^k as the model-consistent solutions for inflation and a firm's common reset price, respectively, under the assumption that all firms have planning horizons of length k . More generally, for these variables, the following relationships hold:

$$\mathbb{E}_t^k \pi_{t+k-j} = E_t \pi_{t+k-j}^j \quad \text{and} \quad \mathbb{E}_t^k p_{t+k-j} = E_t p_{t+k-j}^j \quad \text{for } 0 \leq j \leq k. \quad (9)$$

Expression (9) provides a mapping between the subjective expectations operator of a firm with a k -period planning horizon and the model-consistent expectations operator. It reflects that a firm formulates its plans in period $t + k - j$ using its full knowledge of the model's structural equations and contingencies under the assumption that other firms have j periods left in their planning horizon like themselves. Thus, while E_t is the model-consistent expectations operator, the endogenous variables π_{t+k-j}^j and p_{t+k-j}^j still reflect an agent's limited planning horizon and subjective beliefs regarding the planning horizons of other firms.

2.3 Planning Decision

We now turn to characterizing a firm’s choice regarding its planning horizon—a decision it makes whenever it has the opportunity to reset its price. We model this choice as a sequential stopping problem governed by a one-stage look-ahead rule. Such rules are standard in the stopping-rule literature: the decision maker compares stopping at the current stage or horizon with continuing a horizon further and then stopping, and in monotone stopping problems the one-stage look-ahead rule achieves the global optimum (see [Chow et al., 1971](#); [Ferguson, 2008](#)).

In our setting, the stage is the planning horizon. The firm begins at $k = 0$, where its price takes into account current contingencies while future outcomes are evaluated using the continuation value function. It then asks whether extending the horizon by one stage improves its pricing decision enough to justify the additional cognitive cost. If the answer is yes, it moves to $k = 1$ and repeats the same comparison; if the answer is no, it stops and uses the reset price associated with its current horizon.

To formalize the sequential rule, we assume that a firm faces a cost function $\tau(k)$ that reflects the cognitive effort associated with choosing its price over future contingencies. This cost function is strictly increasing in the planning horizon, and we assume the incremental cost

$$\Delta\tau_{k+1} \equiv \tau(k+1) - \tau(k)$$

is weakly increasing in k . Accordingly, planning further into the future is potentially more costly for a firm. We also assume that $\tau(0) = 0$, implying that it is costless for a firm to evaluate future outcomes using its continuation value function. This assumption is consistent with the idea that it requires more cognitive effort to explicitly reason through future contingencies than to rely on a coarse continuation value that only reflects a firm’s past, average profitability.

For a firm that has already chosen to set its price taking into account contingencies through horizon k , re-solving its pricing problem over $k + 1$ contingencies yields a net value of

$$H_t^{k+1}(p_t^{k+1}, S_t) - \tau(k+1).$$

A firm could instead choose not to re-solve the pricing problem and avoid the additional cognitive effort. In that case, it re-uses p_t^k , which has been chosen for contingencies through

period $t + k$, yielding a net value of

$$H_t^{k+1}(p_t^k, S_t) - \tau(k).$$

In choosing not to re-solve the pricing problem, the firm acknowledges that there are future contingencies at date $t + k + 1$ and thus evaluates losses using the $k + 1$ objective function. While p_t^k does not generally solve the $k + 1$ objective function, it is the natural benchmark because it is the reset price chosen for the current horizon. In this sense, p_t^k plays the role of a reference price for measuring the private value loss from not solving the expanded-horizon pricing problem, in the spirit of near-rational behavior such as in [Akerlof and Yellen \(1985\)](#).⁶ Retaining this price only costs the firm $\tau(k)$ in terms of cognitive effort, while re-solving the pricing problem requires the firm to pay the incremental planning cost.

The difference between these two objects captures the marginal benefit of extending the planning horizon by one additional period. A firm evaluates these options sequentially at each horizon until the incremental cost $\Delta\tau_{k+1} \equiv \tau(k+1) - \tau(k)$ exceeds the marginal benefit of extending the horizon:

$$k_t^* = \min \{k \in \{0, 1, 2, \dots\} \mid H_t^{k+1}(p_t^{k+1}, S_t) - H_t^{k+1}(p_t^k, S_t) \leq \Delta\tau_{k+1}\}. \quad (10)$$

Equation (10) characterizes the planning horizon selected by the one-stage rule. Because the marginal benefit of looking further into the future depends on the current realization of S_t , the selected planning horizon is inherently state dependent.

2.4 Learning about the Value Function

After the shocks are realized, a firm makes its pricing decision and its decision about how far ahead to plan. It also updates its continuation value function over time following the approach in [Woodford \(2019\)](#). In particular, a firm forms a new estimate of its continuation value function, $H_t^E(p_{ft})$, which is obtained as part of its pricing problem. This estimate comes from a firm evaluating its k_t^* objective at the realized state, S_t , for any choice of p_{ft} :

$$H_t^E(p_{ft}) = H_t^{k_t^*}(p_{ft}, S_t)$$

⁶The comparison that a firm makes is analogous to that of menu-cost pricing models, such as in, for example, [Mankiw \(1985\)](#). In those models, a firm compares the gain from changing its nominal price from its current price with the cost of doing so. In our setting, the cost is not a physical cost of changing the current price, but the cognitive cost of extending the planning horizon; the benchmark is the horizon- k reset price rather than an inherited nominal price.

A firm uses this estimate to update the continuation value function to be used in its pricing problem at date $t + 1$:

$$H_{t+1}(p_{ft}) = (1 - \gamma)H_t(p_{ft}) + \gamma H_t^E(p_{ft}), \quad (11)$$

where $0 \leq \gamma \leq 1$. As in [Woodford \(2019\)](#), a firm uses an average of its past value functions to form its beliefs about its continuation value function. In effect, a firm accumulates knowledge about its continuation value function from its pricing problem with γ determining how much weight it puts on recent experience versus knowledge acquired in the distant past. A notable difference from [Woodford \(2019\)](#) is that in his setup a firm's new estimate of its value function uses the same k -period objective function at each date, while equation (11) implies that a firm's new estimate can vary over time depending on how far in the future a firm chooses to plan.

While learning introduces an additional parameter, γ , it has both theoretical and empirical appeal. Without learning, a firm's longer-run beliefs (i.e., those outside of their planning horizon) would remain fixed in response to structural changes in their environment such as a permanent increase in technology. In contrast, with learning, a firm's longer-run beliefs about inflation evolve and eventually reflect such changes. On the empirical side, [Gust et al. \(2022\)](#) show that an FHP model with learning gives rise to empirically realistic movements in longer-run inflation expectations as well as persistent inflation dynamics arising from these movements in longer-run inflation expectations. Accordingly, a firm's longer-run beliefs about inflation as reflected in its continuation value function play a particularly important role in inflation dynamics.

3 Analytical Results

We characterize the dynamics of inflation analytically by taking a second order approximation to a firm's maximization problem. Appendix A provides the details of our approach, where we derive an analytical representation between inflation and a firm's choice of its planning horizon. We also decompose inflation into a cyclical component, reflecting the effects of shocks, and a trend component, which reflects movements in inflation due to a firm's updating of its continuation value function. This decomposition allows us to express inflation as:

$$\hat{\pi}_t \equiv \hat{\pi}_t^{k_t^*} = \kappa A_y(k_t^*) \hat{y}_t + A_\varepsilon(k_t^*) \hat{\varepsilon}_{st} + \beta^{k_t^*+1} \bar{\pi}_t, \quad (12)$$

where $\hat{\pi}_t$ denotes realized inflation expressed in log-deviations from its steady state value, $\bar{\pi}_t$ is trend inflation (also expressed in log-deviations from the steady state inflation rate), and $\hat{\varepsilon}_{st}$ is the normalized supply disturbance induced by fluctuations in the elasticity of substitution, so that a positive realization of $\hat{\varepsilon}_{st}$ raises inflation. Movements in trend inflation can be interpreted as changes in a firm's longer-term beliefs about inflation (i.e., those outside its planning horizon), and thus the model relates trend inflation to firms' longer-term inflation expectations. The sensitivity of inflation to the output gap also depends on the reduced-form parameter, κ , which in turn depends on the frequency with which a firm is able to reset its price.⁷

Equation (12) highlights that inflation is linear in aggregate demand and supply disturbances as well as trend inflation, with coefficients that vary with k_t^* . In particular,

$$A_y(k_t^*) = \frac{1 - (\beta\rho_y)^{k_t^*+1}}{1 - \beta\rho_y}, \quad A_\varepsilon(k_t^*) = \frac{1 - (\beta\rho_\varepsilon)^{k_t^*+1}}{1 - \beta\rho_\varepsilon}.$$

Because $A_y(k)$ and $A_\varepsilon(k)$ increase as k increases, in periods when a firm chooses to plan further into the future, inflation is more sensitive to demand and supply disturbances. Moreover, as a firm plans further into the future, $\beta^{k_t^*+1}$ declines and the trend becomes less important in determining inflation. As shown in Appendix A, trend inflation evolves according to:

$$\bar{\pi}_{t+1} = (1 - \gamma)\bar{\pi}_t + \gamma\hat{\pi}_t. \quad (13)$$

Accordingly, a firm's belief about trend inflation next period is an exponentially weighted average of past realized inflation. And, as implied by expression (12), realized inflation embeds the time-varying effects of aggregate demand, supply disturbances, and trend inflation through a firm's selected planning horizon.

3.1 Analytical Characterization of the Planning-Horizon Decision

Using a second-order approximation of the profit function, Appendix A shows that we can rewrite equation (10) as:

$$k_t^* \equiv k^*(\hat{y}_t, \hat{\varepsilon}_{st}, \bar{\pi}_t) = \min \left\{ k \in \{0, 1, 2, \dots\} : \frac{\mathcal{K}}{2} (\hat{p}_t^{k+1} - \hat{p}_t^k)^2 \leq \Delta\tau_{k+1} \right\}, \quad (14)$$

where $\hat{p}_t^k$ denotes the log-deviation of the common reset price chosen by a firm with planning horizon k , and $\mathcal{K}$ is a reduced-form parameter determined by the curvature of a firm's demand

⁷As shown in Appendix A, $\kappa = \frac{(1-\beta\theta)(1-\theta)}{\theta}\xi$ where $\xi = \frac{a}{(1+a)\bar{\varepsilon}-1}$ and $a = \frac{\alpha}{1-\alpha}$.

curve and frequency of price setting:

$$\mathcal{K} = \left(\frac{\bar{\varepsilon}}{1 - \alpha} - 1 \right) \frac{\bar{\varepsilon} - 1}{1 - \beta\theta} > 0.$$

In expression (14) a firm's marginal benefit to extending its planning horizon by one period is governed by the one-stage reset-price increment from re-solving its pricing problem to take into account additional future contingencies:

$$\hat{p}_t^{k+1} - \hat{p}_t^k = \beta^{k+1} \left(\frac{1}{1-\theta} \right) \left[\kappa \rho_y^{k+1} \hat{y}_t + \rho_\varepsilon^{k+1} \hat{\varepsilon}_{st} - (1 - \beta) \bar{\pi}_t \right]. \quad (15)$$

A firm's selected planning horizon is therefore state dependent. The benefit of a one-period extension of its planning horizon varies with the size of the disturbances to aggregate demand and aggregate supply as well as the level of trend inflation relative to the steady state inflation rate. To derive closed-form solutions for the optimal horizon k_t^* , we assume a linear planning cost function:

$$\tau(k) = \frac{\mathcal{K}}{2} \tau_0^2 k, \quad (16)$$

where $\tau_0 > 0$. The linear planning cost function is a convenient benchmark but it is not required, even for analytical tractability: analogous expressions can be derived for other planning-cost schedules, including exponential costs, with the solutions adjusted to reflect the relevant marginal costs. In what follows, we focus on the role of the demand disturbance in determining k_t^* , setting the supply disturbance $\hat{\varepsilon}_{st} = 0$. However, a similar analysis focusing on the role of the supply disturbance follows if the demand disturbance is instead set to zero. Finally, we use the tie-breaking convention that when a firm is indifferent between stopping and extending the planning horizon, it selects the shorter horizon. We begin by establishing that a firm's choice of its planning horizon is optimal absent learning considerations—that is, where trend inflation is zero. In this case, the expression for k_t^* has a simple form.

Proposition 3.1 (k_t^* under no-learning). *Assume $0 < \rho_y \leq 1$, $\bar{\pi}_t = 0$, and $\hat{\varepsilon}_{st} = 0$. Then $k_t^* = 0$ iff*

$$|\kappa \hat{y}_t| \leq \frac{(1 - \theta) \tau_0}{\beta \rho_y},$$

and for $k \geq 1$,

$$k_t^* = k \iff \frac{(1 - \theta) \tau_0}{(\beta \rho_y)^k} < |\kappa \hat{y}_t| \leq \frac{(1 - \theta) \tau_0}{(\beta \rho_y)^{k+1}}.$$

Moreover, the marginal value of extending the planning horizon is weakly decreasing in k , so the one-stage look-ahead rule selects the unique global solution to the sequential planning

problem.

See Appendix B for the proof. The planning horizon depends on the state only through $|\hat{y}_t|$ and is weakly increasing in $|\hat{y}_t|$. A firm's horizon choice is symmetric in that positive and negative demand disturbances of equal magnitude will result in identical choices for k_t^* .

When $|\hat{y}_t|$ is small—that is, when aggregate demand is close to its steady state value—the reset price chosen under a short planning horizon is close to the reset price if the Calvo price-setting firm had infinite-horizon rational expectations (i.e., $k \rightarrow \infty$). Without learning, a firm relies on a continuation value function that reflects its steady state profits, and the mispricing from relying on this continuation value function is not large for small values of $\hat{y}_t$. In that case, the marginal benefit of extending the planning horizon is low, and the firm stops immediately. In contrast, when the demand disturbance is far from its steady state value, the marginal benefit of extending the planning horizon can be high. In that case, a short planning horizon omits important future consequences of a persistent demand disturbance. Moreover, the resulting mispricing relative to a firm with infinite-horizon rational expectations can be large, and the marginal benefit of planning further can be high, which creates an incentive for firm to extend its planning horizon.⁸

When trend inflation and supply disturbances are absent, the marginal benefit of planning further ahead declines monotonically at rate $(\beta\rho_y)^2$. The rule stops at the first horizon at which the marginal benefit of extending the horizon, net of the incremental cost, is nonpositive. Because the marginal benefit is monotonically decreasing, the selected horizon is the unique global optimum.

We next allow trend inflation to enter the planning-horizon decision. To isolate its role most cleanly, we first consider the unit-root demand case (again setting the supply disturbance to zero). In this case, the reset-price increment depends on the single state variable

$$x_t \equiv \kappa\hat{y}_t - (1 - \beta)\bar{\pi}_t, \quad (17)$$

which measures the demand disturbance relative to the trend-inflation component of the firm's continuation value. Proposition 3.2 gives the resulting closed-form solution.

Proposition 3.2 (k_t^* under unit-root demand). *Assume $\rho_y = 1$ and $\hat{\varepsilon}_{st} = 0$. Then $k_t^* = 0$*

⁸By comparing an FHP firm to one with infinite-horizon rational expectations, we focus on the mispricing arising from the truncated planning horizon of an FHP firm. Accordingly, we isolate this source of mispricing from that arising from Calvo price setting and monopolistic competition.

iff $|x_t| \leq (1 - \theta)\tau_0/\beta$, and for $k \geq 1$,

$$k_t^* = k \iff \frac{(1 - \theta)\tau_0}{\beta^k} < |x_t| \leq \frac{(1 - \theta)\tau_0}{\beta^{k+1}}.$$

Moreover, the marginal value of extending the planning horizon is weakly decreasing in k , so the one-stage look-ahead rule selects the unique global solution to the sequential planning problem.

See Appendix B for the proof. With $\rho_y = 1$ and $\hat{\varepsilon}_{st} = 0$, the optimal planning decision depends on how far the demand disturbance pushes inflation away from its trend-adjusted value, which reflects a firm's beliefs about its continuation value function. When $|x_t|$ is small—that is, when aggregate demand is close to its trend *reference point* $\hat{y}_t^R = (1 - \beta)\bar{\pi}_t/\kappa$ —a firm's marginal benefit of extending its planning horizon is low and a firm tends to choose a short planning horizon. In contrast, when the demand disturbance is far from its trend reference point, $|x_t|$ is large. In this case, a short planning horizon can omit important future consequences of a permanent demand disturbance and the marginal benefit of planning further ahead can be high, creating an incentive for a firm to extend its planning horizon. The unit-root setting clarifies the role of firms' longer-run beliefs—via trend inflation—in determining a firm's optimal planning horizon.

3.2 When is the One-Stage Rule Optimal and Symmetric?


The two propositions establish the optimality of a firm's choice of its planning horizon in simplified settings. In both cases, the marginal benefit of extending the planning horizon is proportional to the square of a single state variable multiplied by a term that declines geometrically with k . In Proposition 3.1, that state variable is $\kappa\hat{y}_t$; in Proposition 3.2, it is $x_t = \kappa\hat{y}_t - (1 - \beta)\bar{\pi}_t$. Hence the planning rule is symmetric around the relevant reference point and the marginal benefit of extending the horizon is monotone in k . The one-stage rule therefore stops at the first horizon at which the marginal benefit falls below the marginal cost, and this coincides with the global solution to the sequential planning problem.

The symmetry and monotonicity of a firm's marginal benefit to extending its planning horizon need not survive once both learning and stationary shocks are present, or once more than one disturbance enters the model. First, with stationary disturbances, trend inflation changes the symmetry of the decision rule. If demand is stationary and $\bar{\pi}_t > 0$ —and again supply disturbances are set to zero—the marginal benefit of extending the horizon from k to $k + 1$ is proportional to $(\kappa\rho_y^{k+1}\hat{y}_t - (1 - \beta)\bar{\pi}_t)^2$. A positive demand disturbance will be consistent with firms' longer-run inflationary beliefs (i.e., trend inflation), and so the marginal

benefit will be smaller than otherwise. In contrast to a positive demand shock, a negative demand disturbance will not offset the (positive) trend inflation term in the marginal benefit expression. This means that with positive trend inflation the selected planning horizon is therefore no longer symmetric around a single trend-adjusted reference point: firms tend to choose longer horizons for negative demand shocks and shorter horizons for positive demand shocks that can partially offset the positive trend component on the one-stage reset price increment.

Figure 1 illustrates this loss of symmetry. The horizontal axis measures the demand gap relative to the trend-adjusted reference point. When $\bar{\pi}_t = 0$, this reference point is zero and the planning rule is symmetric. When $\bar{\pi}_t > 0$ and demand is stationary, the red dashed line is no longer symmetric around zero: negative deviations from $\hat{y}_t^R$ call for longer planning horizons than positive deviations of the same size.

Figure 1: Planning horizon and the trend-adjusted reference point.



Note: The figure solves the endogenous-planning stopping rule under the stylized linear planning-cost schedule ($\tau_0 = 0.012$), with the supply disturbance $\hat{\epsilon}_{s,t} = 0$. The black line sets $\bar{\pi}_t = 0$; the red dashed line sets $\bar{\pi}_t = 0.50$. In both cases $\beta = 1.03^{-1/4}$, $\theta = 0.60$, $\rho_y = 0.90$, and $\kappa = 0.0104$.

The other property that need not hold in a more general setup is that the one-stage marginal benefit need not be monotone in the planning horizon. This issue is distinct from the asymmetry shown in Figure 1. Conditional on a state the sequence of reset price increments $\{\hat{p}_t^{k+1} - \hat{p}_t^k\}_{k=0,1,\dots}$ can fall and then rise. This non-monotonicity can arise when demand and supply disturbances have opposite signs and different persistence. If, for example, $\hat{y}_t > 0$,

$\hat{\varepsilon}_{st} < 0$, and $\rho_\varepsilon < \rho_y$, then the supply component may initially offset the demand component in the reset-price increment. Because the supply component decays more quickly, this offset can disappear at longer horizons. The marginal benefit can fall below the marginal cost and later rise above it again as k increases.

This non-monotonicity matters for interpretation of the stopping rule. The one-stage rule remains well defined: it selects the first horizon at which the marginal benefit of extending the horizon is no larger than the marginal cost. However, when marginal benefits are non-monotone, this first crossing need not coincide with the horizon that would be chosen by a firm that globally searched over all horizon lengths. We therefore interpret the rule as a myopic planning rule, following the stopping-rule terminology in [Ferguson \(2008\)](#). This interpretation is consistent with the bounded rationality of FHP firms. A global search over all planning horizons would require an additional layer of cognitive effort, since the firm would have to evaluate and compare the value of many possible horizon lengths before deciding how far ahead to plan. The one-stage rule keeps this planning problem sequential and incremental: at each stage, the firm asks whether extending the horizon by one stage improves its pricing decision enough to justify the incremental cognitive cost. The analytical results above establish conditions when the sequential, one-stage rule coincides with the unique global solution.

4 Calibration and Quantitative Results

This section calibrates the model and organizes the quantitative results around four points. First, endogenous planning preserves the main empirical success of the FHP model: it generates forecast-error predictability in the range of the data and persistent fluctuations in inflation through movements in trend inflation. Second, endogenous planning extends the fixed planning framework by making inflation’s sensitivity to shocks state dependent, which generates endogenous movements in inflation uncertainty. Third, endogenous planning can account for the positive relationship between inflation uncertainty and the size of forecast revisions observed in the data while the model with fixed planning horizons or rational expectations cannot. We show that the relationship between inflation uncertainty and forecast revisions provides a way of disciplining the cost of firms extending their planning horizon. Fourth, while the calibrated mechanism accounts for only a modest part of the total volatility of measured inflation uncertainty, it lines up with increases in inflation uncertainty observed during major historical episodes.

4.1 Calibration and Model Variants

We calibrate both fixed- and endogenous-planning variants of the model using quarterly U.S. data. Following [Christiano et al. \(2005\)](#), we set $\beta = 1.03^{-\frac{1}{4}}$ and $\alpha = 0.36$. Consistent with their estimates, θ is set to 0.6, which implies that the average price duration is 2.5 quarters. For the elasticity of substitution between goods, we set $\bar{\varepsilon} = 10$, in line with [Atkeson and Burstein \(2008\)](#). For the updating of a firm’s continuation value function, the gain parameter, γ is set to 0.2, which is consistent with the estimated value in [Gust et al. \(2026\)](#).

The AR(1) process for aggregate demand, y_t , is estimated using the output-gap series as measured by the Congressional Budget Office (CBO) over the 1960:Q2-2024:Q2 sample period. Estimation using this series implies that $\rho_y = 0.9$ with $\sigma_y = 1.011$. We use observed inflation over this sample period, measured as PCE core inflation, to infer the aggregate supply shock from equation (12) and estimate its AR(1) process. We do this for four different model variants shown in Table 1. This procedure lets each variant match the same historical inflation series while making transparent how much persistence must be assigned to the residual supply disturbance under each expectations/planning structure.

Table 1: Model Variants, Planning Costs, and Aggregate-Supply Shocks

Model	Planning-cost specification	ρ_ε	σ_ε
Endogenous Planning, Linear Cost	$\underline{k} = 2, \tau_0 = 0.038$	0.57	0.093
Fixed Planning	$\underline{k} = k^* = 4, \tau_0 \rightarrow \infty$	0.54	0.101
Rational Expectations	$\underline{k} = 0, \tau_0 = 0$	0.93	0.018
Endogenous Planning, Exponential Cost	$b_0 = 0.030, b_1 = 0.25$	0.56	0.094

Notes: The Rational Expectations variant is approximated by setting $k = 1000$; $\underline{k}$ is irrelevant when $\tau_0 = 0$.

The first row of the table shows our baseline calibration of the endogenous-planning model, which features a slight extension of the cost function in the analytical results.

$$\tau(k) = \frac{\mathcal{K}}{2} \left(\frac{\tau_0}{100} \right)^2 \max\{k - \underline{k}, 0\}. \quad (18)$$

This is a piecewise linear cost function. This specification allows firms to plan up to $\underline{k}$ periods at no cost and then face a constant marginal cost of extending the planning horizon for $k > \underline{k}$. We include the free-planning threshold $\underline{k}$ so that the endogenous-planning variant can match the average planning horizon estimated in [Gust et al. \(2026\)](#) and has the same average planning horizon as the fixed-planning variant. The constant marginal cost parameter, τ_0 , is

then determined from the empirical relationship between inflation uncertainty and forecast revisions as discussed below.

Existing evidence is useful for disciplining the plausibility of planning costs, but it does not provide a direct estimate of the marginal cost of extending a firm’s planning horizon for our planning problem. [Reis \(2006\)](#) calibrates information-processing costs that generate infrequent producer planning, [Zbaracki et al. \(2004\)](#) provide direct evidence that managerial and customer costs of price adjustment are economically meaningful, and [Alvarez et al. \(2011\)](#) separate observation costs from menu costs in a structural pricing model. These objects are closely related to the idea that price setting involves costly managerial attention, but they are not the same as our marginal planning-cost parameter. We therefore calibrate τ_0 using the empirical relationship between forecast revisions and inflation uncertainty, while using the implied average horizon and planning-cost share as checks on economic plausibility.

The second row of Table 1 shows the standard FHP model with fixed planning discussed extensively in [Woodford \(2019\)](#) and [Gust et al. \(2026\)](#), who estimate $k^* = 4$. While this variant is implemented by fixing $k^* = 4$, it can be interpreted as a limiting case of the linear cost function with no cost for horizons up to k^* and prohibitive costs beyond that horizon. The third row shows the NK Phillips curve under rational expectations, which is the costless-planning limit with $\tau_0 = 0$ and $k^* \rightarrow \infty$. Finally, for robustness, we consider an endogenous-planning variant with an exponential cost function:

$$\tau(k) = \frac{\mathcal{K}}{2} \left(\frac{b_0}{100} \right)^2 [\exp(b_1 k) - 1]. \quad (19)$$

Similar to the piecewise linear function, we choose the parameters b_0 and b_1 to match an average planning horizon of 4 quarters and the empirical relationship between inflation uncertainty and forecast revisions.


As shown in Table 1, the calibrated aggregate-supply shock process is quite similar in the baseline endogenous-planning model and the fixed-planning model. In both cases, the model generates persistent movements in aggregate inflation through firms’ longer-run beliefs about inflation ($\bar{\pi}_t$) while higher-frequency movements in inflation are captured by a less persistent aggregate-supply shock. By contrast, the rational expectations variant has little endogenous persistence through expectations formation. To generate persistent movements in inflation, the inferred aggregate-supply shock must therefore be much more persistent.

The top panel of Figure 2 varies the cost parameter, τ_0 , to show how it affects the average planning horizon in the model. Starting from the rational expectations limit, $\tau_0 = 0$, planning

is costless and firms choose a very long average planning horizon. As τ_0 rises, sophisticated forward-looking planning becomes more costly and the average horizon declines. At high values of the cost parameter, the horizon approaches the fixed-horizon limit associated with the free-planning threshold, which for the baseline calibration is 2 quarters. The baseline calibration of τ_0 lies between these two extremes and implies $E[k_t^*] \approx 4$, matching the average planning horizon estimated in [Gust et al. \(2026\)](#).

The bottom panel reports average planning costs paid by resetting firms as a share of their steady-state per-period profits. These costs are zero when $\tau_0 = 0$, since planning is costless, and initially rise as higher values of τ_0 make forward planning costly. The relationship is not monotone. As τ_0 rises further, firms economize on how far ahead they plan, and average costs start to fall. In the fixed-planning limit, average planning costs are again zero even though the marginal cost of planning beyond $k_t^* = \underline{k}$ is prohibitive. Quantitatively, average planning costs are very small, remaining below 0.05 percent of steady-state per-period profits. This reflects that most of the benefit from planning comes from setting the reset price “right on average”, as summarized by the experience-based continuation value, which is costless for a firm to update. Relative to the benefit of setting the reset price “right on average”, the incremental gain from setting prices to reflect expected future contingencies one additional period ahead is small.

Figure 2: Planning horizon and planning costs.



Note: The figure varies τ_0 while holding the other parameters and the aggregate-supply shock process at their baseline values. Average planning costs are costs paid by firms that receive a Calvo price-reset opportunity, expressed as a share of those firms' steady-state per-period profits.

4.2 Forecast Predictability


As discussed in [Gust, Herbst, and López-Salido \(2022\)](#) and [Gust, Herbst, and López-Salido \(2026\)](#), the fixed-planning model generates movements in longer-run inflation expectations that are in line with empirical measures and accounts for survey evidence that inflation forecast errors are predictable. We replicate this latter result using the predictability regression of [Coibion and Gorodnichenko \(2015\)](#) and show that the endogenous-planning model is also consistent with this evidence. In particular, Figure 3 shows the estimate from regressing the average SPF forecast error of inflation on the average SPF forecast revision.⁹ The coefficient on forecast revisions in the forecast-error regression, β_{CG} , from this data is 1.25. This positive relationship implies that forecasts underreact to new information: when forecasters revise up expected inflation, realized inflation tends to rise by more than the average SPF

⁹Appendix C describes the empirical CG regression and the construction of the model-based forecasts, forecast errors, and forecast revisions. For comparability with [Gust et al. \(2026\)](#), we use the 1969Q4-2007Q4 sample period for the empirical CG regression.

forecast.

The FHP model generates this positive correlation under either fixed or endogenous planning because firms’ forecasts are initially muted relative to realized inflation. This muted effect reflects that firms’ forecast do not fully reflect the effects of persistent shocks outside of their planning horizons. While the regression coefficients implied by the endogenous-planning variants of the model are somewhat low relative to the empirical point estimate, they are both well within the 95 percent confidence interval. In contrast, in the rational expectations version of the model, firms’ planning horizons are very long and fully reflect the effects of persistent shocks, making forecast errors unpredictable. As a result, under rational expectations, $\beta_{CG} = 0$. In sum, endogenous planning preserves the forecast-predictability evidence that supports the FHP paradigm.

Figure 3: Inflation forecast predictability across models.



Note: The data bar reports the point estimate from regressing the average SPF forecast error on the average SPF forecast revision over the 1969:Q4–2007:Q4 sample period. The point estimate is $\hat{\beta}_{CG} = 1.25$ with an approximate 95 percent confidence interval of $[0.28, 2.22]$. Model bars report simulated population moments using 200,000 observations.

4.3 Forecast Revisions and Inflation Uncertainty

A key prediction of the endogenous planning model is that uncertainty about inflation is positively related to the size of a firm’s forecast revisions. To assess this prediction, we use an empirical measure of inflation uncertainty estimated from the predictive density of De Polis et al. (2024) from the U.S. core PCE prices. De Polis et al. (2024) use a score-driven skew- t density model for inflation, building on the score-driven framework of Harvey (2013) and Creal et al. (2013) and the macroeconomic downside-risk specification in Delle Monache

et al. (2024). This is a density-based uncertainty measure rather than a cross-sectional measure of forecast disagreement.¹⁰ We align this measure with the absolute value of same-target SPF four-quarter forecast revisions, which captures the size of the aggregate surprise to forecasters' inflation outlooks. Appendix E summarizes the construction of the empirical uncertainty measure.

In the model, as in the data, the uncertainty measure is a one-quarter-ahead predictive standard deviation. The model counterpart is the one-quarter-ahead conditional volatility of inflation, defined as follows:

$$\sigma_{\pi,t} = \sqrt{\text{Var}_t(\pi_{t+1})}, \quad \text{Var}_t(\pi_{t+1}) \equiv E_t [(\pi_{t+1} - E_t \pi_{t+1})^2].$$


Under fixed planning, the horizon is constant and the conditional variance of inflation is constant. Under endogenous planning, large and persistent shocks have two effects. First, they shift firms' inflation outlooks, generating large absolute forecast revisions. Second, they raise the marginal benefit of planning further ahead. Firms then choose longer horizons, increasing the sensitivity of inflation to shocks and raising the one-quarter-ahead conditional volatility of inflation.¹¹ Accordingly, the endogenous-planning model predicts a positive correlation between absolute forecast revisions and the conditional volatility of inflation.

Figure 4 shows this relationship in the data and in the baseline endogenous-planning model. The left panel plots empirical inflation uncertainty against the absolute value of SPF forecast revisions, while the right panel plots the same objects in the model. In both panels, inflation uncertainty and forecast revisions are positively correlated. This positive relationship is the key moment we use to identify the marginal planning cost, τ_0 .

¹⁰The distinction between disagreement and uncertainty is emphasized by D'Amico and Orphanides (2008) and Giordani and Söderlind (2003). The broader empirical literature also documents substantial time variation in model-based and density-based uncertainty measures; see, for example, Stock and Watson (2007), Primiceri (2005), Justiniano and Primiceri (2008), and Jurado et al. (2015).

¹¹Appendix D describes how we compute the one-quarter-ahead conditional volatility of inflation in the endogenous-planning model.

Figure 4: Forecast revisions and inflation uncertainty in the data and model.



Note: The left panel uses the empirical one-quarter-ahead density-uncertainty series for core PCE inflation and the average SPF four-quarter forecast revision over the 1969:Q1–2007:Q1 survey-quarter sample. The right panel uses simulated data from the baseline endogenous-planning model. The red line in the data panel is an OLS fitted line; the red line in the model panel reports equal-count binned means. ρ denotes the Pearson correlation between the two series.

Table 2 compares the empirical moments with the baseline linear endogenous-planning model, the fixed-planning model, and the exponential-cost endogenous-planning model. In the data, larger forecast revisions in the SPF are associated with a higher level of inflation uncertainty, as the two series have a correlation a little higher than 0.4. The two endogenous-planning models are able to match this moment while the fixed-planning model cannot, since inflation uncertainty does not vary when the planning horizon is fixed.

Table 2 also reports the implied distribution of planning horizons and the size of model-implied stochastic volatility relative to the empirical density-uncertainty series. The baseline linear model captures a modest part of the total volatility in inflation uncertainty: the standard deviation of model-implied uncertainty is about 14 percent of the empirical value. This relatively low volatility reflects that the endogenous-planning model generates this volatility endogenously from standard aggregate demand and supply shocks. If we allowed the volatility of these shocks to vary stochastically as well, the model could match the fluctuations in the empirical measure of inflation uncertainty.

Table 2: Selected Moments Across Model Variants

Object	Data	Baseline Linear	Fixed	Exponential
<i>Targeted moments</i>				
$\text{Corr}(\sigma_{\pi,t}, REV_{t,t+4})$	0.43	0.44	n.a.	0.40
$E[k_t^*]$	4.00	4.09	4.00	4.00
<i>Inflation uncertainty</i>				
$E[\sigma_{\pi,t}]$	0.279	0.224	0.212	0.220
$\text{sd}(\sigma_{\pi,t})$	0.166	0.023	0.000	0.015
<i>Planning horizon distribution</i>				
$\text{mode}(k_t^*)$	n.a.	2	4	3
$\text{sd}(k_t^*)$	n.a.	3.13	0.00	2.52
$\text{Pr}(k_t^* > \underline{k})$	n.a.	0.57	0.00	n.a.

Notes: Data moments use the empirical one-quarter-ahead density-uncertainty series for core PCE inflation and the average SPF same-target four-quarter forecast revision over the 1969:Q1–2007:Q1 survey-quarter sample. Model entries report simulated population moments using 200,000 observations. The fixed-planning Pearson correlation is not reported because inflation uncertainty is constant in that model.


4.4 Historical Movements in Inflation Uncertainty

With the endogenous-planning model generating time-varying inflation uncertainty endogenously, Figure 5 compares the model-implied historical uncertainty series with the empirical one-quarter-ahead uncertainty measure. The model-based measure is smaller than the empirical measure, consistent with the moments in Table 2. Nevertheless, in several important episodes, uncertainty in the model rises in line with the empirical measure. In particular, the model-implied measure increases sharply during the high-inflation period of the 1970s and early 1980s, with smaller increases occurring around the Great Financial Crisis and the COVID/post-pandemic inflation episode. The model’s uncertainty measure rises during these episodes because the aggregate demand and supply shocks are relatively large relative to trend inflation at those times, inducing firms to plan further ahead, as shown in the middle panel. This deeper planning then raises the sensitivity of inflation to shocks, making it more volatile.

The bottom panel of the figure shows the endogenous-planning model’s implications for firms’ expectations of longer-run inflation ($\bar{\pi}_t$) during these episodes. Firms’ longer-run

inflation beliefs rose through the 1970s and early 1980s and then fell during the disinflation that followed. During the COVID/post-pandemic inflation episode, the model's measure of inflation uncertainty rises only temporarily, while firms' expectations of longer-run inflation, which depend on lagged realized inflation, remained elevated in 2023 and 2024.

Figure 5: Empirical and model-based measures of inflation uncertainty.



Note: The middle panel reports the selected planning horizon k_t^* , with the dashed line showing its sample mean. The trend-inflation line is $\pi^{Avg} + \bar{\pi}_t$, where $\bar{\pi}_t$ is the demeaned model trend state and π^{Avg} is the sample mean of quarterly core PCE inflation.

5 Conclusion

This paper extends the finite-horizon planning framework by making firms' planning horizons endogenous. The endogenous-planning model preserves the main empirical strengths of fixed finite-horizon planning models while generating time-varying inflation uncertainty. In the model, forecast revisions and inflation uncertainty are positively related, as in the data, and this relationship provides a way to discipline the marginal cost associated with forward-looking planning.

A natural next step is to embed endogenous planning horizons in a richer general equilibrium environment. In such a model, both households and firms could choose how far ahead to plan, allowing the planning horizon to shape not only price setting but also consumption, saving, and expectations about future policy. This extension would provide a useful framework for studying the monetary-policy implications of endogenous planning, including how alternative monetary policies and forward guidance affect the planning horizons of households and firms.

References

- AKERLOF, G. A. AND J. L. YELLEN (1985): “A Near-Rational Model of the Business Cycle, with Wage and Price Inertia,” *The Quarterly Journal of Economics*, 100, 823–838.
- ALVAREZ, F. E., F. LIPPI, AND L. PACIELLO (2011): “Optimal Price Setting With Observation and Menu Costs,” *The Quarterly Journal of Economics*, 126, 1909–1960.
- ANGELETOS, G.-M., Z. HUO, AND K. A. SASTRY (2020): “Imperfect Macroeconomic Expectations: Evidence and Theory,” Working Paper 27308, National Bureau of Economic Research.
- ATKESON, A. AND A. BURSTEIN (2008): “Pricing-to-Market, Trade Costs, and International Relative Prices,” *American Economic Review*, 98, 1998–2031.
- BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77, 623–685.
- BOUTROS, M. (2026): “Windfall income shocks with finite planning horizons,” *Journal of Financial Economics*, 176, 104174.
- CALVO, G. (1983): “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 383–398.
- CHOW, Y. S., H. ROBBINS, AND D. SIEGMUND (1971): *Great Expectations: The Theory of Optimal Stopping*, Boston: Houghton Mifflin.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.
- COIBION, O. AND Y. GORODNICHENKO (2015): “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” *American Economic Review*, 105, 2644–78.
- CREAL, D., S. J. KOOPMAN, AND A. LUCAS (2013): “Generalized Autoregressive Score Models with Applications,” *Journal of Applied Econometrics*, 28, 777–795.
- D’AMICO, S. AND A. ORPHANIDES (2008): “Uncertainty and Disagreement in Economic Forecasting,” Tech. Rep. 2008-56, Board of Governors of the Federal Reserve System.
- DE POLIS, A., L. MELOSI, AND I. PETRELLA (2024): “The Taming of the Skew: Asymmetric Inflation Risk and Monetary Policy,” December 2024 draft.
- DELLE MONACHE, D., A. DE POLIS, AND I. PETRELLA (2024): “Modeling and Fore-

- casting Macroeconomic Downside Risk,” *Journal of Business & Economic Statistics*, 42, 1010–1025.
- FERGUSON, T. S. (2008): *Optimal Stopping and Applications*, UCLA, available at <https://www.math.ucla.edu/~tom/Stopping/Contents.html>.
- GABAIX, X. (2020): “A Behavioral New Keynesian Model,” *American Economic Review*, 110, 2271–2327.
- GIORDANI, P. AND P. SÖDERLIND (2003): “Inflation Forecast Uncertainty,” *European Economic Review*, 47, 1037–1059.
- GUST, C., E. HERBST, AND D. LÓPEZ-SALIDO (2026): “Inflation Expectations with Finite Horizon Planning,” *Journal of Economic Dynamics and Control*, 182, 105222.
- GUST, C., E. HERBST, AND D. LÓPEZ-SALIDO (2022): “Short-Term Planning, Monetary Policy, and Macroeconomic Persistence,” *American Economic Journal: Macroeconomics*, 14, 174–209.
- HARVEY, A. C. (2013): *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*, Econometric Society Monographs, Cambridge: Cambridge University Press.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105, 1177–1216.
- JUSTINIANO, A. AND G. E. PRIMICERI (2008): “The Time-Varying Volatility of Macroeconomic Fluctuations,” *American Economic Review*, 98, 604–641.
- MAĆKOWIAK, B. AND M. WIEDERHOLT (2009): “Optimal Sticky Prices under Rational Inattention,” *American Economic Review*, 99, 769–803.
- MANKIW, N. G. (1985): “Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly,” *The Quarterly Journal of Economics*, 100, 529–537.
- MANKIW, N. G. AND R. REIS (2002): “Sticky information versus sticky prices: A proposal to replace the new Keynesian Phillips curve,” *Quarterly Journal of Economics*, 117, 1295–1328.
- MCCALL, J. J. (1970): “Economics of Information and Job Search,” *The Quarterly Journal of Economics*, 84, 113–126.
- PACIELLO, L. AND M. WIEDERHOLT (2014): “Exogenous Information, Endogenous Information, and Optimal Monetary Policy,” *Review of Economic Studies*, 81, 356–388.

- PRIMICERI, G. E. (2005): “Time Varying Structural Vector Autoregressions and Monetary Policy,” *Review of Economic Studies*, 72, 821–852.
- REIS, R. (2006): “Inattentive Producers,” *Review of Economic Studies*, 73, 793–821.
- SIMS, C. A. (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50, 665–690.
- STOCK, J. H. AND M. W. WATSON (2007): “Why Has U.S. Inflation Become Harder to Forecast?” *Journal of Money, Credit and Banking*, 39, 3–33.
- STOKEY, N. L. (2009): *The Economics of Inaction: Stochastic Control Models with Fixed Costs*, Princeton, NJ: Princeton University Press.
- VAVRA, J. (2014): “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” *The Quarterly Journal of Economics*, 129, 215–258.
- WOODFORD, M. (2009): “Information-Constrained State-Dependent Pricing,” *Journal of Monetary Economics*, 56, S100–S124.
- (2019): “Monetary Policy Analysis When Planning Horizons Are Finite,” *NBER Macroeconomics Annual*, 33, 1–50.
- ZBARACKI, M. J., M. RITSON, D. LEVY, S. DUTTA, AND M. BERGEN (2004): “Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets,” *Review of Economics and Statistics*, 86, 514–533.

Appendix

Appendix A derives the equations that characterize inflation dynamics and a firm's planning-horizon rule. Appendix B proves the propositions in the main text, Appendix C describes the empirical CG regression and the forecast objects used to compute the model-implied CG regressions, Appendix D describes the model-based measure of inflation uncertainty, and Appendix E describes the empirical measure.

A Derivation of Inflation Dynamics

In this appendix, we derive the equations that characterize inflation in the model using a second order approximation of a firm's objective function. We show that this approximation allows us to characterize inflation dynamics and a firm's planning-horizon rule analytically.

A.1 Second Order Approximation of Per-Period Profits

We begin by specifying the exogenous processes for aggregate demand and the elasticity of substitution and then taking a second-order approximation of per-period profits. Let $r_{ft} \equiv \frac{P_{ft}}{P_t}$ denote firm f 's relative price. Per-period profits are given by

$$D(r_{ft}, Y_t, \varepsilon_t) = \left(r_{ft} - \bar{w} Y_{ft}^a \right) Y_{ft}, \quad (\text{A.1})$$

where

$$a \equiv \frac{\alpha}{1 - \alpha}, \quad \bar{w} = \frac{1 - \alpha}{\mu}, \quad \mu \equiv \frac{\bar{\varepsilon}}{\bar{\varepsilon} - 1}.$$

Firm-level demand depends on a firm's relative price as well as Y_t and ε_t :

$$Y_{ft} = r_{ft}^{-\varepsilon_t} Y_t. \quad (\text{A.2})$$

Evaluated at the steady state, per-period profits are

$$\bar{D} \equiv D(1, 1, \bar{\varepsilon}) = 1 - \frac{1 - \alpha}{\mu} = \frac{1 + \alpha(\bar{\varepsilon} - 1)}{\bar{\varepsilon}} > 0.$$

Define $\hat{y}_t \equiv \log(Y_t)$ and $\hat{r}_{ft} \equiv \log r_{ft}$ so that $r_{ft} = e^{\hat{r}_{ft}}$ along with $s_t = (\hat{y}_t, \hat{\varepsilon}_t)$. With these definitions, let

$$\hat{D}(\hat{r}, \hat{y}, \hat{\varepsilon}) \equiv D(e^{\hat{r}}, e^{\hat{y}}, \bar{\varepsilon} e^{\hat{\varepsilon}}),$$

so that $\hat{D}(\hat{r}_{ft}, \hat{y}_t, \hat{\varepsilon}_t) = D(r_{ft}, Y_t, \varepsilon_t)$. A second-order expansion of $\hat{D}$ around $(\hat{r}, \hat{y}, \hat{\varepsilon}) =$

$(0, 0, 0)$ yields

$$\begin{aligned} \hat{D}(\hat{r}_{ft}, s_t) - \bar{D} \approx \tilde{D}(\hat{r}_{ft}, s_t) &= \frac{1}{2} \hat{D}_{\hat{r}\hat{r}} \hat{r}_{ft}^2 + \hat{D}_{\hat{r}\hat{y}} \hat{r}_{ft} \hat{y}_t + \hat{D}_{\hat{r}\hat{\varepsilon}} \hat{r}_{ft} \hat{\varepsilon}_t \\ &+ \hat{D}_{\hat{y}} \hat{y}_t + \frac{1}{2} \hat{D}_{\hat{y}\hat{y}} \hat{y}_t^2, \end{aligned} \quad (\text{A.3})$$

where derivatives of $\hat{D}$ are evaluated at $(0, 0, 0)$.

The coefficients in (A.3) are:

$$\begin{aligned} \hat{D}_{\hat{r}\hat{r}} &= -\left[(a+1)\bar{\varepsilon} - 1\right](\bar{\varepsilon} - 1) < 0, \\ \hat{D}_{\hat{r}\hat{y}} &= a(\bar{\varepsilon} - 1) > 0, \quad \hat{D}_{\hat{r}\hat{\varepsilon}} = -1, \\ \hat{D}_{\hat{y}} &= \frac{1}{\bar{\varepsilon}} > 0, \quad \hat{D}_{\hat{y}\hat{y}} = \frac{1 - a(\bar{\varepsilon} - 1)}{\bar{\varepsilon}}. \end{aligned}$$

A.2 Continuation Value Function

We write the continuation value function as a quadratic in the firm's relative price parameterized by belief coefficients (c_t, v_t, χ_t) :

$$H_t(\hat{r}_{ft}) = \mathcal{K} \left[v_t \hat{r}_{ft} - \frac{1}{2} \chi_t \hat{r}_{ft}^2 + c_t \right], \quad (\text{A.4})$$

where the scaling parameter, $\mathcal{K}$, is given by:

$$\mathcal{K} = \frac{-\hat{D}_{\hat{r}\hat{r}}}{1 - \beta\theta} = \left[(a+1)\bar{\varepsilon} - 1\right] \frac{\bar{\varepsilon} - 1}{1 - \beta\theta} > 0.$$

The subscript t indicates dependence on the date- t belief parameters (c_t, v_t, χ_t) , which are taken as fixed when solving a firm's date- t problem.

A.3 Pricing Decision

We use the second-order approximation to characterize a firm's pricing decision and planning-horizon decision. We begin with the pricing choice for $k = 0$ and then turn to a firm's pricing decision for $k > 0$. In the firm's problem, we scale a firm's nominal price by lagged inflation so that $p_{ft} = \frac{P_{ft}}{\bar{\pi} P_{t-1}}$ and $\hat{p}_{ft} = \hat{r}_{ft} + \hat{\pi}_t$.

A firm with planning horizon $k = 0$ solves

$$\max_{\hat{p}_{ft}} H_t^0(\hat{p}_{ft}, s_t),$$

subject to

$$H_t^0(\hat{p}_{ft}, s_t) = \tilde{D}(\hat{p}_{ft} - \hat{\pi}_t^0, s_t) + \beta\theta H_t(\hat{p}_{ft} - \hat{\pi}_t^0). \quad (\text{A.5})$$

In equation (A.5), $\hat{\pi}_t^0$ denotes the beliefs about aggregate inflation of a firm with a $k = 0$ horizon and the dependence of $H_t^0(\hat{p}_{ft}, s_t)$ on t reflects its dependence on the parameters governing the continuation value function, (c_t, v_t, χ_t) .

Differentiating a firm's value function equations with respect to $\hat{p}_{ft}$ yields

$$\begin{aligned} H_{\hat{p},t}^0(\hat{p}_{ft}, s_t) \equiv \mathcal{K}[v_t(0) - \chi_t(0)\hat{p}_{ft}] &= (1 - \beta\theta)\mathcal{K}[\xi\hat{y}_t - \zeta\hat{\varepsilon}_t - (\hat{p}_{ft} - \hat{\pi}_t^0)] + \\ &\quad \beta\theta\mathcal{K}[v_t - \chi_t(\hat{p}_{ft} - \hat{\pi}_t^0)], \end{aligned} \quad (\text{A.6})$$

where

$$\xi \equiv \frac{a}{(a+1)\bar{\varepsilon} - 1} > 0, \quad \zeta \equiv \frac{1}{[(a+1)\bar{\varepsilon} - 1](\bar{\varepsilon} - 1)} > 0,$$

Also, we have defined a firm with horizon $k = 0$ marginal value function, $H_{\hat{p},t}^0(\hat{p}_{ft}, s_t)$, as a function of the parameters, $(\chi_t(0), v_t(0))$. Equating terms in expression (A.6), these coefficients are given by:

$$\chi_t(0) = (1 - \beta\theta) + \beta\theta\chi_t = 1 - \beta\theta(1 - \chi_t), \quad v_t(0) = (1 - \beta\theta)(\xi\hat{y}_t - \zeta\hat{\varepsilon}_t) + \beta\theta v_t + \hat{\pi}_t^0.$$

Note that $\chi_t = 1$ implies $\chi_t(0) = 1$. Later we show that a firm's updating of its continuation value function implies that $\chi_t = 1$ is a fixed point in which $\chi_t = 1$ for all $t \geq 0$. Imposing this fixed point in the first-order condition for $k = 0$ horizon and solving for a firm's reset price implies

$$\hat{p}_t^0 = v_t(0) = (1 - \beta\theta)(\xi\hat{y}_t - \zeta\hat{\varepsilon}_t) + \beta\theta v_t + \hat{\pi}_t^0, \quad (\text{A.7})$$

where $\hat{p}_t^0$ denotes the common reset price of a firm with $k = 0$ horizon.

As discussed in the main text, a firm forms beliefs under the assumption that other firms have the same planning horizon as themselves. Applying this assumption in expression (6) of the main text implies that the beliefs for aggregate inflation of a k -horizon firm satisfy:

$$\left(\frac{\pi_t^k}{\bar{\pi}}\right)^{1-\varepsilon_t} = \theta + (1 - \theta)(p_t^k)^{1-\varepsilon_t}.$$

Log-linearizing this expression yields:

$$\hat{\pi}_t^k = (1 - \theta)\hat{p}_t^k, \quad k \geq 0. \quad (\text{A.8})$$

Substituting this expression into equation (A.7) implies that the beliefs for a firm with horizon $k = 0$ satisfy:

$$\hat{\pi}_t^0 \equiv \Pi^0(v_t, s_t) = \beta(1 - \theta)v_t + \kappa\hat{y}_t + \hat{\varepsilon}_{st}, \quad (\text{A.9})$$

where

$$\kappa \equiv \tilde{\kappa}\xi, \quad \tilde{\kappa} \equiv \frac{(1 - \beta\theta)(1 - \theta)}{\theta},$$

and the aggregate supply disturbance is normalized as $\hat{\varepsilon}_{st} \equiv -\tilde{\kappa}\zeta\hat{\varepsilon}_t$. With this normalization, expression (3) can be rewritten as:

$$\hat{\varepsilon}_{st} = \rho_\varepsilon \hat{\varepsilon}_{st-1} + \sigma_s e_{st}, \quad e_{st} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \quad (\text{A.10})$$

where $\sigma_s \equiv \tilde{\kappa}\zeta\sigma_\varepsilon$.

Expression (A.9) highlights that a firm with a $k = 0$ horizon believes that inflation at date t depends on its current estimate of its marginal continuation value function (v_t) and the shocks at date t . We can also use this expression to show how $H_{\hat{p},t}^0(\hat{p}_{ft}, s_t)$ depends on v_t and the date t shocks:

$$H_{\hat{p},t}^0(\hat{p}_{ft}, s_t) = \mathcal{K}[v_t(0) - \hat{p}_{ft}] = \mathcal{K}[\hat{p}_t^0 - \hat{p}_{ft}] = \mathcal{K}\left[\frac{1}{1 - \theta}\Pi^0(v_t, s_t) - \hat{p}_{ft}\right]. \quad (\text{A.11})$$

Pricing for $k > 0$. For planning horizons $k \geq 1$, a firm chooses its reset price according to

$$\max_{\hat{p}_{ft}} H_t^k(\hat{p}_{ft}, s_t),$$

$$H_t^k(\hat{p}_{ft}, s_t) = \tilde{D}(\hat{p}_{ft} - \hat{\pi}_t^k, s_t) + \beta\theta E_t H_t^{k-1}(\hat{p}_{ft}, s_{t+1}). \quad (\text{A.12})$$

Differentiating (A.12) with respect to $\hat{p}_{ft}$ yields

$$H_{\hat{p},t}^k(\hat{p}_{ft}, s_t) = (1 - \beta\theta)\mathcal{K}\left[\xi\hat{y}_t - \zeta\hat{\varepsilon}_t - (\hat{p}_{ft} - \hat{\pi}_t^k)\right] + \beta\theta E_t H_{\hat{p},t}^{k-1}(\hat{p}_{ft}, s_{t+1}), \quad (\text{A.13})$$

For $k = 1$, we can use (A.11) to write

$$E_t H_{\hat{p},t}^0(\hat{p}_{ft}, s_{t+1}) = \mathcal{K}\left[\frac{1}{1 - \theta}E_t \hat{\pi}_{t+1}^0 - \hat{p}_{ft}\right], \quad (\text{A.14})$$

where

$$E_t \hat{\pi}_{t+1}^0 = E_t \Pi^0(v_t, s_{t+1}) = \beta(1 - \theta)v_t + \kappa \rho_y \hat{y}_t + \rho_\varepsilon \hat{\varepsilon}_{st}.$$

Using expression (A.14) in equation (A.13) and solving for a firm's reset price yields:

$$\hat{p}_t^1 - \hat{\pi}_t^1 = (1 - \beta\theta)(\xi \hat{y}_t - \zeta \hat{\varepsilon}_t) + \frac{\beta\theta}{1 - \theta} E_t \hat{\pi}_{t+1}^0.$$

Using equation (A.8), we can write a firm's beliefs for inflation as:

$$\hat{\pi}_t^1 = \beta E_t \hat{\pi}_{t+1}^0 + \kappa \hat{y}_t + \hat{\varepsilon}_{st}$$

We can also define $H_{\hat{p},t}^1(\hat{p}_{ft}, s_t)$ as:

$$H_{\hat{p},t}^1(\hat{r}_{ft}, s_t) \equiv \mathcal{K} [v_t(1) - \chi_t(1)\hat{p}_{ft}], \quad (\text{A.15})$$

and note that the coefficients $v_t(1)$ and $\chi_t(1)$ satisfy:

$$v_t(1) = \hat{p}_t^1, \quad \chi_t(1) = 1,$$

where we have $\chi_t(1) = 1$, because $\chi_t = 1$ implies $\chi_t(0) = \chi_t(1) = 1$. Thus, we can write:

$$H_{\hat{p},t}^1(\hat{p}_{ft}, s_t) = \mathcal{K} [\hat{p}_t^1 - \hat{p}_{ft}] = \mathcal{K} \left[\frac{1}{1 - \theta} \hat{\pi}_t^1 - \hat{p}_{ft} \right] = \mathcal{K} \left[\frac{1}{1 - \theta} \Pi^1(v_t, s_t) - \hat{p}_{ft} \right]$$

Solving the firm's problem sequentially for $k > 1$, one can show that $\chi_t(k)$ coefficients satisfies the recursion:

$$\chi_t(0) = (1 - \beta\theta) + \beta\theta \chi_t, \quad \chi_t(k) = (1 - \beta\theta) + \beta\theta \chi_t(k - 1), \quad k \geq 1. \quad (\text{A.16})$$

It follows from this recursion that $\chi_t = 1$ implies $\chi_t(k) = 1$ for all $k \geq 0$ at the fixed point $\chi = 1$. Using this recursion, we can write

$$H_{\hat{p},t}^{k-1}(\hat{p}_{ft}, s_t) = \mathcal{K} [v_t(k - 1) - \hat{p}_{ft}] = \mathcal{K} [\hat{p}_t^{k-1} - \hat{p}_{ft}] = \mathcal{K} \left[\frac{1}{1 - \theta} \hat{\pi}_t^{k-1} - \hat{p}_{ft} \right]. \quad (\text{A.17})$$

Using equation (A.17) in expression (A.13) and using it to solve for a firm's reset price yields:

$$\hat{p}_t^k - \hat{\pi}_t^k = (1 - \beta\theta) (\xi \hat{y}_t - \zeta \hat{\varepsilon}_t) + \frac{\beta\theta}{1 - \theta} E_t \hat{\pi}_{t+1}^{k-1}. \quad (\text{A.18})$$

Using (A.8), we can write a k -horizon firm's inflation beliefs as:

$$\hat{\pi}_t^k = \beta E_t \hat{\pi}_{t+1}^{k-1} + \kappa \hat{y}_t + \hat{\varepsilon}_{st}, \quad k \geq 1. \quad (\text{A.19})$$

Iterating (A.19) forward therefore implies

$$\hat{\pi}_t^k = A_y(k) \kappa \hat{y}_t + A_\varepsilon(k) \hat{\varepsilon}_{st} + \beta^{k+1} (1 - \theta) v_t, \quad (\text{A.20})$$

where

$$A_y(k) = \frac{1 - (\beta \rho_y)^{k+1}}{1 - \beta \rho_y}, \quad A_\varepsilon(k) = \frac{1 - (\beta \rho_\varepsilon)^{k+1}}{1 - \beta \rho_\varepsilon}.$$

Dividing by $(1 - \theta)$ gives a firm's reset price:

$$\hat{p}_t^k = \frac{\kappa A_y(k)}{1 - \theta} \hat{y}_t + \frac{A_\varepsilon(k)}{1 - \theta} \hat{\varepsilon}_{st} + \beta^{k+1} v_t. \quad (\text{A.21})$$

A.4 Trend Inflation Derivation

The pricing derivation also implies the law of motion for the trend component in the inflation equation. Recall from (11) that a firm updates its continuation value function according to

$$H_{t+1}(\hat{p}_{ft}) = (1 - \gamma) H_t(\hat{p}_{ft}) + \gamma H_t^E(\hat{p}_{ft}), \quad (\text{A.22})$$

where a firm's new estimate of its value function is the value function associated with the realized planning horizon (k_t^*) and realized value of S_t at date t :

$$H_t^E(\hat{p}_{ft}) = H_t^{k_t^*}(\hat{p}_{ft}, \underline{S}_t).$$

Differentiating (A.22) with respect to $\hat{p}_{ft}$ gives

$$H_{\hat{p},t+1}(\hat{p}_{ft}) = (1 - \gamma) H_{\hat{p},t}(\hat{p}_{ft}) + \gamma H_{\hat{p},t}^E(\hat{p}_{ft}). \quad (\text{A.23})$$

Under the fixed point $\chi_t = 1$, the marginal continuation value function is

$$H_{\hat{p},t}(\hat{p}_{ft}) = \mathcal{K} [v_t - \hat{p}_{ft}].$$

Using (A.17) at the selected horizon, the marginal estimate of the value function satisfies

$$H_{\hat{p},t}^E(\hat{p}_{ft}) = H_{\hat{p},t}^{k_t^*}(\hat{p}_{ft}, \underline{S}_t) = \mathcal{K}[v_t(k_t^*) - \hat{p}_{ft}] = \mathcal{K}\left[\frac{\hat{\pi}_t^{k_t^*}}{1-\theta} - \hat{p}_{ft}\right]. \quad (\text{A.24})$$

The last equality uses $v_t(k) = \hat{p}_t^k$ together with the log-linear price index relation $\hat{\pi}_t^k = (1-\theta)\hat{p}_t^k$.

Substituting these two marginal value functions into (A.23) gives

$$H_{\hat{p},t+1}(\hat{p}_{ft}) = \mathcal{K}\left[(1-\gamma)v_t + \gamma\frac{\hat{\pi}_t^{k_t^*}}{1-\theta} - \hat{p}_{ft}\right].$$

Comparing this expression with $H_{\hat{p},t+1}(\hat{p}_{ft}) = \mathcal{K}[v_{t+1} - \hat{p}_{ft}]$ implies

$$v_{t+1} = (1-\gamma)v_t + \gamma\frac{\hat{\pi}_t^{k_t^*}}{1-\theta}. \quad (\text{A.25})$$

Finally, define trend inflation as the scaled marginal continuation value,

$$\bar{\pi}_t \equiv (1-\theta)v_t. \quad (\text{A.26})$$

Multiplying (A.25) by $(1-\theta)$ yields

$$\bar{\pi}_{t+1} = (1-\gamma)\bar{\pi}_t + \gamma\hat{\pi}_t^{k_t^*}.$$

Since realized inflation is $\hat{\pi}_t \equiv \hat{\pi}_t^{k_t^*}$, trend inflation evolves according to

$$\bar{\pi}_{t+1} = (1-\gamma)\bar{\pi}_t + \gamma\hat{\pi}_t. \quad (\text{A.27})$$

A.5 Planning-Horizon Decision

Planning is costless at $k = 0$ but costly to extend further into the future. Let $\tau(k)$ denote the total cost of planning k periods ahead, with marginal cost $\Delta\tau_{k+1} \equiv \tau(k+1) - \tau(k) > 0$. To determine k_t^* , a firm starts at $k = 0$ and proceeds sequentially through $k = 0, 1, 2, \dots$, extending its planning horizon from k to $k+1$ whenever the marginal benefit exceeds the marginal cost:

$$MB(k+1)(v_t, s_t) \geq \Delta\tau_{k+1}, \quad k = 0, 1, 2, \dots, \quad (\text{A.28})$$

and stops otherwise.

The marginal benefit of extending the horizon from k to $k + 1$ is the increase in the horizon- $(k + 1)$ value from re-optimizing its price taking into account contingencies through period $t + k + 1$:

$$MB(k + 1)(v_t, s_t) = H_t^{k+1}(\hat{p}_t^{k+1}, s_t) - H_t^{k+1}(\hat{p}_t^k, s_t), \quad (\text{A.29})$$

Expression (A.17) implies that $H_t^{k+1}(\hat{p}_t^k, s_t)$ satisfies:

$$H_t^{k+1}(\hat{p}_{ft}, s_t) = \mathcal{K} \left[\hat{p}_t^{k+1} \hat{p}_{ft} - \frac{1}{2} p_{ft}^2 + c_t(k + 1) \right]$$

Because (A.29) compares the same horizon- $(k + 1)$ objective evaluated at two different prices while holding $(k + 1, s_t, v_t)$ fixed, the term $c_t(k + 1)$ drops out of the marginal-benefit calculation and we can write:

$$MB(k + 1)(v_t, s_t) = \mathcal{K} \left[\hat{p}_t^{k+1} (\hat{p}_t^{k+1} - \hat{p}_t^k) - \frac{1}{2} (\hat{p}_t^{k+1})^2 + \frac{1}{2} (\hat{p}_t^k)^2 \right]$$

which simplifies to:

$$MB(k + 1)(v_t, s_t) = \frac{\mathcal{K}}{2} (\hat{p}_t^{k+1} - \hat{p}_t^k)^2 \geq 0. \quad (\text{A.30})$$

Using (A.20), the one-stage reset-price increment can be written as

$$\hat{p}_t^{k+1} - \hat{p}_t^k = \frac{\beta^{k+1}}{1 - \theta} [\kappa \rho_y^{k+1} \hat{y}_t + \rho_\varepsilon^{k+1} \hat{\varepsilon}_{st} - (1 - \beta) \bar{\pi}_t], \quad (\text{A.31})$$

where we have defined trend inflation as $\bar{\pi}_t \equiv (1 - \theta)v_t$. Substituting (A.31) into (A.30) yields the stopping rule

$$k_t^* = \min \left\{ k \in \{0, 1, 2, \dots\} \mid \frac{\mathcal{K}}{2} (\hat{p}_t^{k+1} - \hat{p}_t^k)^2 \leq \Delta \tau_{k+1} \right\}. \quad (\text{A.32})$$

Thus, the planning horizon is determined by the point at which the squared one-stage change in the firm's desired reset price becomes too small to justify the additional planning cost.

B Proofs of Propositions

This section proves the two propositions in the main text. The proofs use the stopping rule (A.32) and the one-stage reset-price increment (A.31) derived in Appendix A. We first state a simple monotonicity result under which the one-stage look-ahead rule solves the full sequential planning problem.

Lemma B.1 (Monotone one-stage look-ahead rule). *For a fixed state, let*

$$D_t(k+1) \equiv MB_t(k+1) - \Delta\tau_{k+1}$$

denote the net marginal value of extending the planning horizon from k to $k+1$. Suppose $D_t(k+1)$ is weakly decreasing in k and that a finite stopping horizon exists. Then the one-stage look-ahead rule

$$k_t^* = \min\{k \geq 0 : D_t(k+1) \leq 0\}$$

solves the sequential planning problem generated by the one-stage net marginal values. Under the convention that the firm selects the shortest horizon among ties, k_t^ is the unique selected global solution.*

Proof. For every $j < k_t^*$, $D_t(j+1) > 0$, so each extension before k_t^* raises the cumulative net value of planning. Because $D_t(k+1)$ is weakly decreasing, $D_t(j+1) \leq 0$ for every $j \geq k_t^*$. Therefore any extension beyond k_t^* adds only weakly negative net marginal value. Thus no horizon below k_t^* or above k_t^* yields a higher cumulative net value, so k_t^* solves the sequential planning problem. If a longer horizon ties with k_t^* , every intervening net marginal value must be zero; the shortest-horizon tie-breaking convention selects k_t^* uniquely. $\square$

B.1 Proof of Proposition 3.1

Assume $0 < \rho_y < 1$, $\bar{\pi}_t = 0$, and $\hat{\varepsilon}_{st} = 0$. Then (A.31) implies

$$\hat{p}_t^{k+1} - \hat{p}_t^k = \frac{\beta^{k+1}}{1-\theta} \kappa \rho_y^{k+1} \hat{y}_t.$$

Under the linear planning cost $\tau(k) = (\mathcal{K}/2)\tau_0^2 k$, the stopping condition is

$$\frac{\mathcal{K}}{2} \left(\frac{(\beta\rho_y)^{k+1}}{1-\theta} \kappa \hat{y}_t \right)^2 \leq \frac{\mathcal{K}}{2} \tau_0^2.$$

Because $\mathcal{K} > 0$, this is equivalent to

$$|\kappa \hat{y}_t| \leq \frac{(1-\theta)\tau_0}{(\beta\rho_y)^{k+1}} = \frac{\Theta}{(\beta\rho_y)^{k+1}}, \quad \Theta \equiv (1-\theta)\tau_0.$$

Immediate stopping occurs at $k = 0$ iff $|\kappa \hat{y}_t| \leq \Theta/(\beta\rho_y)$. For $k \geq 1$, the firm chooses horizon k iff it does not stop at $k-1$ but does stop at k . Since $0 < \beta\rho_y < 1$, these two conditions

are

$$|\kappa \hat{y}_t| > \frac{\Theta}{(\beta \rho_y)^k} \quad \text{and} \quad |\kappa \hat{y}_t| \leq \frac{\Theta}{(\beta \rho_y)^{k+1}},$$

which implies

$$k_t^* = k \iff \frac{\Theta}{(\beta \rho_y)^k} < |\kappa \hat{y}_t| \leq \frac{\Theta}{(\beta \rho_y)^{k+1}}.$$

For each candidate horizon k , the marginal benefit is

$$MB_t(k+1) = \frac{\mathcal{K}}{2} \left(\frac{\beta^{k+1}}{1-\theta} \kappa \rho_y^{k+1} \right)^2 \hat{y}_t^2.$$

Because $0 < \beta \rho_y < 1$, $MB_t(k+1)$ is weakly decreasing in k . The marginal cost is constant, so the net marginal value $D_t(k+1) = MB_t(k+1) - \Delta \tau_{k+1}$ is weakly decreasing in k . Lemma B.1 then implies that the one-stage look-ahead rule selects the unique global solution under the tie-breaking convention. Away from cutoff states, no net marginal value is exactly zero, so the global solution is unique without invoking tie-breaking.

The marginal benefit is therefore a function of the state only through $|\hat{y}_t|$. This establishes symmetry: $k^*(\hat{y}_t, 0, 0) = k^*(-\hat{y}_t, 0, 0)$. It also establishes monotonicity. If $|\hat{y}_1| \leq |\hat{y}_2|$, then $MB_t(k+1; \hat{y}_1) \leq MB_t(k+1; \hat{y}_2)$ for every k . Therefore any stopping condition that is satisfied at $\hat{y}_2$ is also satisfied at $\hat{y}_1$, so the first stopping horizon cannot be larger at $\hat{y}_1$ than at $\hat{y}_2$. Thus $k^*(\hat{y}_t, 0, 0)$ is weakly increasing in $|\hat{y}_t|$.

B.2 Proof of Proposition 3.2

Assume $\rho_y = 1$ and $\hat{\varepsilon}_{st} = 0$. Then (A.31) becomes

$$\hat{p}_t^{k+1} - \hat{p}_t^k = \frac{\beta^{k+1}}{1-\theta} [\kappa \hat{y}_t - (1-\beta)\bar{\pi}_t] = \frac{\beta^{k+1}}{1-\theta} x_t,$$

where $x_t \equiv \kappa \hat{y}_t - (1-\beta)\bar{\pi}_t$. Under the linear planning cost $\tau(k) = (\mathcal{K}/2)\tau_0^2 k$, the marginal cost is constant:

$$\Delta \tau_{k+1} = \frac{\mathcal{K}}{2} \tau_0^2.$$

Substituting into the stopping rule (A.32) gives

$$\frac{\mathcal{K}}{2} \left(\frac{\beta^{k+1}}{1-\theta} x_t \right)^2 \leq \frac{\mathcal{K}}{2} \tau_0^2.$$

Because $\mathcal{K} > 0$, this is equivalent to

$$|x_t| \leq \frac{(1-\theta)\tau_0}{\beta^{k+1}} = \frac{\Theta}{\beta^{k+1}}, \quad \Theta \equiv (1-\theta)\tau_0.$$

Immediate stopping occurs at $k = 0$ iff $|x_t| \leq \Theta/\beta$. For $k \geq 1$, the firm chooses horizon k iff it does not stop at $k - 1$ but does stop at k . Since $0 < \beta < 1$, these two conditions are

$$|x_t| > \frac{\Theta}{\beta^k} \quad \text{and} \quad |x_t| \leq \frac{\Theta}{\beta^{k+1}},$$

which implies

$$k_t^* = k \iff \frac{\Theta}{\beta^k} < |x_t| \leq \frac{\Theta}{\beta^{k+1}}.$$

It remains to show that the one-stage look-ahead rule solves the sequential planning problem in this case. The marginal benefit is

$$MB_t(k+1) = \frac{\mathcal{K}}{2} \left(\frac{\beta^{k+1}}{1-\theta} x_t \right)^2.$$

Because $0 < \beta < 1$, $MB_t(k+1)$ is weakly decreasing in k . The marginal cost $\Delta\tau_{k+1} = (\mathcal{K}/2)\tau_0^2$ is constant, so the net marginal value $D_t(k+1) = MB_t(k+1) - \Delta\tau_{k+1}$ is weakly decreasing in k . Lemma B.1 then implies that the one-stage look-ahead rule selects the unique global solution under the tie-breaking convention. Away from cutoff states, no net marginal value is exactly zero, so the global solution is unique without invoking tie-breaking.

C Forecast Construction and CG Regressions

This section describes the empirical CG regression used in Figure 3 and how the corresponding model-implied forecasts, forecast errors, and forecast revisions used in the quantitative results are constructed. The notation for the model objects follows the inflation equation in the main text and the derivations in Appendix A.

C.1 Empirical CG Regression

The empirical coefficient in Figure 3 is constructed from Philadelphia Fed SPF GDP price deflator level forecasts and real-time GDP price deflator vintages. The SPF level workbook reports forecasts in columns denoted PGDP1, ..., PGDP6. For each survey origin, we use the

one-year inflation forecast

$$F_t^{SPF} = 100 [\log(\text{PGDP}5_t) - \log(\text{PGDP}1_t)], \quad (\text{C.1})$$

and compare it with the previous survey's forecast for the same calendar inflation window,

$$F_{t-1}^{SPF} = 100 [\log(\text{PGDP}6_{t-1}) - \log(\text{PGDP}2_{t-1})]. \quad (\text{C.2})$$

The same-target forecast revision is therefore

$$REV_t^{SPF} = F_t^{SPF} - F_{t-1}^{SPF}. \quad (\text{C.3})$$

The realized inflation rate is constructed from the real-time GDP price deflator for the same start and end quarters, using the first vintage available at least four quarters after the target quarter. The forecast error is

$$FE_t^{SPF} = \pi_t^{RT} - F_t^{SPF}. \quad (\text{C.4})$$

We then estimate

$$FE_t^{SPF} = a + \beta_{CG} REV_t^{SPF} + u_t. \quad (\text{C.5})$$

The sample is 1969:Q4 through 2007:Q4, with dates assigned by the target quarter. Standard errors are Newey-West HAC standard errors with a Bartlett kernel and lag truncation of four quarters. This procedure gives $\hat{\beta}_{CG} = 1.248$ with HAC standard error 0.494 and 148 observations, which is the empirical bar reported in Figure 3.

C.2 Forecasts of FHP Firms

The model forecasts are constructed from a firm with planning horizon k_t^* . Define the number of remaining planned periods relevant for a forecast of inflation h quarters ahead as

$$\iota_t(h) \equiv \max\{0, k_t^* - h\}. \quad (\text{C.6})$$

For $j \in \{y, \varepsilon\}$, define the h -step loading

$$A_{j,h}(\iota) \equiv \rho_j^h \frac{1 - (\beta \rho_j)^{\iota+1}}{1 - \beta \rho_j}. \quad (\text{C.7})$$

The date- t subjective forecast of quarterly inflation h quarters ahead is

$$E_t^{k_t^*} \hat{\pi}_{t+h} = \kappa A_{y,h}(\iota_t(h)) \hat{y}_t + A_{\varepsilon,h}(\iota_t(h)) \hat{\varepsilon}_{st} + \beta^{\iota_t(h)+1} \bar{\pi}_t. \quad (\text{C.8})$$

This expression is valid for both fixed and endogenous planning. In the fixed-planning model, k_t^* is constant. In the endogenous-planning model, forecasts made at date t condition on the planning horizon chosen at that forecast origin.

The model analogue of the CG regression uses forecasts of average inflation over the next four quarters. Define realized four-quarter average inflation from $t + 1$ through $t + 4$ as

$$\hat{\pi}_{t,t+4}^A \equiv \frac{1}{4} \sum_{h=1}^4 \hat{\pi}_{t+h}. \quad (\text{C.9})$$

The date- t forecast of this object is

$$F_{t,t+4}^4 \equiv E_t^{k_t^*} \hat{\pi}_{t,t+4}^A = \frac{1}{4} \sum_{h=1}^4 E_t^{k_t^*} \hat{\pi}_{t+h}. \quad (\text{C.10})$$

C.3 Forecast Errors, Revisions, and the CG Regression

The four-quarter-ahead forecast error is realized average inflation minus the date- t forecast:

$$FE_{t,t+4}^4 \equiv \hat{\pi}_{t,t+4}^A - F_{t,t+4}^4. \quad (\text{C.11})$$

The same-target forecast revision compares the date- t forecast with the forecast for the same inflation window made one quarter earlier:

$$REV_{t,t+4}^4 \equiv F_{t,t+4}^4 - F_{t-1,t+4}^4. \quad (\text{C.12})$$

The earlier forecast is

$$F_{t-1,t+4}^4 = \frac{1}{4} \sum_{h=2}^5 E_{t-1}^{k_{t-1}^*} \hat{\pi}_{t-1+h}, \quad (\text{C.13})$$

so it uses quarterly inflation forecasts at horizons two through five from the perspective of date $t - 1$.

The model analogue of the CG coefficient is estimated from simulated data using

$$FE_{t,t+4}^4 = a + \beta_{CG}^4 REV_{t,t+4}^4 + u_{t,t+4}. \quad (\text{C.14})$$

For fixed planning, the regression sets $k_t^* = k^*$ at every forecast origin. For endogenous planning, the regression uses the date-specific k_t^* chosen by the model at the forecast origin. Realized inflation uses future realized planning horizons, while the forecast made at date t

uses the planning horizon chosen at date t .

D Model-Based Inflation Uncertainty

In this section of the appendix, we describe how we compute the one-quarter-ahead conditional volatility of inflation in the endogenous-planning model. Because next-period's planning horizon is a nonlinear function, we use quadrature to approximate the conditional volatility of inflation.

We begin by defining the date- t state as

$$s_t = (\hat{y}_t, \hat{\varepsilon}_{st}, \bar{\pi}_t).$$

Given date- t information, current inflation is determined by (12), and the learning recursion (13) pins down next period's trend component:

$$\bar{\pi}_{t+1} = (1 - \gamma)\bar{\pi}_t + \gamma\hat{\pi}_t. \quad (\text{D.1})$$

Thus, when computing the one-step-ahead conditional variance, uncertainty comes from the next-period innovations to aggregate demand and the supply disturbance:

$$\hat{y}_{t+1}(z_y) = \rho_y \hat{y}_t + \sigma_y z_y, \quad z_y \sim \mathcal{N}(0, 1), \quad (\text{D.2})$$

$$\hat{\varepsilon}_{s,t+1}(z_s) = \rho_\varepsilon \hat{\varepsilon}_{st} + \sigma_s z_s, \quad z_s \sim \mathcal{N}(0, 1), \quad (\text{D.3})$$

where z_y and z_s are independent.

For any realization (z_y, z_s) , the next-period planning horizon is obtained from the same stopping rule as in (A.32). Specifically, define

$$\Delta p_{t+1}^{k+1}(z_y, z_s) \equiv \frac{\beta^{k+1}}{1 - \theta} [\kappa \rho_y^{k+1} \hat{y}_{t+1}(z_y) + \rho_\varepsilon^{k+1} \hat{\varepsilon}_{s,t+1}(z_s) - (1 - \beta)\bar{\pi}_{t+1}]. \quad (\text{D.4})$$

Then

$$k_{t+1}^*(z_y, z_s) = \min \left\{ k \in \{0, 1, 2, \dots\} \mid \frac{\mathcal{K}}{2} [\Delta p_{t+1}^{k+1}(z_y, z_s)]^2 \leq \Delta \tau_{k+1} \right\}. \quad (\text{D.5})$$

Given this horizon, the deterministic component of next-quarter inflation is

$$\hat{\pi}_{t+1}^{det}(z_y, z_s) = \kappa A_y(k_{t+1}^*) \hat{y}_{t+1}(z_y) + A_\varepsilon(k_{t+1}^*) \hat{\varepsilon}_{s,t+1}(z_s) + \beta^{k_{t+1}^*+1} \bar{\pi}_{t+1}, \quad (\text{D.6})$$

where $k_{t+1}^* = k_{t+1}^*(z_y, z_s)$.

Let $\{x_i, w_i\}_{i=1}^N$ denote order- N Gauss-Hermite nodes and weights. For any integrable function g of two independent standard-normal shocks,

$$\mathbb{E}_t[g(z_y, z_s)] \approx \frac{1}{\pi} \sum_{i=1}^N \sum_{j=1}^N w_i w_j g(\sqrt{2}x_i, \sqrt{2}x_j). \quad (\text{D.7})$$

Applying (D.7) to (D.6), the quadrature approximation to the conditional mean is

$$\hat{m}_t = \frac{1}{\pi} \sum_{i=1}^N \sum_{j=1}^N w_i w_j \hat{\pi}_{t+1}^{det}(\sqrt{2}x_i, \sqrt{2}x_j), \quad (\text{D.8})$$

and the corresponding conditional variance is

$$\hat{v}_t = \frac{1}{\pi} \sum_{i=1}^N \sum_{j=1}^N w_i w_j \left[\hat{\pi}_{t+1}^{det}(\sqrt{2}x_i, \sqrt{2}x_j) \right]^2 - \hat{m}_t^2. \quad (\text{D.9})$$

The model-based conditional volatility is therefore

$$\sigma_{\pi,t} = \sqrt{\hat{v}_t}. \quad (\text{D.10})$$

In the numerical implementation used for the simulations, we use $N = 5$ nodes in each shock dimension. The small nonnegativity adjustment in (D.10) only guards against roundoff error.

The fixed-planning benchmark is a useful special case. If the planning horizon is fixed at k , the next-period inflation rule is linear in the two innovations and the planning horizon does not move with (z_y, z_s) . In that case,

$$(\sigma_{\pi}^{fixed}(k))^2 = [\kappa A_y(k) \sigma_y]^2 + [A_{\varepsilon}(k) \sigma_s]^2. \quad (\text{D.11})$$

Hence fixed planning implies a constant conditional volatility. Endogenous planning generates time-varying volatility because the mapping from (z_y, z_s) to k_{t+1}^* changes with the current state $(\hat{y}_t, \hat{\varepsilon}_{st}, \bar{\pi}_t)$.

E Empirical Inflation Uncertainty

This section summarizes the density-based inflation uncertainty measure used in the quantitative exercise. The object is the one-quarter-ahead predictive standard deviation of U.S. core PCE inflation implied by the skew- t density model in [De Polis et al. \(2024\)](#).

Let y_t denote annualized quarterly core PCE inflation. The empirical model specifies the conditional predictive density as

$$y_t \mid Y_{t-1} \sim \text{Skt}_\nu(\mu_t, \sigma_t^2, \varrho_t),$$

where μ_t is location, σ_t is scale, $\varrho_t \in (-1, 1)$ governs asymmetry, and ν controls tail thickness. The asymmetry parameter tilts the density: positive values put more mass in the right tail and negative values put more mass in the left tail. Since the density is asymmetric, the conditional mean is

$$E_{t-1}y_t = \mu_t + g(\nu)\sigma_t\varrho_t,$$

where $g(\nu)$ is determined by the skew- t family. The term $g(\nu)\sigma_t\varrho_t$ is therefore the contribution of the balance of risks to the mean inflation forecast.

Time variation is introduced through a score-driven recursion for

$$f_t = (\mu_t, \log \sigma_t, \text{atanh}(\varrho_t))'.$$

These transformations keep the scale positive and the asymmetry parameter inside its admissible interval. Each component is decomposed into persistent and transitory parts, $f_t = \bar{f}_t + \tilde{f}_t$. After observing inflation at date t , the model updates both parts with the scaled score of the log predictive likelihood:

$$\bar{f}_{t+1} = \bar{f}_t + A s_t, \quad \tilde{f}_{t+1} = \Phi \tilde{f}_t + B s_t.$$

Here s_t is the likelihood score premultiplied by the inverse square root of the Fisher information matrix. In words, the location score responds to signed forecast errors, the scale score responds to excess squared forecast errors, and the asymmetry score responds differently to positive and negative errors. The Student- t tails dampen the influence of isolated large observations, so the update is less fragile than a Gaussian score update.

[De Polis et al. \(2024\)](#) estimate the static parameters in these updating equations using Bayesian methods. The priors put persistent dynamics on the transitory components and conservative inverse-gamma priors on the score loadings. Posterior simulation is carried out with an adaptive Metropolis-Hastings algorithm.