Beyond the Mean: Endogenous Expectation Weighting and the Shape of the Phillips Curve

Tibor Pál

4 November 2026

Abstract

This paper shows that the Phillips curve slope varies across the inflation distribution and evolves over time. We propose and apply a smoothed dynamic multiple quantile framework to the hybrid New Keynesian Phillips curve, in which the relative weight of survey-based inflation expectations and a recursively evolving inflation benchmark varies over time, while the Phillips curve slope varies across quantiles and over time. Using quarterly U.S. core PCE inflation data from 1960Q1 to 2025Q4, we find pronounced heterogeneity in the Phillips curve. The relative importance of survey-based expectations and the inflation benchmark changes markedly over time. The estimated slopes are asymmetric across quantiles, with a marked flattening in the upper tail. Lower-quantile slopes vary substantially over time, including pronounced steepening around the first quartile in the early 2010s, whereas the median slope remains relatively stable after the mid-1980s. Relatively elevated inflation expectations and strong inflation persistence, together with temporary flattening of the lower-tail Phillips curve slope, are consistent with the missing deflation following the Global Financial Crisis and during the COVID-19 pandemic. Finally, the filtered conditional inflation quantiles and their higher-order moments provide a quantitative measure of inflation tail risk, including persistent upside risks following the recent surge in inflation.

Keywords: Phillips curve; Inflation expectations; Inflation dynamics; Dynamic multiple quantiles; Score-driven models; Inflation-at-risk

JEL Classification: E31, C14, C32, E52

1 Introduction

"The literature has reached a limit on how much can be learned about the New Keynesian Phillips curve from aggregate macroeconomic time series."

— Mavroeidis, Plagborg-Møller, and Stock (2014)

Reconciling the canonical New Keynesian Phillips curve with the empirical behavior of inflation has been notoriously difficult. Since the seminal contribution of Phillips (1958), growing attention has been devoted to identifying the key drivers of inflation dynamics, particularly the relative roles of forward- and backward-looking price-setting behavior and the strength of the relationship between inflation and aggregate activity. Despite substantial methodological progress, empirical findings using limited information methods remain mixed, and a consensus has yet to emerge.

This study revisits the Phillips curve within a distributional framework by examining two dimensions of its dynamics for which empirical findings remain conflicting. First, we propose a tractable updating mechanism for the time-varying relative weight of inflation expectations and inflation inertia, consistent with recent theoretical findings on endogenous expectation switching (Cornea-Madeira et al., 2019). Second, we allow the Phillips curve slope to vary over time within a structure motivated by models of state-dependent pricing behavior (Harding et al., 2022), and examine its heterogeneity across the conditional inflation distribution.

Understanding these dynamics is central to monetary policy, as Phillips curve relationships are typically used to help identify unobserved macroeconomic quantities, such as the output gap and the natural rate of unemployment (Kuttner, 1994; Staiger et al., 1997; Laubach and Williams, 2003; Lewis and Vazquez-Grande, 2019; Pál and Storti, 2025). However, estimating the Phillips curve with limited-information methods is sensitive to the measurement of inflation expectations (Mavroeidis et al., 2014), while the stability of its slope remains debated (e.g., Hazell et al., 2022; Rossi et al., 2024). In addition, understanding how the conditional inflation distribution evolves is important for assessing inflation risks and their implications for monetary policy López-Salido and Loria (2024).

Recent research shows that it is important to look beyond the conditional mean and account for state dependence to understand inflation dynamics better. López-Salido and Loria (2024) finds that key distributional features are missed when focusing only on the conditional mean. Cornea-Madeira et al. (2019) shows that the relative importance of different expectation-formation rules can evolve endogenously. This is especially relevant in light of the structural-break evidence from Smith et al. (2025), which motivates a framework that captures nonlinearities, state dependence, and distributional heterogeneity in inflation dynamics.

Against this background and considering the long-standing empirical difficulty of reconciling the New Keynesian Phillips curve (NKPC) with U.S. inflation data,

this paper contributes to the inflation-dynamics literature by jointly modeling multiple conditional quantiles of U.S. inflation. Building on the Dynamic Multiple Quantile (DMQ) approach of Catania and Luati (2023), we introduce observation-driven parameter dynamics that admit an economic interpretation consistent with existing theory.

The DMQ framework serves as a natural starting point because it jointly models a broad set of conditional quantiles and utilizes information from various regions of the conditional distribution through its updating mechanism. The proposed extension introduces two economically interpretable time-varying parameter processes: updating the relative expectation weight offers a distributional analogue to performance-based expectation reweighting, and the Phillips curve slopes adjust to new distributional information according to their sensitivity to aggregate demand conditions.

Introducing quantile-specific parameters across a large set of conditional quantiles, however, substantially increases model dimensionality. To maintain a parsimonious parameterization, we propose a methodological extension of the DMQ framework, the *smoothed Dynamic Multiple Quantile* (sDMQ) model, which estimates the relevant parameters at a reduced set of target quantiles and recovers their values at the remaining quantiles through smooth interpolation. This construction preserves the distributional information used by the updating mechanism while reducing the number of freely estimated parameters. The resulting framework accommodates flexible, heterogeneous, and time-varying Phillips curve relationships within a tractable observation-driven specification, while allowing the evolution of expectation weights and slope parameters to be interpreted economically.

Our empirical analysis of quarterly U.S. core inflation reveals substantial nonlinearity in both dimensions. The relative importance of survey-based inflation expectations and inflation inertia changes markedly over time, while the Phillips curve slope exhibits an asymmetric profile that flattens toward the upper tail. Time variation in the slope is concentrated in the lower and middle regions of the distribution, with both its shape and magnitude shifting across macroeconomic regimes. These findings remain robust to alternative measures of inflation expectations, although the persistence and variability of the estimated slopes differ across specifications.

The historical evolution of these relationships provides further insights into U.S. inflation dynamics. During the 1970s, an increase in the relative weight on inflation expectations is consistent with weakly anchored expectations and the high-inflation regime of the Great Inflation. Following the Global Financial Crisis, the flattening of the Phillips curve in the lower tail, combined with strong inflation persistence, is consistent with the *missing deflation* observed during the recovery. In the post-COVID period, further changes in the slope profile and a renewed role for expectations highlight the importance of state dependence and time variation in inflation dynamics.

2 A Dynamic Distributional Phillips Curve

Studies relying on conventional quantile regression provide important insights into U.S. inflation dynamics that are often obscured when the focus is on the conditional mean (e.g., Xu et al., 2015; López-Salido and Loria, 2024). Typically these applications, however, estimate conditional quantiles in isolation and assume constant parameters within estimation samples, potentially limiting the exploitation of cross-quantile information and the ability to accommodate evolving relationships. Moreover, specifications incorporating autoregressive (AR) dynamics require selecting a lag structure ex ante, introducing another potential source of misspecification.

These limitations are particularly relevant for aggregate NKPC specifications, which have long faced difficulties in identifying the effects of inflation expectations and aggregate demand conditions on inflation (Mavroeidis et al., 2014). We therefore build on the Dynamic Multiple Quantile (DMQ) framework of Catania and Luati (2023), which jointly models a large set of conditional quantiles and allows them to evolve over time. The DMQ model belongs to the class of observation-driven models (Cox, 1981; Creal et al., 2013; Harvey, 2013), in which conditional quantiles are updated recursively through a quasi score-driven framework (Blasques et al., 2023). Importantly, its forcing variables exploit information from the position of observed inflation relative to a dense grid of conditional quantiles, allowing the updating mechanism to incorporate signals from all regions of the distribution.

We extend the DMQ framework in two important directions. First, we introduce *endogenous expectation weighting*, in which the relative importance of survey-based inflation expectations and past inflation dynamics evolves over time. Second, we allow the Phillips curve slope to vary both across quantiles and over time. Each parameter process follows an observation-driven updating scheme that, as we show, provides a statistically grounded and economically interpretable basis for their evolution.

Assigning a separate time-varying Phillips curve slope process to every modeled quantile would introduce a large number of parameters, particularly when accounting for time variation. In the proposed smoothed DMQ (sDMQ) model, slope parameters are instead specified at a reduced set of economically meaningful *target quantiles*. Their values at the remaining probability levels are obtained through smooth interpolation, making them deterministic functions of the target parameters. The associated *auxiliary quantiles* are retained in the construction of the forcing variables, allowing the target slope parameters to be updated using information from the full quantile grid without requiring a separate slope process at each probability level. This provides a parsimonious and economically interpretable specification of time-varying Phillips curve slopes.

Let $0 < \tau_1 < \cdots < \tau_J < 1$ denote the probability levels used to model the conditional inflation distribution. Among these, $P < J$ levels are selected as target quantiles, indexed by $p = 1, \ldots, P$, while the remaining $J - P$ levels are treated

as auxiliary quantiles. At each target probability level τ_p, the conditional inflation quantile is specified as

$$q_t^{\tau_p}(\pi) = \tilde{q}_t^{\tau_p}(\pi) + \kappa_{p,t}\, x_{t-1}, \tag{2.1}$$

where $q_t^{\tau_p}(\pi)$ denotes the conditional inflation quantile at probability level τ_p, given the information available at time $t-1$. These quantiles are obtained recursively through the observation-driven filtering mechanism and represent one-step-ahead quantile forecasts. Throughout the paper, the terms *filtered quantiles*, *predicted quantiles*, and *one-step-ahead quantile forecasts* are used interchangeably to refer to these quantities, which are determined before observing inflation at time t. The term x_{t-1} represents a proxy for real marginal cost, with a quantile-specific and time-varying slope parameter $\kappa_{p,t}$. The component $\tilde{q}_t^{\tau_p}(\pi)$ denotes the dynamic quantile process inherited from the DMQ framework, in which quantiles are constructed around a reference conditional median using positive, time-varying quantile spacings.

The two extensions are detailed in the following subsections. Subsection 2.1 modifies the reference median process to incorporate survey-based inflation expectations and allows their relative weight to evolve over time. Subsection 2.2 introduces the time-varying slope processes and the interpolation scheme used to construct the auxiliary parameters. The baseline DMQ recursion, forcing-variable construction, and related properties inherited from Catania and Luati (2023) are summarized in Appendix A.

2.1 Endogenous Expectation Weighting

The assumption that agents form expectations rationally has long been challenged in the macroeconomic literature. The empirical difficulties associated with the purely forward-looking New Keynesian Phillips Curve (NKPC) have made this tension particularly salient, as its conventional specification under rational expectations has difficulty matching the strong persistence observed in U.S. inflation dynamics (Fuhrer and Moore, 1995; Galí and Gertler, 1999; Rudd and Whelan, 2005). These empirical shortcomings have contributed to the prevalent use of hybrid versions of the NKPC. These specifications introduce an explicit backward-looking component to capture inflation persistence (Fuhrer and Moore, 1995; Rudd and Whelan, 2005), consistent with models in which a fraction of price setters follow rules based on past inflation (Galí and Gertler, 1999). More generally, the empirical and theoretical limitations associated with model-consistent rational expectations have motivated a broader literature on alternative and heterogeneous expectation-formation mechanisms (Mavroeidis et al., 2014).

While incorporating information on past inflation substantially improved the empirical fit of the new Phillips curve, empirical studies have not reached a con-

sensus on the relative importance or specification of the forward- and backward-looking components (Galí and Gertler, 1999; Sbordone, 2005; Lindé, 2005; Rudd and Whelan, 2006). In addition, evidence suggests that expectation formation is heterogeneous and that the degree of disagreement among forecasters varies considerably over time (Mankiw et al., 2003). Together, these findings motivate models in which the relative weight on inflation expectations and inflation inertia is allowed to evolve over time. Notably, Cornea-Madeira et al. (2019) develop a heterogeneous-expectations Phillips curve in which agents switch between forward- and backward-looking expectation-formation rules according to their recent relative forecasting performance, generating substantial time variation in the weights assigned to the two components. Related models with endogenous weights similarly emphasize that the relative importance of alternative expectation-formation rules need not remain static over time (Branch and Evans, 2006; Cornea-Madeira and Madeira, 2022).

Motivated by these findings and building on the idea that relative weights may respond endogenously to predictive performance, we allow the relative expectation weight to evolve within a quasi-score-driven framework (Blasques et al., 2023). The dynamic multiple quantile (DMQ) framework of Catania and Luati (2023) provides a natural basis for this extension, as its time-varying conditional quantiles evolve through a deterministic, observation-driven updating mechanism based on the gradient of the quantile check loss function. The resulting mechanism is conceptually related to performance-based reweighting. In the special case of median, the underlying quantile loss is proportional to absolute forecast error, whereas the general specification exploits information from multiple regions of the conditional inflation distribution. The updating signal reflects both the discrepancy between realized and predicted inflation and the sensitivity of the conditional quantiles to the relative expectation weight. Thus, unlike explicit behavioral switching models, the proposed approach does not model individual agents' choice of forecasting rules, but provides an economically interpretable econometric mechanism through which their relative importance adjusts to incoming information.

To introduce the time-varying relative weight, we specify the reference quantile at the conditional median in Eq. 2.1 as

$$\tilde{q}_t^* = (1 - \beta_t)\pi_{t-1}^E + \alpha u_{t-1}^* + \beta_t \tilde{q}_{t-1}^*, \tag{2.2}$$

where π_{t-1}^E denotes the survey-based inflation expectation available at time $t-1$. The forcing variable u_{t-1}^* enters the reference-quantile dynamics with loading parameter α. We use lagged survey expectations because they are predetermined relative to current inflation, which helps mitigate endogeneity concerns (Mavroeidis et al., 2014). This timing convention is also consistent with models of sticky information and inattentive price setting, in which firms do not instantaneously acquire and incorporate all contemporaneous macroeconomic information into their pricing decisions (Mankiw and Reis, 2002; Reis, 2006).

Before turning to the dynamics of the weight parameter β_t, it is useful to clarify the interpretation of the two components it governs. First, the component $\tilde{q}^*_{t-1}$ offers more than a conventional backward-looking term typically specified through a finite lag operator. Given the dynamic nature of the framework, it summarizes information embedded in the conditional inflation distribution and accumulated through the history of the inflation process, with $\tilde{q}^*_t = Q_{0.50}(\pi_t \mid \mathcal{I}_{t-1})$. It therefore reflects not only past inflation realizations but also previously incorporated expectations, aggregate demand conditions, and past innovations transmitted through the dynamics of the model. Consequently, the model's recursive structure reduces reliance on an explicitly imposed finite lag structure and may mitigate misspecification associated with restrictive lag specifications (Mavroeidis et al., 2014). We therefore interpret $\tilde{q}^*_{t-1}$ as a distributionally informed, model-consistent *inflation benchmark* that incorporates a broader information set than a conventional lagged-inflation term.

Second, π^E_{t-1} provides an external, survey-based signal of expected inflation that is not generated by the internal dynamics of the model. The specification therefore combines an observed expectations anchor with an endogenous benchmark summarizing information already incorporated into the inflation process. This distinction is analogous in spirit to Malmendier and Nagel (2016), who combine an experience-based component of inflation expectations with an unobserved common information component that may reflect the influence of professional forecasts.

The two components are also linked over time. Since survey-based expectations enter the reference quantile with weight $1 - \beta_t$, their influence can persist through the recursive evolution of the endogenous inflation benchmark. This provides an explicit channel through which expectations incorporated in earlier periods affect subsequent reference quantiles.

We model the relative weight between these two components using the logistic transformation

$$\beta_t = \frac{1}{1 + \exp(-\bar{\beta}_t)},$$

which ensures that $0 < \beta_t < 1$ and that the weights of the two components are strictly positive and sum to unity. Specifically, $1 - \beta_t$ is the weight on explicit inflation expectations, while β_t is the weight on the inflation benchmark q^*_{t-1}. The unrestricted parameter $\bar{\beta}_t$ evolves according to

$$\bar{\beta}_t = (1 - \phi^\beta)\bar{\beta}_0 + \phi^\beta \bar{\beta}_{t-1} + \rho_\beta u^\beta_{t-1}, \tag{2.3}$$

where ϕ^β denotes the AR parameter, ρ_β is the loading coefficient on the forcing variable u^β_{t-1}, and $\bar{\beta}_0$ determines the long-run relative weight of the two components. Since β_t represents a time-varying relative transmission weight, it shares some conceptual similarities with the fraction of firms following different price-setting rules in hybrid NKPC models.

We construct the forcing variable from the gradient of the aggregate quantile loss. Since β_t enters the reference median and shifts the remaining quantiles through

the common reference process, the marginal effect of β_t on each conditional quantile is given by

$$\frac{\partial q_t^{\tau_j}}{\partial \beta_t} = \tilde{q}_{t-1}^{*} - \pi_{t-1}^{E}.$$

Using $\partial \beta_t / \partial \bar{\beta}_t = \beta_t(1-\beta_t)$, the forcing variable for expectation weighting is therefore

$$\begin{aligned} u_t^{\beta} &= -\frac{\partial}{\partial \bar{\beta}_t} \sum_{j=1}^{J} \rho_{\tau_j}\left(y_t - q_t^{\tau_j}\right) \\ &= -\sum_{j=1}^{J} \frac{\partial \rho_{\tau_j}}{\partial q_t^{\tau_j}} \frac{\partial q_t^{\tau_j}}{\partial \beta_t} \frac{\partial \beta_t}{\partial \bar{\beta}_t} \\ &= \underbrace{(-z_t)}_{\text{distributional inflation surprise}} \times \underbrace{\beta_t(1-\beta_t)}_{\text{reweighting capacity}} \times \underbrace{\left(\tilde{q}_{t-1}^{*} - \pi_{t-1}^{E}\right)}_{\text{expectations–inertia wedge}}, \end{aligned} \tag{2.4}$$

where $z_t = \sum_{j=1}^{J} z_{j,t}$ denotes the aggregate quantile-hit signal, which captures distributional inflation news relative to the model-implied conditional distribution. It therefore summarizes unexpected movements in realized inflation and may reflect unobserved cost-push and other inflationary shocks. Combined with the divergence between the inflation benchmark and survey expectations, the sign and magnitude of this innovation determine the direction and strength of the reweighting response. The forcing variable thus depends jointly on (i) the distributional inflation surprise, (ii) the sensitivity of the relative weight to changes in its latent state, and (iii) the divergence between the endogenous inflation benchmark and explicit inflation expectations. Consequently, the forcing variable is zero if there is no distributional inflation surprise or if the two benchmarks coincide.

This updating mechanism is closely related to the performance-based expectation switching of Cornea-Madeira et al. (2019). Their model assigns greater weight to the expectation rule with superior recent forecasting performance, measured using differences in absolute forecast errors. In the present framework, relative predictive performance is evaluated through the marginal effect of changing β_t on the aggregate quantile loss. Hence, the forcing variable provides a distributional analogue of performance-based reweighting rather than relying on an externally specified comparison of historical forecast errors.

The connection becomes particularly clear in the special case where only the conditional median is considered. Setting $J = 1$ and $\tau_{j^*} = 0.50$ gives

$$\rho_{0.50}(y_t - q_t^{*}) = \frac{1}{2}\left|y_t - q_t^{*}\right|,$$

so that the loss underlying the updating mechanism reduces, up to scale, to absolute forecast error, the performance criterion employed by Cornea-Madeira et al. (2019). However, whereas their behavioral switching mechanism evaluates the relative forecasting performance of competing expectation-formation rules, the proposed approach updates their relative contribution according to the sensitivity of

the combined forecast error to the expectation weight. The two mechanisms therefore share an underlying performance-based motivation, while differing in how forecasting performance translates into adjustments of the relative weight.

The general specification extends this idea along three dimensions. First, the inflation benchmark $\tilde{q}^*_{t-1}$ is not restricted to a simple lagged-inflation forecasting rule, but recursively summarizes information accumulated through the history of the inflation process. Second, although the relative weight β_t enters directly only the reference quantile process, its updating mechanism exploits information from a dense set of conditional inflation quantiles, allowing forecast errors across different regions of the distribution to inform the relative importance of the two components. Third, unlike Cornea-Madeira et al. (2019), the model does not require a predetermined window over which past forecast performance is evaluated. Instead, past information is transmitted through the persistence of the latent weight and the recursive evolution of the inflation benchmark. The resulting mechanism is conceptually related to performance-based reweighting, with adjustments to the relative expectation weight driven by incoming information from different regions of the conditional inflation distribution.

The recursive structure also makes the propagation of previously incorporated expectations explicit. Since

$$\tilde{q}^*_t = (1-\beta_t)\pi^E_{t-1} + \beta_t \tilde{q}^*_{t-1} + \cdots$$

and

$$\tilde{q}^*_{t-1} = (1-\beta_{t-1})\pi^E_{t-2} + \beta_{t-1} \tilde{q}^*_{t-2} + \cdots ,$$

recursive substitution yields

$$\tilde{q}^*_t = (1-\beta_t)\pi^E_{t-1} + \beta_t(1-\beta_{t-1})\pi^E_{t-2} + \beta_t\beta_{t-1}\tilde{q}^*_{t-2} + \cdots .$$

Thus, expectations that enter explicitly in one period become incorporated into the reference inflation process and are subsequently propagated through the model dynamics. We refer to this mechanism as *endogenous propagation of expectations*.

We close the section by illustrating the endogenous adjustment of the relative expectation weight through a simple example. Consider an initial steady state in which inflation follows its long-run reference dynamics and the relative weight β_t is close to its long-run level. Suppose that inflation subsequently rises above the conditional median and is realized in the upper region of the conditional inflation distribution. Realized inflation then exceeds a sufficiently large share of the predicted quantiles, causing the aggregate quantile-hit signal z_t to become negative and $-z_t$ to capture positive distributional inflation news. If survey expectations and the updated benchmark subsequently diverge, the innovation generates a nonzero forcing variable,

$$u^{\beta}_{t+1} = (-z_{t+1})\beta_{t+1}(1-\beta_{t+1})\left(\tilde{q}^*_t - \pi^E_t\right).$$

The direction of reweighting then depends jointly on the sign of the distributional inflation news and on the relative position of the two benchmarks. For example,

if $-z_{t+1} > 0$ and survey expectations exceed the benchmark, so that $\tilde{q}_t^* - \pi_t^E < 0$, the forcing variable becomes negative and, for $\rho_\beta > 0$, reduces β_{t+2}, thereby increasing the relative weight placed on survey expectations. Reweighting therefore arises from the interaction between new distributional information and disagreement between the inflation benchmark and survey expectations. At the same time, expectations previously incorporated into $\tilde{q}_t^*$ continue to propagate through the reference state process.

2.2 Time-Varying Phillips Curve Slopes Across Quantiles

The seemingly contradictory presence of persistent real effects of monetary policy and relatively frequent price adjustments at the disaggregated level has long challenged efforts to reconcile macroeconomic and microeconomic evidence (Bils and Klenow, 2004; Nakamura and Steinsson, 2008; Gopinath and Itskhoki, 2010; Harding et al., 2022). While conventional time-dependent pricing models such as Calvo (1983) offer a tractable explanation for aggregate nominal rigidity, they impose exogenous price-adjustment opportunities. In contrast, standard state-dependent pricing models such as that of Dotsey et al. (1999) allow endogenous pricing decisions but typically generate smaller real effects of monetary shocks than comparable time-dependent models (Golosov and Lucas, 2007).

Gertler and Leahy (2008) make an important advance on this matter. They show that combining real rigidities with state-dependent pricing can reconcile frequent individual price adjustments with substantial aggregate nominal price stickiness. Real rigidities, including strategic complementarities in price setting, attenuate firms' responses to changes in marginal costs (Gertler and Leahy, 2008; Gopinath and Itskhoki, 2010).

Incorporating real rigidities into a state-dependent pricing framework helps reconcile microeconomic evidence on firm-level price adjustments with aggregate nominal rigidity. Yet the Phillips curve derived by Gertler and Leahy (2008) still relies on a log-linear approximation around a zero-inflation steady state. This approximation may break down during large economic disturbances, such as the Global Financial Crisis and the COVID-19 pandemic, motivating consideration of nonlinear aggregate inflation dynamics (Harding et al., 2022, 2023). Specifically, Harding et al. (2022) demonstrate that real rigidities arising from Kimball demand aggregation produce a nonlinear and asymmetric Phillips curve. Their model implies relatively flat inflation responses during deep recessions and steeper responses during expansions, reflecting differences in firms' incentives to adjust prices across economic states. Harding et al. (2023) further examine nonlinear inflation transmission in the context of post-pandemic inflation dynamics. Related empirical evidence is provided by Forbes et al. (2021), who find that the Phillips curve becomes flatter when output falls below potential in low-inflation environments, consistent with downward nominal wage and price rigidities.

An important implication of this literature is that the sensitivity of inflation to real activity need not remain constant over the business cycle. Notably, the nonlinearities in Harding et al. (2022) arise even with constant structural parameters because underlying price-setting incentives depend on prevailing economic conditions. While these mechanisms provide an economic motivation for allowing the Phillips curve slope to vary, an empirical model need not impose their specific structural implications directly. Instead, a flexible specification can accommodate changing inflation sensitivity and allow the data to determine its direction and magnitude.

Motivated by these considerations, we allow the Phillips curve slope to vary over time in a state-dependent fashion. To clarify the economic mechanism underlying the proposed updating process, we first consider a special case in which the slope is common across all conditional inflation quantiles. Specifically, the conditional quantiles are given by

$$q_t^{\tau_j}(\pi) = \tilde{q}_t^{\tau_j}(\pi) + \kappa_t x_{t-1}, \tag{2.5}$$

where x_{t-1} denotes the lagged output gap and κ_t is the common slope parameter, entering the model through an exponential link function to ensure positivity, i.e., $\kappa_t = \exp(\bar{\kappa}_t)$. The component $\tilde{q}_t^{\tau_j}(\pi)$ summarizes the remaining determinants of the conditional inflation quantiles, including the expectation dynamics discussed in the previous subsection. Thus, the specification distinguishes adjustments in the inflation benchmark from changes in inflation's sensitivity to real activity.

We specify the evolution of the common slope parameter as

$$\bar{\kappa}_t = (1 - \phi^\kappa)\bar{\kappa}_0 + \phi^\kappa \bar{\kappa}_{t-1} + \rho^\kappa u_{t-1}^\kappa, \tag{2.6}$$

where ϕ^κ governs inertia, ρ^κ is the loading coefficient on the forcing variable u_{t-1}^κ, and $\bar{\kappa}_0$ determines the long-run reference log slope around which the process evolves in a shock-free environment.

Analogously to the endogenous expectation weight, we construct the forcing variable based on the gradient of the aggregate quantile loss with respect to the time-varying slope parameter. While the exact gradient reads as

$$u_{t,\text{grad}}^\kappa = -\frac{\partial}{\partial \bar{\kappa}_t} \sum_{j=1}^{J} \rho_{\tau_j}\big(\pi_t - q_t^{\tau_j}(\pi)\big) = (-z_t) x_{t-1} \exp(\bar{\kappa}_t), \tag{2.7}$$

we depart from this original formulation in one important respect. Motivated by models of state-dependent price adjustment and nonlinear inflation transmission (Gertler and Leahy, 2008; Harding et al., 2022), we modify the gradient-based forcing variable so that the magnitude, rather than the level, of the output gap affects how the slope adjusts over time. In addition, we use the output gap at time t as u_t^κ determines the slope parameter for the following period. Since u_t^κ determines the following period's slope parameter, this also aligns the slope update

with the output gap entering the subsequent period's conditional inflation forecasts. The modification is consistent with the broader quasi score-driven framework of Blasques et al. (2023), which allows the updating function to differ from the exact score. As shown in the empirical analysis, the considerable improvement in model fit further supports our modeling choice.

Accordingly, the preferred forcing variable is specified as

$$u_t^{\kappa} = \underbrace{(-z_t)}_{\text{distributional inflation news}} \times \underbrace{|x_t|}_{\text{real activity}} \times \underbrace{\exp(\bar{\kappa}_t)}_{\text{slope sensitivity}} , \tag{2.8}$$

where the aggregate quantile-hit signal captures distributional inflation forecast errors, while the absolute output gap scales their influence on slope updating.

These features reflect the idea that larger deviations of output from potential amplify adjustments in the sensitivity of inflation to real activity. The proposed specification therefore provides a flexible, state-dependent updating mechanism without imposing a specific recession-expansion asymmetry or identifying the underlying price-setting frictions.

The homogeneous specification allows inflation sensitivity to change over time but assumes a constant Phillips curve slope across the conditional distribution. We now relax this assumption by allowing the slopes to vary across quantiles as well as over time.

As discussed previously, Catania et al. (2026) allow exogenous covariates to have different effects across quantiles while preserving the quantile ordering of the original DMQ framework. To achieve this, the additive covariate component is specified as a noncrossing quantile function. They implement this through a nonlinear location-scale specification that also captures state dependence. Under this specification, marginal effects are monotonic across quantiles for a given covariate state. At the same time, their variation over time reflects changes in the covariate rather than observation-driven updates to quantile-specific covariate-effect parameters.

Since there is no clear theoretical reason to impose a monotonic Phillips curve slope profile across the inflation distribution, we take a different approach in the proposed sDMQ framework. We allow slopes to vary non-monotonically across quantiles while preserving their economic interpretation and the option to impose theoretically motivated restrictions directly on them.

Accommodating such flexibility, however, comes with several practical challenges. Because the DMQ updating mechanism relies on a dense grid of quantiles, estimating a separate time-varying slope process for each probability level would enlarge the parameter space. Using common slopes for groups of adjacent quantiles would reduce this number but create discontinuities in the slope profile. Moreover, allowing slopes to differ across quantiles can lead to quantile crossing, as their ordering is no longer guaranteed by construction.

To address these limitations, the proposed sDMQ framework reduces the number of parameters while retaining heterogeneity and smoothness across quantiles.

We specify a smaller set of time-varying slope parameters, $\hat{\kappa}_{p,t}$, at economically meaningful *target probability levels*. We consider $\tau \in 0.05, 0.25, 0.75, 0.95$, used in quantile-based studies of macroeconomic risk (e.g., Adrian et al., 2019; López-Salido and Loria, 2024). Additionally, we include the median $\tau = 0.50$, as it provides a natural reference for economic interpretation. We obtain the parameters at the remaining *auxiliary probability levels* through smooth interpolation of the target parameters, yielding a complete, regularized slope profile without requiring a separate parameter process at each probability level.

At this point, it is important to note that the key role of the auxiliary quantiles is to construct the forcing variables, allowing the updating mechanism to exploit information from a dense grid of conditional quantiles. While all quantiles enter the estimation criterion, potential crossings are penalized only at the target quantiles. Although this does not guarantee global quantile monotonicity, the estimation procedure follows established approaches to noncrossing quantile regression (Fakoor et al., 2023; Tang et al., 2026). Further details are provided in Subsection 2.3.

Building on the proposed modifications to the modeling framework, Eq. 2.1 is reformulated as

$$q_t^{\tau_j}(\pi) = \tilde{q}_t^{\tau_j}(\pi) + \exp(\kappa_{j,t}) x_{t-1}, \tag{2.9}$$

where $\kappa_{j,t}$ denotes the quantile-specific log slope and the full set of conditional inflation quantiles $q_t^{\tau_j}(\pi)$ comprises the directly modeled target quantiles $\hat{q}_t^{\tau_p}(\pi)$ and the auxiliary quantiles $\bar{q}_t^{\tau_j}(\pi)$ at the remaining probability levels. The complete parameter vector $\boldsymbol{\kappa}_t = (\kappa_{1,t}, \ldots, \kappa_{J,t})'$ is constructed from the vector of target parameters $\hat{\boldsymbol{\kappa}}_t = (\hat{\kappa}_{1,t}, \ldots, \hat{\kappa}_{P,t})'$, with the remaining elements obtained through interpolation described in Appendix B.

Turning to the specification of the target parameters, their time-varying processes follow autoregressive dynamics:

$$\hat{\kappa}_{p,t} = (1 - \phi^{\kappa})\hat{\kappa}_{p,0} + \phi^{\kappa}\hat{\kappa}_{p,t-1} + \rho_p u_{p,t-1}^{\kappa}, \tag{2.10}$$

where ϕ^{κ} denotes the common AR parameter, ρ_p is the loading coefficient on the forcing variable $u_{p,t-1}^{\kappa}$, and $\hat{\kappa}_{p,0}$ determines the long-run reference slope parameter at target quantile p.

Analogously to the homogeneous case, we start from the gradient of the aggregate quantile loss with respect to the target log-slope parameters. The effect of each target parameter on the conditional quantiles is given by

$$\frac{\partial q_t^{\tau_j}}{\partial \hat{\kappa}_{p,t}} = \exp(\kappa_{j,t}) x_{t-1} \frac{\partial \kappa_{j,t}}{\partial \hat{\kappa}_{p,t}}.$$

As in the homogeneous case, we replace the lagged output gap with its contemporaneous absolute value. The preferred forcing variable is therefore

$$u_{p,t}^{\kappa} = \underbrace{|x_t|}_{\text{real activity}} \sum_{j=1}^{J} \underbrace{(-z_{j,t})}_{\text{distributional news}} \underbrace{\exp(\kappa_{j,t})}_{\text{slope sensitivity}} \underbrace{\frac{\partial \kappa_{j,t}}{\partial \hat{\kappa}_{p,t}}}_{\text{interpolation sensitivity}}. \tag{2.11}$$

The forcing variable links inflation news to the magnitude of contemporaneous deviations from potential output, as in the homogeneous specification. However, two additional sources of heterogeneity arise. First, the exponential term reflects differences in prevailing Phillips curve slopes across quantiles, so the same output gap and inflation forecast error can contribute differently to target-parameter updates. Second, the interpolation derivative determines how information from different parts of the distribution enters the update of a given target parameter. Consequently, each target parameter is adjusted using a weighted combination of distributional inflation forecast errors, with weights determined by the prevailing slope profile and interpolation structure. For brevity, the derivative of the interpolation term is presented in Appendix B.

With this construction, each time-varying target parameter responds not only to its own quantile forecast error but also to errors from the auxiliary quantiles affected by interpolation. This allows information from the relevant regions of the distribution to contribute to the parameter updates while keeping the updating mechanism tractable.

The resulting mechanism complements the endogenous expectation weighting discussed in the previous subsection, but operates through a different channel. The expectation-weight mechanism adjusts the relative influence of survey expectations and the inflation benchmark according to their divergence and the distributional forecast signal. The slope mechanism instead adjusts the contribution of real activity to predicted inflation. Consequently, inflation forecast errors can influence both dimensions of the model, allowing changes in the inflation benchmark and changes in inflation sensitivity to contribute jointly to the evolution of the conditional inflation distribution.

2.3 Estimation and Identification

Unlike the original DMQ and DMQ-X specifications, where quantile ordering is ensured by construction (Catania and Luati, 2023; Catania et al., 2026, respectively), the proposed sDMQ specification does not guarantee global quantile monotonicity. We address potential crossings at the target quantiles differently in the static and time-varying specifications. In the static sDMQ specification, noncrossing is imposed through a hard constraint on the ordering of the target quantiles, consistent with the constrained quantile regression approach of Bondell et al. (2010). In the time-varying specification, quantile-specific parameter updates can generate crossings even when the slope processes are initialized at the constrained static estimates. Imposing hard constraints throughout the resulting dynamic paths would further restrict the evolution of the slope processes. We therefore incorporate a soft crossing penalty into the estimation objective consistent with penalized noncrossing quantile estimation approaches (Fakoor et al., 2023; Shen et al., 2024; Tang et al., 2026).

Our primary objective is to analyze inflation dynamics at the target probability levels and estimate their associated Phillips curve slope parameters. Although the estimation objective could be restricted to the target quantiles, we include both target and auxiliary quantiles. The empirical results indicate that including auxiliary quantiles improves model fit. Their inclusion allows information from the full quantile grid to contribute to both parameter estimation and the construction of the forcing variables. The updating mechanism nevertheless differs from the exact gradient of the estimation loss, consistent with the quasi score-driven framework of Blasques et al. (2023).

We follow the estimation approach of Catania and Luati (2023), which builds on the multi-quantile estimation method of White et al. (2015). Specifically, estimation relies on the quantile check loss function, which admits an interpretation as the negative log-likelihood of an asymmetric Laplace working density (Koenker and Machado, 1999; Poiraud-Casanova and Thomas-Agnan, 2000; Kotz et al., 2001). More generally, the asymmetric Laplace density belongs to the tick-exponential family introduced by Komunjer (2005). Thus, minimizing the quantile check loss admits a quasi-maximum likelihood (QML) interpretation without requiring the asymmetric Laplace distribution to represent the true conditional distribution. The observation-driven updating mechanism is consistent with the quasi score-driven (QSD) framework of Blasques et al. (2023), which allows the updating function to differ from the score implied by the estimation likelihood.

The model parameters are estimated by minimizing the aggregate quantile check loss over both target and auxiliary probability levels. Accounting for the crossing penalty in the time-varying specification, the estimation objective reads as

$$\hat{\theta}_T = \arg\min_{\theta\in\Theta} \left\{ \sum_{t=1}^{T} \sum_{j=1}^{J} \rho_{\tau_j}\big(\pi_t - q_t^{\tau_j}(\theta)\big) + \lambda \mathcal{C}_T(\theta) \right\}, \tag{2.12}$$

where $q_t^{\tau_j}(\theta)$ denotes the conditional inflation quantile at probability level τ_j, including both target and auxiliary quantiles, and θ collects the parameters to be estimated. The term $\mathcal{C}_T(\theta)$ denotes a nonnegative penalty for crossings among the target quantiles, with λ governing its strength. In the static specification, $\lambda = 0$ and the admissible parameter space Θ incorporates hard noncrossing constraints on the target quantiles. In the time-varying specification, $\lambda > 0$ and potential crossings among the target quantiles are penalized through the objective function. Thus, both target and auxiliary quantiles contribute to the estimation loss, while crossing restrictions are imposed only on the target quantiles.

Estimation proceeds in two stages. We first estimate the static sDMQ specification, allowing the Phillips curve slope to differ across target quantiles while keeping the slope parameters and the expectations weight constant over time. These estimates are then used as fixed long-run reference values and initial values in the time-varying sDMQ specification, where the expectations weight and target slope parameters follow the quasi score-driven updating mechanisms introduced above.

To improve numerical stability and mitigate weak identification of the quantile-specific Phillips curve parameters, we impose a positive lower bound of 0.025 on the slopes in the static sDMQ specification, such that $\exp(\hat{\kappa}_p^S) \geq 0.0025$. Lastly, to avoid extrapolating beyond the extreme target probability levels, we hold the slope parameters constant below $\tau = 0.05$ and above $\tau = 0.95$, using the values at the corresponding boundary quantiles.

3 Data

To analyze the dynamics of US inflation, the seasonally adjusted annualized quarterly growth rate of the core Personal Consumption Expenditures (PCE) price index is used over the period 1960Q1 to 2025Q4. Core PCE inflation serves as the key measure of inflation for the Federal Reserve (Fed), capturing underlying price trends by excluding volatile food and energy components and providing a comprehensive measure of consumer prices that reflects evolving spending patterns.

In our baseline specification, the output gap is defined as the difference between the logarithm of real GDP and the logarithm of potential output, which serves as a proxy for real economic slack. The output gap series is based on estimates from the Congressional Budget Office (CBO). Our choice of the CBO series is motivated by its structural nature, as potential output is derived from a production function framework that explicitly accounts for trend labor input, capital services, and total factor productivity. In addition, the CBO output gap is widely used in empirical macroeconomic analysis, ensuring comparability with the existing literature.

Additionally, we consider two alternative measures of aggregate economic activity: the vacancy-to-unemployment ratio and the negative unemployment gap. The former is constructed following Benigno and Eggertsson (2023), while the latter is defined as the difference between the NAIRU and the observed unemployment rate. All series are obtained from the Federal Reserve Economic Data (FRED) database.

We proxy short-run forward-looking inflation expectations using two alternative measures. In the baseline analysis, we use the one-year-ahead CPI inflation forecasts from the Survey of Professional Forecasters (SPF), provided by the Federal Reserve Bank of Philadelphia at a quarterly frequency and available from 1981Q3 onward. As an alternative, we use the one-year inflation expectations series from the Federal Reserve Bank of Cleveland, obtained from FRED and available from 1982Q1 onward. Following the approach of Benigno and Eggertsson (2023), both series are extended backward to 1960Q1 using 12-month inflation expectations from the Livingston Survey. Since the latter is available semiannually, missing quarterly observations are interpolated using a shape-preserving spline method.

Figure 1 displays the inflation series and Phillips curve determinants used in the analysis. The upper panel shows US core PCE inflation alongside the baseline SPF expectations over the sample period. The lower panel reports the CBO output gap over the sample period. The series exhibits pronounced cyclical dynamics, with

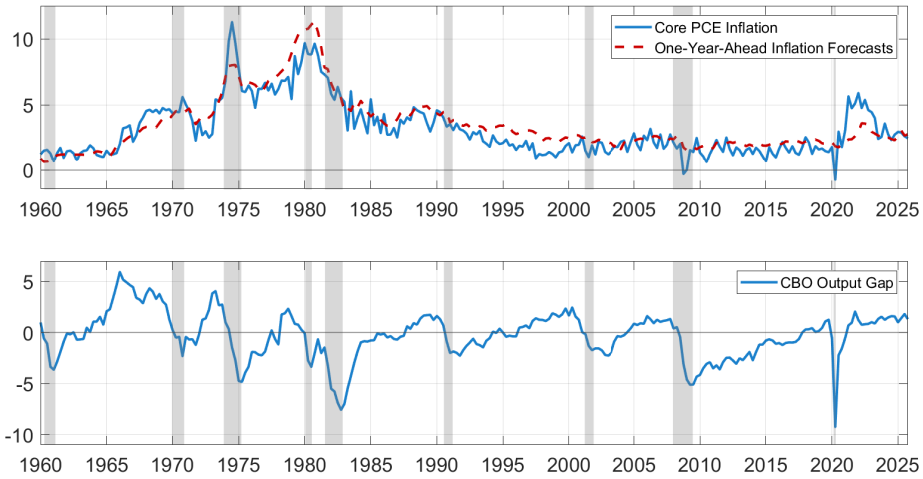


Figure 1: US core PCE inflation and one-year-ahead CPI inflation expectations from the Survey of Professional Forecasters (SPF) (upper panel), and the CBO output gap (lower panel), 1960Q1–2025Q4. Shaded areas indicate recessions as identified by the National Bureau of Economic Research (NBER).

large negative gaps during major recessions, notably in the early 1980s, the Great Recession, and the COVID-19 shock in 2020. Overall, the figure highlights the joint evolution of inflation, expectations, and real economic slack, which motivates the empirical Phillips curve specification.

4 Empirical Results

This section presents the main empirical findings. Subsection 4.1 first examines the time variation in inflation expectation weights, followed by an analysis of the Phillips curve slope across the conditional inflation distribution in Subsection 4.2. Subsection 4.3 then investigates the historical evolution of the slope profile, and finally, Subsection 4.4 examines the dynamics of conditional inflation quantiles and associated inflation risks. Unless otherwise specified, the results are based on the sDMQ-TV specification using the CBO output gap and one-year-ahead SPF inflation expectations. Results for the alternative specifications and model fit comparisons are reported in Appendix C.

4.1 Expectation Weights Dynamics

We begin the empirical discussion by examining the evolution of the relative weight of the backward- and forward-looking terms, as captured by the time-varying weighting process described previously. Our baseline specification uses the one-year-ahead inflation expectation from the Survey of Professional Forecasters (SPF) as a proxy for the forward-looking component. Figure 2 displays the estimated relative weight on inflation expectations, $1-\beta_t$, together with its four-quarter moving average for better visualization. Higher values of $1-\beta_t$ indicate a greater direct con-

tribution of survey-based expectations relative to the inflation benchmark, which captures inflation inertia and backward-looking information incorporated in earlier periods.

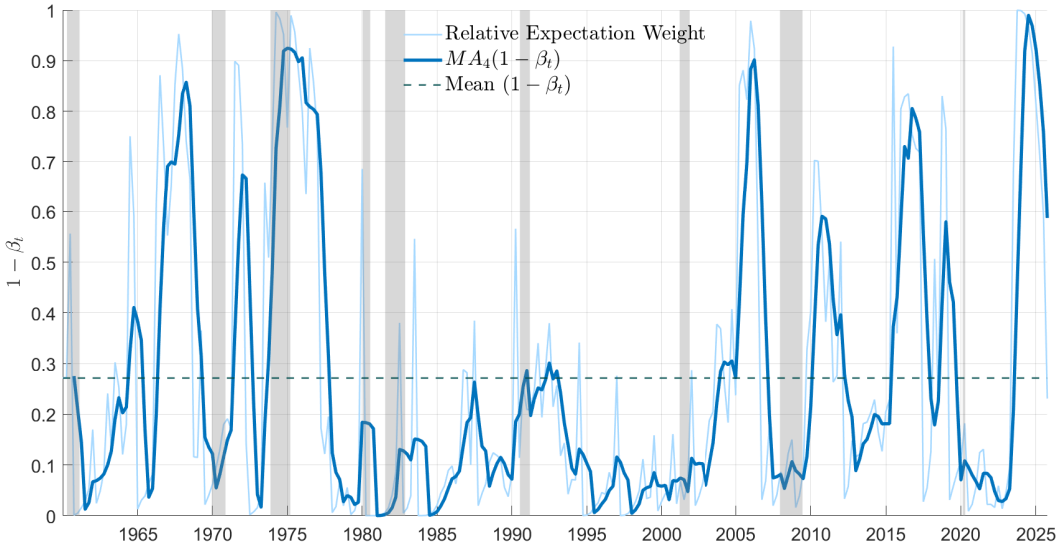


Figure 2: Estimated relative weight on survey-based inflation expectations, $1-\beta_t$ and its four-quarter moving average. Shaded areas indicate recessions as identified by the National Bureau of Economic Research (NBER).

The figure reveals substantial time variation in the relative importance of the two components, consistent with the moderate persistence of the updating process, as reflected in the estimated AR parameter $\phi^{\beta} = 0.74$. This suggests that inflation dynamics can be dominated by either the forward-looking or backward-looking component. Notably, in several periods β_t is estimated to be close to unity, consistent with a temporarily dominant role for inflation inertia and highly persistent inflation dynamics.

The weight on inflation expectations reaches relatively high levels during parts of the 1960s and the mid-1970s, suggesting increased role for explicit inflation expectations during the Great Inflation. It subsequently declines sharply, reaching low levels around the Volcker disinflation and remaining generally subdued throughout the 1980s and 1990s. The increased reliance on the inflation benchmark may reflect the reduced macroeconomic volatility and more predictable economic outcomes observed during the Great Moderation. This interpretation is consistent with Cornea-Madeira et al. (2019), who attribute the increased reliance on backward-looking forecasting rules during this period to their improved relative predictive performance, drawing on evidence from Stock and Watson (2007) and Madeira and Zafar (2015).

In more recent decades, the estimated weight exhibits several pronounced but relatively short-lived increases, notably around the mid-2000s, the mid-2010s, and toward the end of the sample. Importantly, the substantial increase in reliance on inflation expectations around 2024Q4–2025Q1 coincides with heightened geopolitical uncertainty and renewed volatility in energy markets, including concerns about

potential oil supply disruptions.

As discussed, our framework is conceptually related to Cornea-Madeira et al. (2019), who estimate a behavioral Phillips curve in which the relative importance of forward-looking (*fundamentalist*) and backward-looking (*random-walk*) forecasting rules evolves with their recent predictive performance. A comparison of the time profiles of the estimated weights (see their Figures 2 and 4) reveals notable qualitative similarities. Both sets of results exhibit substantial time variation, with a relatively limited role for the forward-looking component during much of the late 1990s and a marked increase around the mid-2000s. These similarities are noteworthy given the different modeling approaches. While their behaviorally microfounded model attributes changes in expectation weights to switching between forecasting rules based on relative predictive performance, our framework generates comparable dynamics through an observation-driven updating mechanism informed by the conditional inflation distribution, without explicitly modeling heterogeneous agents or their forecasting decisions.

To further examine the evolution of expectation weights, Figure 3 compares their four-quarter moving average with two measures of inflation instability and distributional variation: the four-quarter moving average of absolute changes in inflation, $|\Delta\pi_t|$, and the upper-tail dispersion of the conditional inflation distribution, $q_t^{0.95} - q_t^{0.50}$. The first captures recent fluctuations in realized inflation, while the second measures the distance between the conditional median and the upper tail.

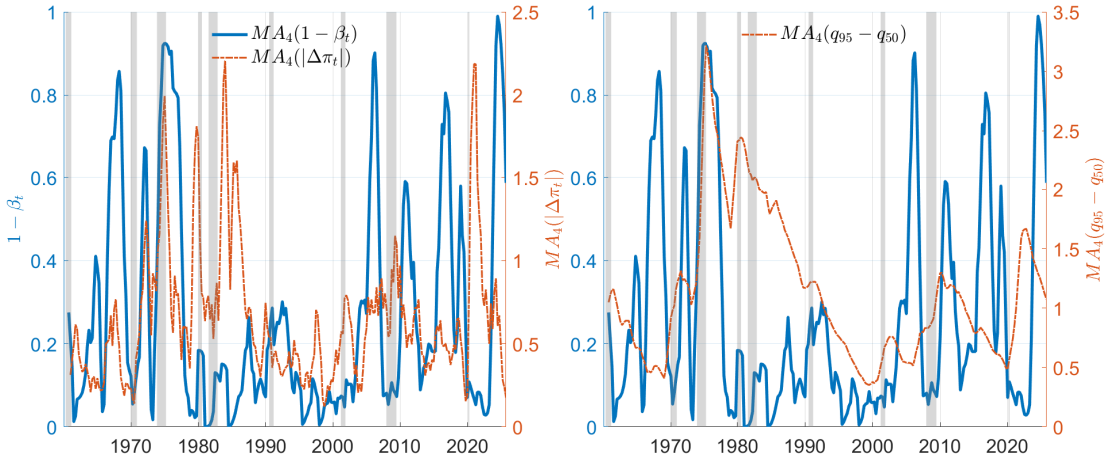


Figure 3: Four-quarter moving average of the estimated relative weight on survey-based inflation expectations, $1-\beta_t$, compared with the absolute quarterly change in inflation (left panel) and upper-tail dispersion (right panel), measured as the difference between the conditional 95th and 50th quantiles. Shaded areas indicate recessions as identified by the National Bureau of Economic Research (NBER).

The figure indicates that periods of relatively high expectations weights often coincide with elevated inflation instability and greater upper-tail dispersion, particularly during the 1970s. Notably, the more stable inflationary environment and lower upper-tail dispersion during the Great Moderation (Stock and Watson, 2002)

are consistent with a greater reliance on backward-looking pricing behavior. However, the onset of the GFC presents a contrasting episode, as heightened inflation instability and upper-tail dispersion were accompanied by a sharp increase in the weight on forward-looking inflation expectations. One possible interpretation is that heightened attention to inflation news strengthens the transmission of professional forecasts to broader expectations, consistent with the information-diffusion mechanism proposed by Carroll (2003).

In summary, these findings are consistent with the structure of the proposed updating mechanism, in which adjustments to the relative expectation weight depend jointly on the distributional inflation surprise and the divergence between survey expectations and the endogenous inflation benchmark. Overall, the results highlight substantial variation in the relative importance of survey-based inflation expectations and the recursively evolving inflation benchmark. Importantly, the results suggest that periods of heightened inflation instability and greater upper-tail dispersion are often accompanied by a stronger reliance on forward-looking inflation expectations.

4.2 The Shape of the Phillips Curve

We next examine the shape of the Phillips curve across the conditional inflation distribution. Figure 4 presents the estimated quantile-specific Phillips curve slopes under an sDMQ specification in which the slope parameters are held time-invariant. Besides the baseline specification with SPF survey expectations, we consider three alternative specifications combining different proxies for real marginal cost and an alternative measure of inflation expectations. Specifically, we estimate the static sDMQ models with two alternative measures of aggregate economic activity: the unemployment-vacancy ratio and the negative unemployment gap, both with SPF expectations. The third specification combines the CBO output gap with Cleveland Fed inflation expectations.

A common feature across all four estimations is the pronounced heterogeneity of the Phillips curve slope across quantiles. The estimated slopes exhibit a broadly hump-shaped pattern, reaching their highest values around the first conditional quartile before declining toward the upper tail of the inflation distribution. This decline is especially pronounced beyond the 75th quantile, with the estimated slopes approaching zero at the upper tail in some specifications. These findings suggest that the sensitivity of conditional inflation outcomes to economic activity is considerably weaker in the upper tail than in the lower and central parts of the distribution.

Although the magnitude of the estimated slopes differs across specifications, their qualitative shapes are remarkably similar. The hump-shaped pattern is preserved when replacing SPF expectations with Cleveland Fed expectations and when employing alternative measures of economic activity. This consistency suggests

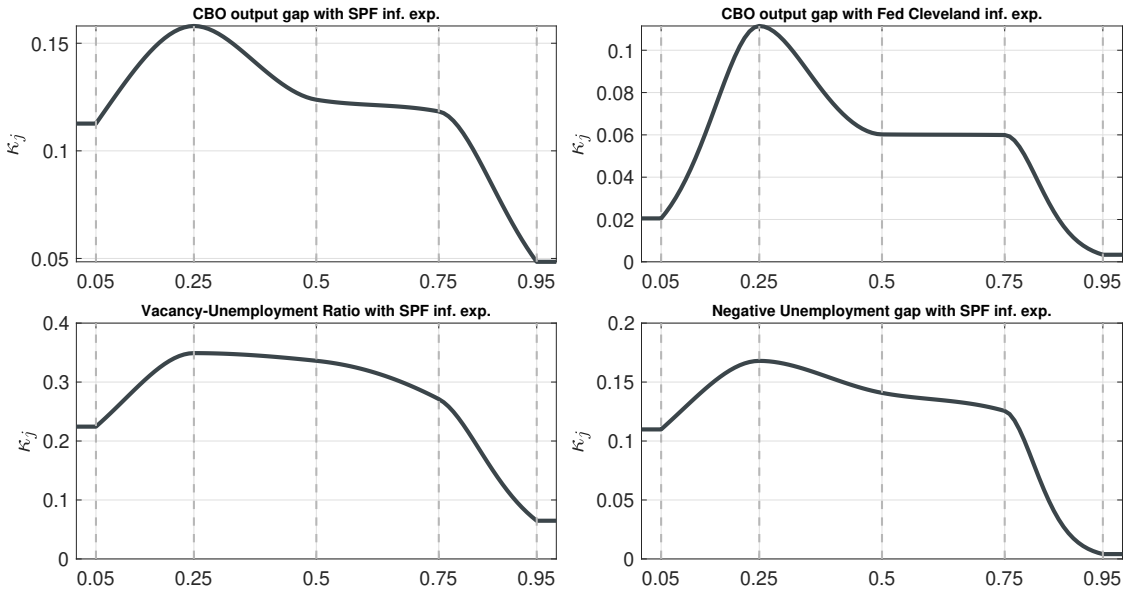


Figure 4: Estimated quantile-specific Phillips curve slopes under four sDMQ specifications with time-invariant slope parameters. The specifications combine the CBO output gap with SPF inflation expectations (top left), the CBO output gap with Cleveland Fed inflation expectations (top right), the vacancy-to-unemployment ratio with SPF expectations (bottom left), and the negative unemployment gap with SPF expectations (bottom right). Vertical dashed lines indicate the target probability levels used in estimation.

that the estimated asymmetry across the inflation distribution is not specific to a particular choice of explanatory variables. Overall, the results indicate that the Phillips curve cannot be adequately characterized by a single homogeneous slope across inflation quantiles, demonstrating the importance of allowing its shape to vary across the conditional distribution.

4.3 Historical Evolution of the Shape of the Phillips Curve

This section analyzes the evolution of the time-varying Phillips curve slope across quantiles of the conditional inflation distribution. By allowing the inflation–output gap trade-off to vary over time and across probability levels, we assess the extent of time variation and distributional asymmetry in the slope dynamics.

Figure 5 displays the time-varying Phillips curve slope estimates obtained from the sDMQ-TV specification with short-run inflation expectations, revealing pronounced heterogeneity across the conditional inflation distribution. For better visualization of the dynamics of the overall slope profile, Figure 7 in Appendix D presents the complete slope shape, including the interpolated auxiliary parameters. The slope is generally steepest and exhibits the greatest variability at the first quartile, remains moderate at the median, and declines toward higher probability levels, becoming flatter and relatively stable in the upper tail.

The results indicate an asymmetric Phillips curve across the conditional inflation distribution, with time-varying patterns that coincide with major macroeco-

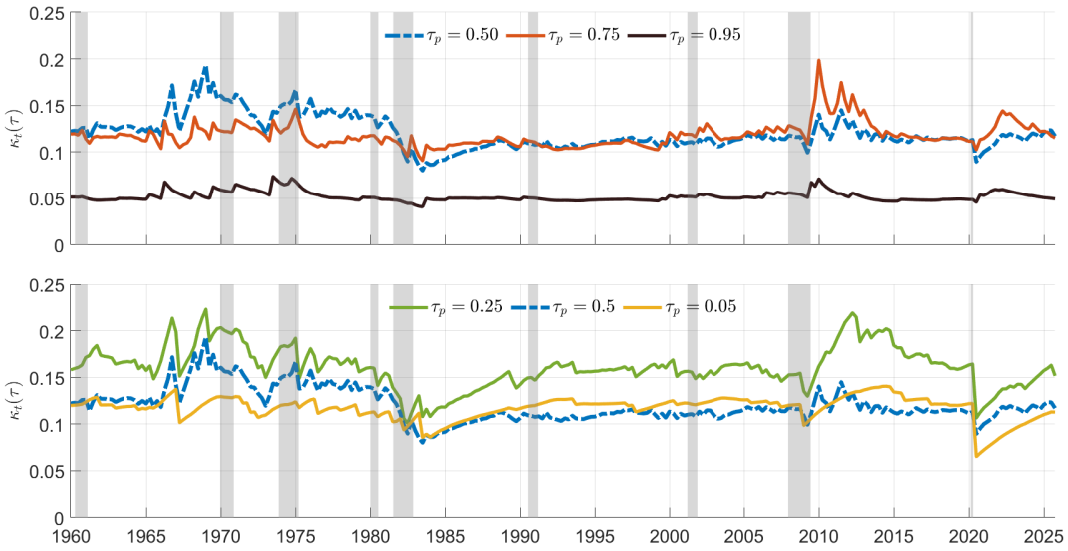


Figure 5: Time-varying Phillips curve slope estimates across the conditional inflation distribution, obtained from the sDMQ-TV specification with short-run inflation expectations. The upper panel displays the slopes at the median, 75th, and 95th quantiles, while the lower panel displays the slopes at the 5th, 25th, and 50th quantiles. Shaded areas indicate recessions as identified by the National Bureau of Economic Research (NBER).

nomic episodes in the US economy. Importantly, this time variation alters the slope profile substantially in certain periods.

During the late 1960s and 1970s, the slopes at the first quartile and median exhibit relatively high levels and pronounced fluctuations. The general steepening of the slope profile during the Great Inflation is consistent with the findings of Demirel and Wilson (2026), who show that supply disruptions and low unemployment can amplify the inflationary effects of increased aggregate demand. The proposed updating mechanism can also capture such dynamics. During episodes of upward inflation surprises, positive innovations in $-z_t$, potentially reflecting, at least in part, inflationary pressures from oil shocks, when combined with a large positive output gap, can lead to upward adjustments in the Phillips curve slope, contributing to its steepening.

In contrast, during the early 1980s, coinciding with the Volcker disinflation, the estimated slopes decline sharply across the lower and central quantiles, indicating a weakening of the inflation–output gap trade-off. Despite this initial decline, the estimated slopes remain relatively stable thereafter, with only modest further flattening. These findings are in line with Hazell et al. (2022), who estimate the Phillips curve using state-level price indices constructed from BLS microdata and find only modest flattening since the 1980s. Their results also suggest that the Volcker disinflation was primarily driven by declining long-run inflation expectations rather than a steep Phillips curve.

Another notable change in the slope profile occurs in the post-GFC period. Coibion and Gorodnichenko (2015) show that the relatively limited disinflation following the GFC can be explained by rising household inflation expectations between

2009 and 2011, without requiring substantial changes in the Phillips curve slope. Our findings provide a complementary perspective, as both SPF expectations and the recursively evolving inflation benchmark, which incorporates information on past inflation dynamics, remained well above realized inflation during this period. This suggests a potential role for persistent inflation expectations in explaining the missing disinflation following the Great Recession, despite substantial economic slack, without requiring a downward adjustment in the Phillips curve slope.

An even more remarkable coincidence is the overlap between the Federal Reserve's quantitative easing (QE) programs and the sharp steepening of the Phillips curve at the first quartile during the early 2010s. With conventional monetary policy constrained by the effective lower bound, large-scale asset purchases were intended to stimulate aggregate demand and support inflation expectations (Krishnamurthy and Vissing-Jorgensen, 2011; Chung et al., 2012). The estimated steepening suggests that lower-quantile inflation outcomes became more sensitive to economic activity during this period. Although this coincidence points to a potential connection between unconventional monetary policy and changes in the Phillips curve slope profile, this relationship warrants further investigation.

Lastly, at the onset of the COVID-19 pandemic, the Phillips curve relationship weakens markedly across the lower and central quantiles of the inflation distribution, with particularly pronounced declines at the first quartile and lower tail. This temporary flattening suggests a substantial reduction in the sensitivity of conditional inflation outcomes to economic slack during a period of exceptional macroeconomic disruption.

The temporary steepening at $\tau_p = 0.75$ over 2022Q1–2024Q4 coincides with a period of exceptionally tight labor markets and disruptions in global oil markets, especially in 2022, when the vacancy-to-unemployment ratio peaked and oil prices surged following Russia's invasion of Ukraine. This finding is qualitatively consistent with Benigno and Eggertsson (2023), who derive and estimate a nonlinear Phillips curve characterized by a kink at a vacancy-to-unemployment ratio near unity, beyond which labor-market tightness and supply shocks have substantially stronger effects on inflation.

Overall, our findings relate to Jorgensen and Lansing (2019), who show that improved anchoring of inflation expectations can explain the flattening of the accelerationist Phillips curve (changes in inflation) and the steepening of the original Phillips curve (level of inflation) without changing the underlying structural slope. Our results offer a complementary perspective: substantial variation in the relative weight on survey-based expectations coexists with a relatively stable median Phillips curve slope over much of the post-1980s period. At the same time, our estimates reveal temporary slope adjustments that differ across the conditional inflation distribution, highlighting an additional dimension of heterogeneity not captured by conventional conditional-mean Phillips curve specifications.

4.4 Distributional Inflation Dynamics

We close the empirical discussion by reporting the dynamic conditional inflation quantiles. By tracking the evolution of selected target quantiles, we analyze time variation in the conditional distribution of US inflation. Focusing on distributional dynamics rather than mean outcomes allows us to identify shifts in asymmetry, dispersion, and tail behavior across macroeconomic regimes.

The filtered target inflation quantiles are reported in Figure 6. A clear pattern emerges in the shape of the distribution when comparing the five target quantiles over time. These distributional shifts are consistent with well-documented episodes in US macroeconomic history. For further insight, Figure 8 in Appendix D presents the implied conditional moments obtained from fitted skewed t-distributions.

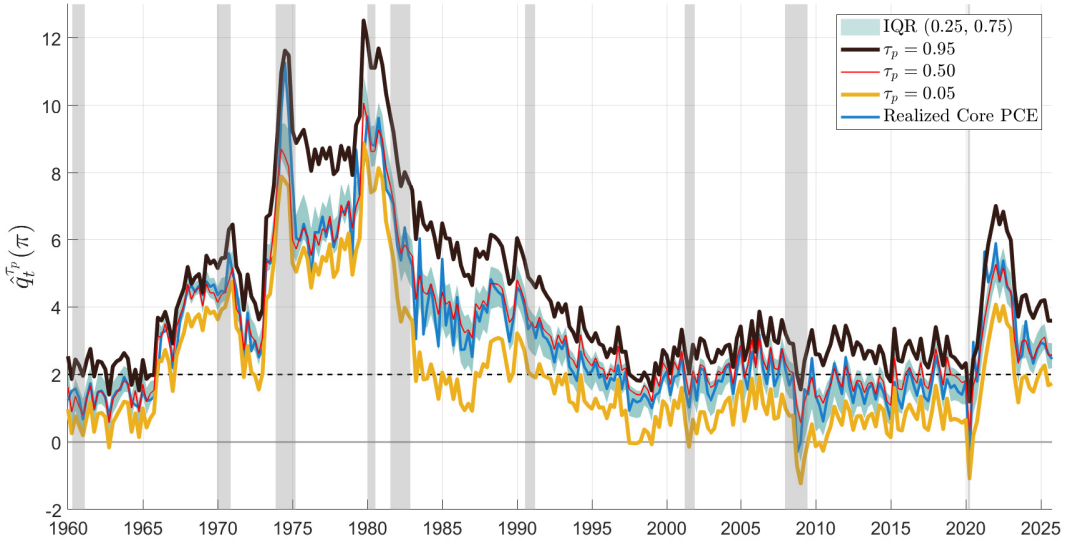


Figure 6: Filtered conditional inflation quantiles with realized core PCE inflation. The horizontal dashed line represents the Fed's 2% PCE inflation target, shown for reference. Vertical shaded areas indicate recessions as identified by the National Bureau of Economic Research (NBER).

During the oil shocks of the 1970s, the right tail of the inflation distribution moved markedly away from the median, consistent with elevated inflationary risk. The right-tail asymmetry remained substantial for an extended period, consistent with persistent inflation and less firmly anchored inflation expectations, alongside a relatively weak sensitivity to output gap fluctuations above the central quantiles of the inflation distribution.

After 1985, during the Great Moderation, the inflation quantiles became more tightly clustered, reflecting lower inflation variability during this interval. However, this period of relative stability was disrupted by the GFC. The widening spread between quantiles around the median indicates greater conditional dispersion during this period.

More recently, a clear contrast emerges between the pre- and post-GFC decades. While realized core PCE inflation after the GFC was only marginally lower than in the preceding period, the lower-tail quantiles at $\tau = 0.05$ periodically became

negative. This indicates that, despite relatively subdued realized inflation and the Fed's 2% PCE inflation target, deflationary risks remained non-negligible in the post-GFC environment, consistent with concerns raised by Fuhrer et al. (2012).

The severe macroeconomic shocks associated with the COVID-19 pandemic marked a turning point in inflation dynamics. After a brief period of downward pressure on prices, the subsequent inflation surge ended a period characterized by relatively low inflationary risk. Importantly, the right-tail quantiles remained close to 4% at the end of the analyzed period, suggesting that upside inflation risk remain a concern. The persistent upper tail is consistent with the relatively weak Phillips curve relationship estimated for the upper tail, suggesting that monetary policy may have limited effectiveness in reducing inflationary risk through the conventional aggregate demand channel during high-inflation regimes, potentially increasing the importance of anchoring inflation expectations.

5 Conclusions

This paper analyses the U.S. inflation dynamics by jointly modeling multiple quantiles within a hybrid New Keynesian Phillips curve framework. To this end, we introduce a smoothed dynamic multiple quantile model that allows quantile-specific, time-varying parameters. The framework uses an observation-driven mechanism to govern the relative weight on survey-based inflation expectations and a recursively evolving inflation benchmark, while allowing the Phillips curve slope to vary across quantiles and over time. By linking the time variation in a reduced set of quantiles to economic theory and to the full conditional distribution within a score-driven framework, the proposed specification captures heterogeneous responses of inflation to aggregate activity across the distribution and over time.

The empirical results reveal pronounced distributional and time-varying nonlinearities in the inflation–output gap trade-off. Specifically, the Phillips curve exhibits an asymmetric profile across quantiles, with relatively steeper slopes at lower and central probability levels and a marked flattening towards the upper tail. This pattern is broadly robust across alternative measures of economic activity and survey expectations. Allowing for time variation reveals a decline in the slope profile around the Volcker disinflation, followed by relatively modest longer-run flattening and substantial stability over subsequent decades, particularly around the median. Nevertheless, temporary slope adjustments remain evident, including a pronounced steepening at the first quartile during the early 2010s and a more moderate increase at the third quartile during the recent inflationary episode. In contrast, the slope at the upper tail remains comparatively low and stable.

The results also highlight substantial variation over time in the relative importance of survey-based inflation expectations and the inflation benchmark process. The estimated expectation weights fluctuate sharply across macroeconomic regimes, suggesting that the relative influence of these two components on inflation

dynamics changes substantially over time. After the GFC, both survey expectations and the recursively evolving inflation benchmark remained above realized inflation, suggesting a possible role for these components in explaining the missing disinflation despite marked economic slack. More broadly, these findings indicate that changes in expectation dynamics may coexist with relatively stable Phillips curve slopes, while temporary slope adjustments differ across the conditional inflation distribution.

The filtered conditional inflation quantiles deliver further insights into the asymmetric evolution of inflationary risks. During the Great Inflation, the pronounced widening of the upper tail reflects elevated inflationary risk and strong persistence. Following the Great Moderation, the conditional distribution became more concentrated, although negative lower-tail quantiles periodically emerged after the GFC, indicating persistent deflationary risks. The COVID-19 pandemic and the subsequent inflation surge marked another substantial shift in the conditional distribution. Notably, the upper-tail inflation quantile remained close to 4% at the end of the analyzed period, suggesting that upside inflation risk remained elevated. The relatively weak Phillips curve relationship in the upper tail further suggests that monetary policy may have limited effectiveness at reducing inflation through the conventional aggregate demand channel during high-inflation regimes, possibly increasing the importance of anchoring inflation expectations.

Overall, the results show that exploiting distributional information, time-varying expectation weights, and quantile-specific Phillips curve slopes provides additional insights into inflation dynamics beyond those available from conventional conditional-mean specifications and limited-information estimation methods.

Appendix

A Baseline DMQ framework

This appendix summarizes the Dynamic Multiple Quantile (DMQ) framework of Catania and Luati (2023) that underlies the specification developed in the main text. The baseline DMQ model jointly describes the dynamics of a predefined set of conditional quantiles through an observation-driven updating mechanism. Three features are especially relevant for the present application. First, the model satisfies the non-crossing quantile criterion by construction. Second, it is consistent with the defining property of fixed quantiles, such that, asymptotically, the proportion of observations falling below the unconditional τ-level quantile converges to τ. Third, the joint modeling of a relatively large set of quantiles provides the distributional information required for the score-driven updating mechanism.

Let $\tau_1 < \cdots < \tau_J$ denote the probability levels of the modeled quantiles and let τ_{j^*} denote the reference probability level. The baseline DMQ recursion is defined as

$$\tilde{q}_t^{\tau_j} = \begin{cases} \tilde{q}_t^{\tau_{j+1}} - \eta_{j,t}, & \text{if } \tau_j < \tau_{j^*}, \\ \tilde{q}_t^{\tau_{j^*}}, & \text{if } \tau_j = \tau_{j^*}, \\ \tilde{q}_t^{\tau_{j-1}} + \eta_{j,t}, & \text{if } \tau_j > \tau_{j^*}, \end{cases} \tag{A.1}$$

where the reference quantile evolves according to

$$\tilde{q}_t^{\tau_{j^*}} = (1-\beta)\tilde{q}_0^{\tau_{j^*}} + \alpha u_{t-1}^{\tau_{j^*}} + \beta \tilde{q}_{t-1}^{\tau_{j^*}}. \tag{A.2}$$

Here, $\tilde{q}_0^{\tau_{j^*}}$ denotes the unconditional quantile at the reference probability level, β controls persistence, and α determines the loading on the corresponding forcing variable.

The remaining quantiles are constructed relative to the reference process through the positive spacing variables

$$\eta_{j,t} = \exp(\xi_{j,t}), \tag{A.3}$$

where

$$\xi_{j,t} = (1-\phi)\xi_{j,0} + \gamma u_{t-1}^{\tau_j} + \phi \xi_{j,t-1}. \tag{A.4}$$

The exponential transformation guarantees positive spacings between adjacent quantiles and therefore preserves their ordering. The parameters ϕ and γ control the persistence of the spacing process and its response to the forcing variable, respectively.

The forcing variables are constructed from quantile-hit information. For the

baseline DMQ specification, they take the form

$$u_t^{\tau_j} = \begin{cases} b_j a_j^{-1} \sum_{i=1}^{j} z_{i,t}, & \text{if } j < j^*, \\ a_j^{-1} \sum_{i=1}^{J} z_{i,t}, & \text{if } j = j^*, \\ b_j a_j^{-1} \sum_{i=j}^{J} z_{i,t}, & \text{if } j > j^*, \end{cases} \tag{A.5}$$

where

$$b_j = \mathbb{1}(\tau_j < \tau_{j^*}) - \mathbb{1}(\tau_j > \tau_{j^*}),$$

and a_j is the normalization term defined in Catania and Luati (2023). The normalization plays a role analogous to information scaling in conventional score-driven models, although it is based on unconditional moments associated with the empirical quantiles rather than on a fully parametric likelihood.

For each probability level τ_j, define the quantile-hit variable as

$$\begin{aligned} z_{j,t} &= \frac{\partial}{\partial q_t^{\tau_j}} \rho_{\tau_j}\left(y_t - q_t^{\tau_j}\right) \\ &= \mathbb{1}\left(y_t \leq q_t^{\tau_j}\right) - \tau_j, \end{aligned} \tag{A.6}$$

where

$$\rho_\tau(\chi) = \chi\left[\tau - \mathbb{1}(\chi < 0)\right]$$

is the quantile check function. Collecting the individual hits gives

$$z_t = (z_{1,t}, \ldots, z_{J,t})'.$$

The vector z_t therefore summarizes the location of the realized observation relative to the collection of modeled conditional quantiles. The forcing variables in (A.5) aggregate the relevant elements of this vector according to the position of each quantile relative to the reference probability level.

Under the regularity conditions established in Catania and Luati (2023), the quantile-hit and corresponding forcing-variable processes satisfy the moment properties required for observation-driven updating. The resulting DMQ specification can therefore be interpreted within the broader class of score-driven models developed by Creal et al. (2013) and Harvey (2013), while retaining the non-crossing structure required for the joint evolution of multiple conditional quantiles.

The framework developed in the main text preserves this baseline structure while modifying the economically relevant components of the model. Specifically, the reference-quantile dynamics are extended to allow endogenous reweighting between survey-based inflation expectations and the endogenous inflation benchmark, while selected semi-structural coefficients are parameterized directly only at a reduced set of target quantiles and recovered at the remaining auxiliary quantiles through smooth interpolation.

B Interpolation procedure and related derivatives

The complete set of quantile-specific log-slope parameters is defined as

$$\kappa_{j,t} = \begin{cases} \hat{\kappa}_{p,t}, & \tau_j = \tau_p, \\ \bar{\kappa}_{j,t}, & \text{otherwise.} \end{cases} \tag{B.1}$$

The auxiliary parameters $\bar{\kappa}_{j,t}$ are obtained through shape-preserving piecewise cubic Hermite interpolation (PCHIP), following Fritsch and Carlson (1980) and Fritsch and Butland (1984). For an auxiliary probability level τ_j between two adjacent target levels, let $p_L(j)$ and $p_U(j)$ denote the corresponding lower and upper target indices, and define

$$\begin{aligned} H_j &= \tau_{p_U(j)} - \tau_{p_L(j)}, \\ b_j &= \frac{\tau_j - \tau_{p_L(j)}}{H_j}. \end{aligned}$$

The interpolated parameters are given by

$$\begin{aligned} \bar{\kappa}_{j,t} = {} & h_{00}(b_j)\hat{\kappa}_{p_L(j),t} + H_j h_{10}(b_j) m_{p_L(j),t} \\ & + h_{01}(b_j)\hat{\kappa}_{p_U(j),t} + H_j h_{11}(b_j) m_{p_U(j),t}, \end{aligned} \tag{B.2}$$

where $m_{p,t}$ denotes the derivative of the interpolated parameter curve with respect to the probability level at target quantile p. The cubic Hermite basis functions are

$$\begin{aligned} h_{00}(b) &= 2b^3 - 3b^2 + 1, & h_{10}(b) &= b^3 - 2b^2 + b, \\ h_{01}(b) &= -2b^3 + 3b^2, & h_{11}(b) &= b^3 - b^2. \end{aligned} \tag{B.3}$$

The knot derivatives $m_{p,t}$ are determined by the PCHIP shape-preserving rules. Let P denote the number of target probability levels, and define the adjacent knot spacings and secant slopes as

$$\begin{aligned} h_p &= \tau_{p+1} - \tau_p, \\ d_{p,t} &= \frac{\hat{\kappa}_{p+1,t} - \hat{\kappa}_{p,t}}{h_p}, \qquad p = 1, \ldots, P-1. \end{aligned} \tag{B.4}$$

For interior knots $p = 2, \ldots, P-1$, the derivatives are given by

$$m_{p,t} = \begin{cases} \dfrac{w_{p,L} + w_{p,R}}{w_{p,L}/d_{p-1,t} + w_{p,R}/d_{p,t}}, & d_{p-1,t} d_{p,t} > 0, \\ 0, & \text{otherwise,} \end{cases} \tag{B.5}$$

where

$$\begin{aligned} w_{p,L} &= 2h_p + h_{p-1}, \\ w_{p,R} &= h_p + 2h_{p-1}. \end{aligned}$$

At the boundaries, the initial one-sided derivatives are

$$\tilde{m}_{1,t} = \frac{(2h_1 + h_2)d_{1,t} - h_1 d_{2,t}}{h_1 + h_2},$$
$$\tilde{m}_{P,t} = \frac{(2h_{P-1} + h_{P-2})d_{P-1,t} - h_{P-1} d_{P-2,t}}{h_{P-1} + h_{P-2}}. \tag{B.6}$$

These derivatives are adjusted to preserve the shape of the interpolated curve. For either boundary $e \in \{1, P\}$, let (a_e, b_e) equal $(d_{1,t}, d_{2,t})$ at the lower boundary and $(d_{P-1,t}, d_{P-2,t})$ at the upper boundary. The adjusted derivative is

$$m_{e,t} = \begin{cases} 0, & \operatorname{sgn}(\tilde{m}_{e,t}) \neq \operatorname{sgn}(a_e), \\ 3a_e, & \operatorname{sgn}(a_e) \neq \operatorname{sgn}(b_e) \text{ and } |\tilde{m}_{e,t}| > 3|a_e|, \\ \tilde{m}_{e,t}, & \text{otherwise.} \end{cases} \tag{B.7}$$

For probability levels below $\tau = 0.05$ and above $\tau = 0.95$, constant extrapolation is imposed using the nearest target parameter.

The derivative required for constructing the forcing variables accounts for both the direct dependence on target parameters and their indirect effects through interpolation. For an auxiliary probability level between adjacent target quantiles, it is given by

$$\frac{\partial \kappa_{j,t}}{\partial \hat{\kappa}_{p,t}} = h_{00}(b_j)\mathbb{1}\{p = p_L(j)\} + h_{01}(b_j)\mathbb{1}\{p = p_U(j)\} + H_j h_{10}(b_j)\frac{\partial m_{p_L(j),t}}{\partial \hat{\kappa}_{p,t}} + H_j h_{11}(b_j)\frac{\partial m_{p_U(j),t}}{\partial \hat{\kappa}_{p,t}}. \tag{B.8}$$

The derivatives of the secant slopes with respect to the target parameters are

$$\frac{\partial d_{k,t}}{\partial \hat{\kappa}_{p,t}} = \frac{\mathbb{1}\{p = k+1\} - \mathbb{1}\{p = k\}}{h_k}. \tag{B.9}$$

For interior knots satisfying $d_{k-1,t} d_{k,t} > 0$, define

$$D_{k,t} = w_{k,L} d_{k,t} + w_{k,R} d_{k-1,t}.$$

Differentiating the weighted harmonic mean gives

$$\frac{\partial m_{k,t}}{\partial \hat{\kappa}_{p,t}} = \frac{(w_{k,L} + w_{k,R}) w_{k,L} d_{k,t}^2}{D_{k,t}^2}\frac{\partial d_{k-1,t}}{\partial \hat{\kappa}_{p,t}} + \frac{(w_{k,L} + w_{k,R}) w_{k,R} d_{k-1,t}^2}{D_{k,t}^2}\frac{\partial d_{k,t}}{\partial \hat{\kappa}_{p,t}}. \tag{B.10}$$

The derivatives are evaluated wherever the PCHIP interpolation mapping is differentiable. At points where the shape-preserving rules switch branches, the corresponding branchwise derivatives are used.

For the boundary knots, differentiation of the initial one-sided expressions yields

$$\frac{\partial \tilde{m}_{1,t}}{\partial \hat{\kappa}_{p,t}} = \frac{2h_1 + h_2}{h_1 + h_2}\frac{\partial d_{1,t}}{\partial \hat{\kappa}_{p,t}} - \frac{h_1}{h_1 + h_2}\frac{\partial d_{2,t}}{\partial \hat{\kappa}_{p,t}},$$
$$\frac{\partial \tilde{m}_{P,t}}{\partial \hat{\kappa}_{p,t}} = \frac{2h_{P-1} + h_{P-2}}{h_{P-1} + h_{P-2}}\frac{\partial d_{P-1,t}}{\partial \hat{\kappa}_{p,t}} - \frac{h_{P-1}}{h_{P-1} + h_{P-2}}\frac{\partial d_{P-2,t}}{\partial \hat{\kappa}_{p,t}}. \tag{B.11}$$

The derivatives of the adjusted boundary slopes follow the active branch of the shape-preserving rule: they are zero when the slope is set to zero, three times the derivative of a_e when the slope is capped at $3a_e$, and equal to the derivative of $\tilde{m}_{e,t}$ otherwise.

At the target probability levels, the Jacobian equals one with respect to the corresponding target parameter and zero otherwise. The same rule applies to the nearest target parameter under constant extrapolation outside the target range. At points where the PCHIP rules switch branches, the interpolation mapping need not be differentiable; the expressions above apply wherever the corresponding derivatives exist.

C Specifications and model fit

Table 1: Model specifications

Model	Description
DMQ(1)	Benchmark DMQ model of Catania and Luati (2023)
DMQ(2)	DMQ model with static and quantile-invariant parameters
sDMQ	sDMQ model with static and quantile-specific PC slope
sDMQ-TV	sDMQ model with time-varying and quantile-specific PC slope and time-varying inflation expectation weights

Table 2: Parameter estimates across model specifications with CBO output gap and SPF one-year-ahead inflation expectations (1960Q1–2025Q4)

	DMQ(1)	**DMQ(2)**	**sDMQ**	**sDMQ-TV**
α	0.526	0.443	0.428	0.423
β	0.9999	0.732	0.727	0.729
γ	0.056	0.067	0.068	0.096
ϕ	0.998	0.998	0.998	0.997
ρ_κ	–	–	–	0.055
ϕ_κ	–	–	–	0.921
ρ_β	–	–	–	0.848
ϕ_β	–	–	–	0.743
$1-\beta$	0.0001	0.268	0.273	0.271
κ	–	0.120	0.112	0.114
$\kappa_{\tau=0.05}$	–	–	0.113	0.117
$\kappa_{\tau=0.25}$	–	–	0.158	0.161
$\kappa_{\tau=0.50}$	–	–	0.124	0.121
$\kappa_{\tau=0.75}$	–	–	0.118	0.118
$\kappa_{\tau=0.95}$	–	–	0.048	0.052
LogLik	5936.7	5701.6	5658.8	5366.7
AIC	11881.5	11413.1	11335.6	10747.4

Note: Under the sDMQ specification, the reported values of κ are averages across quantiles, while under sDMQ-TV they are averaged across both time and quantiles. Under sDMQ-TV, both κ_p and β correspond to sample averages. The values under LogLik correspond to the final negative quasi-maximum likelihood (QML) objective function evaluated at the estimated parameter values.

D Supplementary results

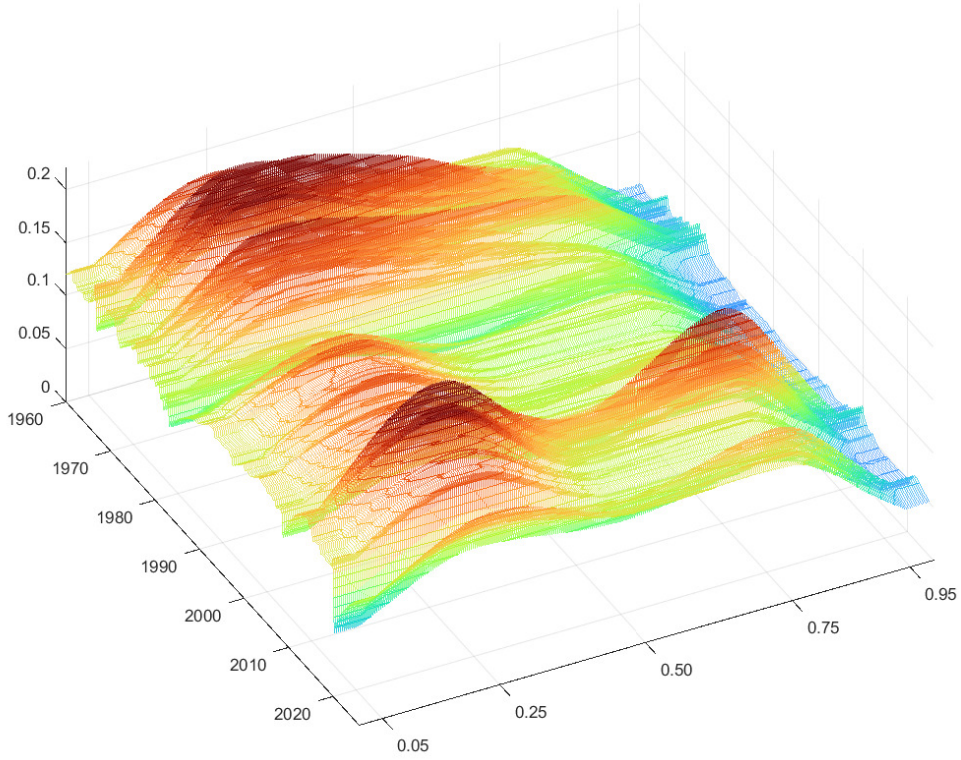


Figure 7: Three-dimensional visualization of the time-varying Phillips curve slope profile across the conditional inflation distribution, including the interpolated auxiliary parameters.

The skew t density takes the form of

$$f(y;\mu,\sigma,\zeta,\nu) = \frac{2}{\sigma}\, d_T\left(\frac{y-\mu}{\sigma};\nu\right)\, T\left(\zeta\,\frac{y-\mu}{\sigma}\sqrt{\frac{\nu+1}{\nu+\left(\frac{y-\mu}{\sigma}\right)^2}}\,;\,\nu+1\right), \qquad \text{(D.1)}$$

where $d_T(\cdot)$ and $T(\cdot)$ denote the probability density function (PDF) and CDF of the skew t distribution, respectively. The parameters are location (μ), scale (σ), degrees of freedom (ν) governing tail thickness, and shape (ζ) governing the skewness. Estimating these parameters at each time t supports the assessment of the tail behavior of inflation by capturing its time-varying asymmetry and kurtosis implied by the estimated target quantiles.

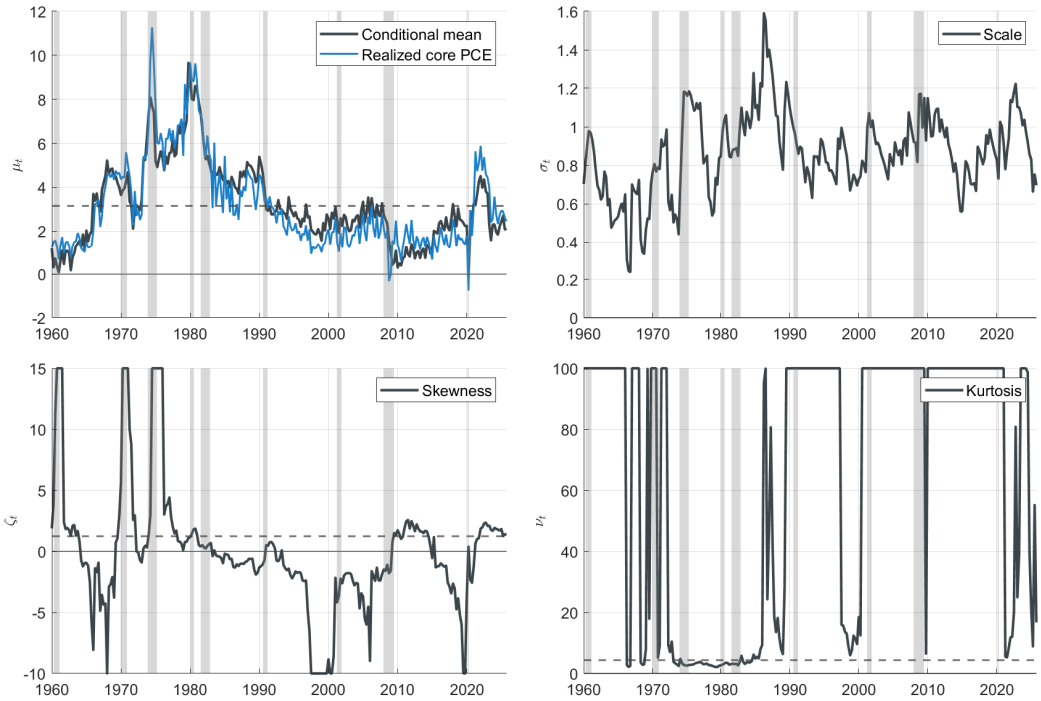


Figure 8: Filtered implied conditional moments (solid lines) obtained by fitting a skewed-t distribution, following Azzalini and Capitanio (2003), at each time t using the target quantiles. The parameters are the location (μ_t, upper left panel), the scale (σ_t, upper right panel), the shape parameter (ζ_t, lower left panel) governing skewness, and the degrees of freedom (ν_t, lower right panel) governing tail thickness. The horizontal dashed lines denote the corresponding unconditional moments. Shaded vertical areas indicate recessions as identified by the National Bureau of Economic Research (NBER).

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019, April). Vulnerable growth. *American Economic Review 109*(4), 1263–89. 13

Azzalini, A. and A. Capitanio (2003, May). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society Series B 65*(2), 367–389. 34

Benigno, P. and G. B. Eggertsson (2023). It's baaack: The surge in inflation in the 2020s and the return of the non-linear phillips curve. Working Paper 31197, National Bureau of Economic Research. 16, 23

Bils, M. and P. Klenow (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy 112*(5), 947–985. 10

Blasques, F., C. Francq, et al. (2023). Quasi score-driven models. *Journal of Econometrics 231*(2), 235–257. 4, 6, 12, 15

Bondell, H. D., B. J. Reich, and H. Wang (2010, 12). Noncrossing quantile regression curve estimation. *Biometrika 97*(4), 825–838. 14

Branch, W. A. and G. W. Evans (2006). Intrinsic heterogeneity in expectation formation. *Journal of Economic Theory 127*(1), 264–295. 6

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*(3), 383–398. 10

Carroll, C. D. (2003, 02). Macroeconomic expectations of households and professional forecasters. *The Quarterly Journal of Economics 118*(1), 269–298. 20

Catania, L. and A. Luati (2023). Semiparametric modeling of multiple quantiles. *Journal of Econometrics 237*(2, Part B), 105365. 3, 4, 5, 6, 14, 15, 27, 28, 32

Catania, L., A. Luati, and P. Vallarino (2026, 03). Economic vulnerability is state dependent. *The Econometrics Journal*. 12, 14

Chung, H., J.-P. Laforte, D. Reifschneider, and J. C. Williams (2012). Have we underestimated the likelihood and severity of zero lower bound events? *Journal of Money, Credit and Banking 44*(s1), 47–82. 23

Coibion, O. and Y. Gorodnichenko (2015, January). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics 7*(1), 197–232. 22

Cornea-Madeira, A., C. Hommes, and D. Massaro (2019). Behavioral heterogeneity in u.s. inflation dynamics. *Journal of Business & Economic Statistics 37*(2), 288–300. 2, 6, 8, 9, 18, 19

Cornea-Madeira, A. and J. Madeira (2022). Econometric analysis of switching expectations in uk inflation. *Oxford Bulletin of Economics and Statistics 84*(3), 651–673. 6

Cox, D. R. (1981). Statistical analysis of time series: Some recent developments. *Scandinavian Journal of Statistics 8*(2), 93–115. 4

Creal, D., S. J. Koopman, and A. Lucas (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics 28*(5), 777–795. 4, 28

Demirel, U. D. and M. Wilson (2026). Supply disruptions, economic slack, and the state-dependent phillips curve. *Economic Inquiry 64*(4), 1260–1280. 22

Dotsey, M., R. G. King, and A. L. Wolman (1999). State-dependent pricing and the general equilibrium dynamics of money and output. *The Quarterly Journal of Economics 114*(2), 655–690. 10

Fakoor, R., T. Kim, J. Mueller, A. J. Smola, and R. J. Tibshirani (2023). Flexible model aggregation for quantile regression. *Journal of Machine Learning Research 24*(162), 1–45. 13, 14

Forbes, K., J. Gagnon, and C. G. Collins (2021, October). Low inflation bends the phillips curve around the world. Working Paper 29323, National Bureau of Economic Research. 10

Fritsch, F. N. and J. Butland (1984). A method for constructing local monotone piecewise cubic interpolants. *SIAM Journal on Scientific and Statistical Computing 5*(2), 300–304. 29

Fritsch, F. N. and R. E. Carlson (1980). Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis 17*(2), 238–246. 29

Fuhrer, J. and G. Moore (1995). Inflation persistence. *The Quarterly Journal of Economics 110*(1), 127–159. 5

Fuhrer, J. C., G. P. Olivei, and G. M. B. Tootell (2012). Inflation dynamics when inflation is near zero. *Journal of Money, Credit and Banking 44*(s1), 83–122. 25

Galí, J. and M. Gertler (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics 44*(2), 195–222. 5, 6

Gertler, M. and J. Leahy (2008). A phillips curve with an ss foundation. *Journal of Political Economy 116*(3), 533–572. 10, 11

Golosov, M. and R. E. Lucas (2007). Menu costs and phillips curves. *Journal of Political Economy 115*(2), 171–199. 10

Gopinath, G. and O. Itskhoki (2010, June). In search of real rigidities. Working Paper 16065, National Bureau of Economic Research. 10

Harding, M., J. Lindé, and M. Trabandt (2022). Resolving the missing deflation puzzle. *Journal of Monetary Economics 126*, 15–34. 2, 10, 11

Harding, M., J. Lindé, and M. Trabandt (2023). Understanding post-covid inflation dynamics. *Journal of Monetary Economics 140*, S101–S118. Inflation: Drivers and Dynamics 2022. 10

Harvey, A. C. (2013). *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. Econometric Society Monographs. Cambridge University Press. 4, 28

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2022). The slope of the Phillips curve: Evidence from u.s. states. *The Quarterly Journal of Economics 137*(3), 1299–1344. 2, 22

Jorgensen, P. and K. Lansing (2019). Anchored inflation expectations and the slope of the phillips curve. Working Paper Series 2019-27, Federal Reserve Bank of San Francisco. 23

Koenker, R. and F. Machado, José A. (1999). Goodness of fit and related inference processes for quantile regression. *Journal of the American Statistical Association 94*(448), 1296–1310. 15

Komunjer, I. (2005). Quasi-maximum likelihood estimation for conditional quantiles. *Journal of Econometrics 128*(1), 137–164. 15

Kotz, S., J. Kozubowski, Tomasz and K. Podgórski (2001). *The Laplace Distribution and Generalizations*. Boston: Birkhäuser. 15

Krishnamurthy, A. and A. Vissing-Jorgensen (2011). The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity 42*(2), 215–287. 23

Kuttner, K. N. (1994). Estimating potential output as a latent variable. *Journal of Business & Economic Statistics 12*(3), 361–368. 2

Laubach, T. and J. C. Williams (2003). Measuring the natural rate of interest. *The Review of Economics and Statistics 85*(4), 1063–1070. 2

Lewis, K. F. and F. Vazquez-Grande (2019). Measuring the natural rate of interest: A note on transitory shocks. *Journal of Applied Econometrics 34*(3), 425–436. 2

Lindé, J. (2005). Estimating new-keynesian phillips curves: A full information maximum likelihood approach. *Journal of Monetary Economics 52*(6), 1135–1149. The econometrics of the New Keynesian price equation. 6

López-Salido, D. and F. Loria (2024). Inflation at risk. *Journal of Monetary Economics 145*, 103570. 2, 4, 13

Madeira, C. and B. Zafar (2015). Heterogeneous inflation expectations and learning. *Journal of Money, Credit and Banking 47*(5), 867–896. 18

Malmendier, U. and S. Nagel (2016, 02). Learning from inflation experiences. *The Quarterly Journal of Economics 131*(1), 53–87. 7

Mankiw, R. Reis, and J. Wolfers (2003). Disagreement about inflation expectations. *NBER Macroeconomics Annual*, 209–248. 6

Mankiw, N. G. and R. Reis (2002, 11). Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics 117*(4), 1295–1328. 6

Mavroeidis, S., M. Plagborg-Møller, and J. H. Stock (2014, March). Empirical evidence on inflation expectations in the new keynesian phillips curve. *Journal of Economic Literature 52*(1), 124–88. 2, 4, 5, 6, 7

Nakamura, E. and J. Steinsson (2008, 11). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics 123*(4), 1415–1464. 10

Pál, T. and G. Storti (2025). Estimating the r-star in the US: A score-driven state-space model with time-varying volatility persistence. Technical Report 125338, Munich Personal RePEc Archive (MPRA). Working Paper. 2

Phillips, A. W. (1958). The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861–1957. *Economica 25*(100), 283–299. 2

Poiraud-Casanova, S. and C. Thomas-Agnan (2000). About monotone regression quantiles. *Statistics & Probability Letters 48*(1), 101–104. 15

Reis, R. (2006, 07). Inattentive producers. *The Review of Economic Studies 73*(3), 793–821. 6

Rossi, B., A. Inoue, and Y. Wang (2024, Jun). Has the phillips curve flattened? French Stata Users' Group Meetings 2024 22, Stata Users Group. 2

Rudd, J. and K. Whelan (2005). New tests of the new-keynesian phillips curve. *Journal of Monetary Economics 52*(6), 1167–1181. The econometrics of the New Keynesian price equation. 5

Rudd, J. and K. Whelan (2006). Can rational expectations sticky-price models explain inflation dynamics? *The American Economic Review 96*(1), 303–320. 6

Sbordone, A. M. (2005). Do expected future marginal costs drive inflation dynamics? *Journal of Monetary Economics 52*(6), 1183–1197. The econometrics of the New Keynesian price equation. 6

Shen, G., Y. Jiao, Y. Lin, J. L. Horowitz, and J. Huang (2024). Nonparametric estimation of non-crossing quantile regression process with deep ReQU neural networks. *Journal of Machine Learning Research 25*(88), 1–75. 14

Smith, S. C., A. Timmermann, and J. H. Wright (2025). Breaks in the phillips curve: Evidence from panel data. *Journal of Applied Econometrics 40*(2), 131–148. 2

Staiger, D., J. H. Stock, and M. W. Watson (1997, March). The nairu, unemployment and monetary policy. *Journal of Economic Perspectives 11*(1), 33–49. 2

Stock, J. H. and M. W. Watson (2002). Has the business cycle changed and why? NBER Working Paper 9127, National Bureau of Economic Research. 19

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. 18

Tang, Q., Y. Gu, and B. Wang (2026). fastkqr: A fast algorithm for kernel quantile regression. *Journal of Computational and Graphical Statistics 35*(1), 395–405. 13, 14

White, H., T.-H. Kim, and S. Manganelli (2015). Var for var: Measuring tail dependence using multivariate regression quantiles. *Journal of Econometrics 187*(1), 169–188. 15

Xu, Q., X. Niu, C. Jiang, and X. Huang (2015). The phillips curve in the us: A nonlinear quantile regression approach. *Economic Modelling 49*, 186–197. 4