

Inflation at Risk

David López-Salido Francesca Loria

Federal Reserve Board

September 8, 2021

Abstract

We investigate how macroeconomic drivers influence the predictive inflation distribution and establish two key findings. First, the recent muted response of the conditional mean of inflation to economic conditions does not convey a complete picture of inflation dynamics. Indeed, we find ample variability in the tails of the inflation outlook that remains even when focusing on the most recent period of stable and low *mean* inflation. Second, we document that tight financial conditions carry substantial downside inflation risks in the United States and in the Euro Area, a feature overlooked by much of the literature but consistent with financial amplification mechanisms. Finally, we show that evidence from financial market quotes, from survey data and from a regime-switching model of inflation is consistent with our findings and use our model to track inflation risks during the Covid-19 crisis.

JEL CLASSIFICATION: C21, C53, E31, E44.

KEYWORDS: Inflation Risks, Quantile Regression.

Special thanks to Ben Bernanke, Danilo Cascaldi-Garcia, Steve Cecchetti, David Cho, Jim Clouse, Olivier Coibion, Deepa Datta, Giovanni Favara, Felix Galbis-Reig, Ed Herbst, Paul Lengermann, Antoine Lepetit, Ed Nelson, Claudia Pacella, Andrea Prestipino, Giorgio Primiceri, Jeremy Rudd, Tatevik Sekhposyan and Srečko Zimic as well as seminar participants of the Research & Statistics Lunch Seminar at the Federal Reserve Board, of the DG-E Research Seminar at the European Central Bank, of the 2020 CEBRA Annual Meeting Session on “Inflation Dynamics and the Phillips Curve”, the “Macro-at-Risk” Session at the ECB Working Group on Econometric Modelling, of the IMF Monetary and Capital Markets (MCM) Policy Forum, the 27th International Conference of Computing in Economics and Finance (CEF), the 2021 International Association for Applied Econometrics (IAAE) Annual Conference, the 2021 Meeting of the Society for Economic Dynamics (SED), and of the International Conference on Economic Modeling and Data Science (EcoMod2021) for their useful comments and suggestions to earlier versions of the paper. We also thank Eugenio Cerutti, Michiel De Pooter, Caitlin Dutta, Joaquin García-Cabo Herrero, Louisa Liles, Luke Lillehaugen, Matteo Luciani, Paul Tran and Riccardo Trezzi for their help in collecting some of the data used in the paper.

Disclaimer: The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

Francesca Loria (corresponding author). Federal Reserve Board, Research & Statistics Division. 20th Street and Constitution Avenue N.W., Washington, D.C. 20551, USA. Email: francesca.loria@frb.gov.

David López-Salido. Federal Reserve Board, Monetary Affairs Division. 20th Street and Constitution Avenue N.W., Washington, D.C. 20551, USA. Email: david.j.lopez-salido@frb.gov.

*“Monetary policy responded first in the summer of 2012 by acting to defuse the sovereign debt crisis, which had evolved from a **tail risk for inflation** into a material threat to price stability.”*

Mario Draghi, *ECB President, Sintra, June 2019.*¹

Introduction

Since the upheavals of the global financial crisis, the emergence of downside risks to the inflation outlook have increasingly become a source of macroeconomic concern. So far, most efforts have been devoted to studying the factors underlying the muted response of the *conditional mean* of inflation to economic and financial conditions. At the same time, much has been said on the inability of labor market conditions to explain recent inflation outcomes. The Phillips curve linkages seem to be breaking down. In this paper, we show that in the presence of tail risks, the conditional inflation mean does not necessarily fully depict the inflation outlook, as reminded by President Draghi’s quote.

Indeed, we show that the contrasting response of the inflation tails and median reveals a more complete picture of the effects that real and financial shocks impinge on its outlook. Specifically, we find that there have been sizeable downside risks to the inflation outlook in the last 20 years, mainly accounted for by financial tightenings. This is consistent with the idea that, due to amplification mechanisms, when financial conditions become tighter firms cut prices disproportionately more, on average. These concepts are related to recent research by [Del Negro, Giannoni, and Schorfheide \(2015\)](#), [Christiano, Eichenbaum, and Trabandt \(2015\)](#), [Christiano, Motto, and Rostagno \(2014\)](#) and [Gilchrist, Schoenle, Sim, and Zakrajšek \(2017\)](#) which shows that financial conditions matter also for inflation dynamics. While these studies focus on the modal outlook, we draw on their insights to study empirically the inflation tails.

Other studies already documented that looking at the entire predictive distribution of economic

¹Mario Draghi, “Twenty Years of the European Central Bank’s Monetary Policy,” speech delivered at the ECB Forum on Central Banking in Sintra on June 18th, 2019 (available at <https://www.bis.org/review/r190618c.htm>).

growth can reveal additional insights into their dynamics. For instance, the deterioration in credit market conditions led to a substantial decline in economic activity as well as a deterioration in the odds of low growth and of high unemployment. Tight financial conditions moved the conditional distribution of real GDP growth to the left (e.g., [Adrian, Boyarchenko, and Giannone, 2019](#) and [Caldara, Cascaldi-Garcia, Cuba Borda, and Loria, 2020](#)) – with its left tail being the most sensitive to macroeconomic shocks (see [Loria, Matthes, and Zhang, 2019](#)) – and implied medium-term upside risks to unemployment (see [Kiley, 2018](#)).

To summarize, in this paper we make three points. First, we offer evidence that some of the macroeconomic factors covered under the “Phillips curve umbrella” – conventionally used to study the conditional mean of inflation – are still at work in the *tails*. Indeed, we find ample variability in the tails of the inflation outlook that remains even when focusing on the most recent period of stable and low *mean* inflation. Second, we show that tight financial conditions carry substantial downside risks to the inflation outlook, an aspect of inflation behavior overlooked by much of the literature but consistent with financial amplification mechanisms. For instance, we replicate our findings within the model by [Gertler, Kiyotaki, and Prestipino \(2019\)](#) - a nonlinear DSGE model which features the possibility of a severe financial crisis. Third, we uncover that these two findings are supported by evidence from inflation probabilities derived from financial contracts, predictive densities from the SPF and a regime-switching model.

Our econometric strategy frames the effects of risk factors on inflation within an “augmented” quantile Phillips curve model using data since the 1970s for the U.S. economy. That is, we extend the standard regression analysis (e.g., [Blanchard, Cerutti, and Summers, 2015](#)) – designed to ascertain the drivers of the conditional mean of inflation – to different inflation quantiles. This setup allows to relate the risks to the inflation outlook to labor market slack, inflation inertia and inflation expectations, as well as relative prices. More importantly, we extend the analysis to consider the effect of financial conditions on the inflation distribution and on the odds of low inflation. To do so, we construct the predictive distribution of the inflation outlook by fitting a flexible distribution on the estimated inflation quantiles.

We check whether probabilities of inflation falling within certain intervals, as priced by financial market quotes, are consistent with some of the conclusions about inflation risks derived from our analysis. We show that these probabilities and the ones obtained from our quantile Phillips curve model share a defining feature, namely that tight financial conditions are associated with higher probabilities of low inflation and that this relationship weakens as one considers higher inflation cutoffs. We also compare the predictive densities obtained from our quantile regression model with those obtained from the SPF and find that they are remarkably similar.

To shed light on the interpretation and on the sources of identification of the inflation tails coming from the quantile regression, we run a regime-switching version of our augmented Phillips curve model. The regime-switching regression is consistent with the main findings from our quantile regression approach. In particular, we show that the Phillips-curve linkages in the left, median and upper inflation tails arising from the quantile regression are comparable to the Phillips-curve linkages informed by regimes of low, moderate and high inflation. In terms of our main result, we show that in regimes of low inflation, credit spreads disproportionately held down inflation as compared to times of moderate and high inflation.

We consider the recent global financial crisis to illustrate how economic and financial headwinds influenced the inflation outlook both in the United States and in the euro area. We show how in the U.S., while average inflation experienced only a modest reduction despite the fall in output triggered by the financial crisis, downside inflation risks moved considerably over this period – mainly a reflection of soaring credit spreads during the financial meltdown. These patterns have been less benign in the euro area, where the sovereign debt crisis triggered a more prolonged increase in the odds of low inflation due to the more limited role of inflation expectations in counteracting downside risks posed by the economic slowdown and financial distress.

Finally, we use our framework to characterize inflation risks in uncertain times. We find that the inflation distributions implied by our model for the recent Covid-19 episode correctly display an increase in downside inflation risks at the onset of the crisis and upside inflation risks in the recovery.

Related Studies Our paper speaks to the literature which studies risks to the inflation outlook. [Andrade, Ghysels, and Idier \(2012\)](#) introduced the nomenclature “inflation-at-risk” when constructing a measure of (left and right) tail risks to inflation using survey-based density forecasts. They also showed that the magnitude and the asymmetry of inflation risks evolves over time and that it is not only related to purely nominal factors but also informed by financial variables, among others. [Kilian and Manganelli \(2007, 2008\)](#) derive inflation risk measures from the private sector agent’s preferences with respect to inflation. In a cross-section of countries, [Cecchetti \(2008\)](#) computes t -distribution approximations to deviations of log GDP and log price level from their trend and documents that asset price booms increase both growth and inflation risks. [Manzan and Zerom \(2013\)](#) find that incorporating macroeconomic variables into quantile regressions improves the accuracy of inflation density forecasts. [Korobilis \(2017\)](#) finds that predictive densities from a quantile regression Bayesian model averaging (QR-BMA) model are superior to and better calibrated than those of the traditional regression BMA model and that this methodology is competitive with popular nonlinear specifications for U.S. inflation. [Galvão and Owyang \(2018\)](#) find that financial conditions have stronger effect on inflation on periods of “financial stress” in their factor-augmented smooth-transition vector autoregressive model (FASTVAR). [Ghysels, Iania, and Striaukas \(2018\)](#) construct measures of inflation risk using a Quantile Autoregressive Distributed Lag Mixed-Frequency Data Sampling (QADL-MIDAS) regression model and find that they contain useful information about future inflation realizations. [Adams, Adrian, Boyarchenko, and Giannone \(2021\)](#) construct risks around consensus forecasts of inflation, among others.

We consider our contribution to this literature as augmenting a well-understood and micro-founded time-series framework, that of a Phillips-curve, to study the risks to the entire distribution of the inflation outlook, coming from both conventional inflation determinants as well as financial conditions. Moreover, we offer a comprehensive review of how evidence on the role of financial conditions for downside risks to the inflation outlook is supported by a comparison with euro-area results as well as alternative frameworks to measure inflation-at-risk: survey evidence, financial market quotes, a regime-switching model and a macroeconomic model with financial panics.

Our approach differs from, and complements, studies that define inflation risks as the chance of lost purchasing power resulting from negative inflation-adjusted returns. These studies evaluate the inflation risk premium associated with the compensation required by investors for future expected inflation or deflation – typically using information contained in financial market quotes (see, e.g., [Boons, Duarte, de Roon, and Szymanowska, 2020](#) among many others). An important departure from our approach is that, in general, they lack an explicit link of these risks to specific macroeconomic outcomes.

Our framework has been taken as a starting point in some recent papers, who confirm our findings that financial conditions are important determinants of inflation risks, and especially of downside risks. Two examples are [Korobilis, Landau, Musso, and Phella \(2021\)](#) in the context of both our semi-structural Phillips curve model as well as other time series models with time-varying parameters, and [Banerjee, Contreras, Mehrotra, and Zampolli \(2020\)](#) in a cross-section of advanced and emerging economies using panel quantile regressions.

Outline In Section 1 we organize ideas by presenting our theoretical framework and empirical strategy. As time-variation emerges in the characterization of the determinants of the inflation distribution, we illustrate subsample results in Section 2 and use them to shed new light on modern inflation linkages. In Section 3 we show supporting evidence for our main findings coming from financial-markets-derived inflation probabilities, survey data and a regime-switching version of our augmented Phillips curve model. We focus on the global financial crisis and compare the United States and euro area inflation experiences in Section 4, while also exploring the role of financial conditions in affecting the odds of low inflation during that period. In Section 5 we show how our model tracks inflation risks during the Covid-19 crisis. Concluding remarks, future research avenues and policy implications are offered in Section 6. A data guide, additional material and robustness exercises are collected in the Appendix. We refer to this appendix material either explicitly in the main text or in footnotes.

1 Characterizing Inflation-at-Risk

In many circumstances the study of the determinants of the conditional mean of inflation may be sufficient to produce a good representation of the modal dynamics of inflation. In other cases, however, studying the response of the tails of the *predictive* inflation distribution is essential for providing a more complete picture. This is likely to be the case, for instance, in the presence of large real or financial shocks, as it aids understanding the effects that these shocks have on inflation. Because of these considerations, we extend the standard regression analysis – designed to ascertain the drivers of the conditional mean of inflation – to the entire inflation distribution.

In this section we describe the econometric specification we use to link economic and financial conditions with risks to the inflation outlook. We first describe conditional inflation quantiles as a function of observed economic and financial variables (risk factors). Second, we use these quantiles to approximate the inflation distribution. Variations in inflation risks are then measured according to how much the tails of the inflation distribution vary with the evolution of economic and financial factors. We refer to these “tail risks” as Inflation-at-Risk (*IaR*).

We frame the effects of different risk factors on inflation within an augmented quantile Phillips curve model. This setup allows us to relate inflation risks to variations in the amount of slack in the labor market, changes in inflation persistence, variations in inflation expectations, as well as movements in relative prices (imported goods and/or oil). Our Phillips curve model is “augmented” as it also incorporates financial conditions (approximated by credit spreads) as an additional factor affecting not just the mean, but mainly the tails of the inflation distribution.

1.1 (Phillips-Curve) Quantile Regressions

Quantile regression models are a flexible tool for studying the determinants of *IaR*.² Our inflation measure of interest is the (annualized) average core CPI inflation rate over the next year (that is, between quarter $t + 1$ and quarter $t + 4$, $\bar{\pi}_{t+1,t+4}$).³ We want to focus on quarterly inflation as this is

²For an introduction to the quantile regression methodology, see [Koenker \(2005\)](#).

³A similar approach is taken in [Adrian, Boyarchenko, and Giannone \(2019\)](#) for the average growth rate of GDP.

typically the main frequency of interest and as it allows to abstract from noisier movements in the monthly series. We consider a linear model for the conditional inflation quantiles whose predicted value

$$\widehat{Q}_\tau(\bar{\pi}_{t+1,t+4}|x_t) = x_t \hat{\beta}_\tau, \quad (1)$$

is a consistent linear estimator of the quantile function of $\bar{\pi}_{t+1,t+4}$ conditional on x_t – where $\tau \in (0, 1)$, x_t is a $1 \times k$ -dimensional vector of conditioning (risk) variables, and $\hat{\beta}_\tau$ is a $k \times 1$ -dimensional vector of estimated quantile-specific parameters.⁴ Accordingly, a determinant x_t may exert non-linear effects on inflation dynamics if it affects differently the median and the tails.

Some observers might wonder about how having overlapping observations in our dependent variable might affect our results. As discussed in [Caldara et al. \(2020\)](#), running a regression of this type, is akin to a “direct” forecast, as opposed to an “iterated” forecast where the dependent variable would be π_{t+1} and where one would iterate the one-step-ahead prediction to obtain multi-horizon forecasts. As shown in that same paper, in simulation direct and iterated forecasts deliver the same results, when they share the same linear (V)AR data-generating process. In empirical data the direct model actually is comparable to if not better than an iterated model in terms of coverage and correct calibration of the predictive density.⁵

Our model for conditional inflation quantiles extends the Phillips-curve model used in the literature. In particular, we closely follow [Blanchard, Cerutti, and Summers \(2015\)](#) which summarized a vast empirical literature on inflation dynamics. Formally, the baseline quantile regression model in (1) can be written as an augmented Phillips curve model:

$$\widehat{Q}_\tau(\bar{\pi}_{t+1,t+4}|x_t) = \hat{\mu}_\tau + (1 - \hat{\lambda}_\tau)\pi_{t-1}^* + \hat{\lambda}_\tau \pi_t^{LTE} + \hat{\theta}_\tau(u_t - u_t^*) + \hat{\gamma}_\tau(\pi_t^R - \pi_t) + \hat{\delta}_\tau F_t, \quad (2)$$

where risk factors affecting the distribution of future inflation can be divided in different blocks.

⁴Formally, the dependence between x_t and a quantile $\tau \in (0, 1)$ of $\bar{\pi}_{t+1,t+4}$ is measured by the coefficient $\hat{\beta}_\tau$:

$$\hat{\beta}_\tau = \arg \min_{\beta_\tau \in \mathbb{R}^k} \sum_{t=1}^{T-h} (\tau \cdot \mathbb{1}_{(\bar{\pi}_{t+1,t+4} \geq x_t \beta)} |\bar{\pi}_{t+1,t+4} - x_t \beta_\tau| + (1 - \tau) \cdot \mathbb{1}_{(\bar{\pi}_{t+1,t+4} < x_t \beta)} |\bar{\pi}_{t+1,t+4} - x_t \beta_\tau|),$$

where $\mathbb{1}_{(\cdot)}$ denotes the indicator function, taking the value one if the condition is satisfied.

⁵This can be rationalized by the fact that the direct model is agnostic about the nature of the evolution of the dependent variable over horizons and thus performs best when the DGP departs from a linear AR process.

A full description of the data is provided in Appendix A.

First, the variables π_{t-1}^* and π_t^{LTE} respectively represent average inflation over the previous four quarters and a measure of long-term inflation expectations. Lagged average inflation captures the role of “intrinsic persistence” or different forms of inertia in the price setting process that could precipitate upward or downward drift in the aggregate inflation rate.⁶ In some models, this variable proxies adaptive or non-rational expectations whereas in others it is used to capture backward-looking or simple rule-of-thumb pricing rules. Long-term inflation expectations approximate the importance of some firms setting prices in a rather forward-looking way. Which of these two elements dominates the persistence observed in the distribution of aggregate inflation depends on the size of the parameter λ_τ . To preserve the notion that inflation persistently deviates from longer-run inflation expectations, we impose the homogeneity constraint in prices by constraining the two coefficients to sum up to one. When $\lambda_\tau = 0$, the model becomes an extension of the accelerationist Phillips curve, where changes in inflation are a function of the unemployment gap. We impose $(1 - \lambda_\tau) + \lambda_\tau = 1$, $0 \leq (1 - \lambda_\tau) \leq 1$ and $0 \leq \lambda_\tau \leq 1$ using the inequality constrained quantile regression method developed by [Koenker and Ng \(2005\)](#).

The second risk factor is linked to variations in the amount of labor market slack – as measured by the unemployment gap $(u_t - u_t^*)$, where u_t is the civilian unemployment rate and u_t^* is the natural rate of unemployment. Most of the recent literature has concentrated on the stability over time of the parameters λ and θ to explain the evolution of average inflation. This literature has focused, for instance, on understanding the failure of average inflation to respond to unemployment – i.e., the flattening of the Phillips curve – and on the increasingly dominant role of inflation expectations in explaining inflation persistence – i.e., the well-anchoring of long-run inflation expectations. In this paper we extend this analysis by exploring the effects of these variables on the tails of the distribution of inflation. The importance of these effects is captured by the variation across quantiles of the parameters λ_τ , $(1 - \lambda_\tau)$ and θ_τ in expression (2).

The third risk factor in (2) is given by $(\pi_t^R - \pi_t)$, which reflects variations in relative prices.

⁶[Wolters and Tillmann \(2015\)](#) use a quantile regression model of core CPI and core PCE inflation which solely conditions on past inflation to study how inflation persistence differs across quantiles.

We use the quarterly change in relative import prices ($\pi_t^I - \pi_t$). As in [Blanchard, Cerutti, and Summers \(2015\)](#), this variable is usually included to capture the pass-through of both nominal exchange rates and oil prices into core inflation measures and is perceived as having been a key driver of the run-up of inflation in the late seventies and the eighties. Lately, this variable has been used to approximate a wide range of risk factors, from changes in global commodity prices, taxes and tariffs to other global influences on domestic inflation. Its effects on the inflation distribution are captured by the cross-quantile variation in the parameters γ_τ in expression (2).

The last, but not least, risk factor that we consider is related to financial conditions. According to conventional wisdom, economic factors – labor market slack, inflation expectations, and relative prices – have been considered as the major sources of variation in the conditional mean of inflation. However, recent research by [Del Negro, Giannoni, and Schorfheide \(2015\)](#), [Christiano, Eichenbaum, and Trabandt \(2015\)](#), [Christiano, Motto, and Rostagno \(2014\)](#) and [Gilchrist, Schoenle, Sim, and Zakrajšek \(2017\)](#) suggests that changes in firms’ financial conditions (proxied by variations in credit spreads) also helps to explain inflation dynamics. After the financial stress of the fall of 2008, these studies aim at explaining how the sharp contraction in economic activity was accompanied by only a modest decline in (average) inflation. However, they mostly discuss the role of financial frictions in amplifying the business cycle and creating adverse feedback loops, while leaving its implications for inflation not fully developed. We thus allow for financial conditions F_t in expression (2), to affect differently the conditional inflation quantiles. This allows a test for the presence of differential effects of financial variables on the mean versus the tails of the inflation distribution (e.g., through the variation in δ_τ). Following these authors, and as recommended by [Gilchrist and Zakrajšek \(2012\)](#), we approximate F_t by the credit spread, cs_t .

1.2 Quantile Function of Inflation

The estimated conditional quantiles are approximations to the so-called “quantile function”, that is, $\mathcal{Q}_\tau(\bar{\pi}_{t+1,t+4}|x_t) = F_{\bar{\pi}_{t+1,t+4}}^{-1}(\tau|x_t)$, where $F^{-1}(\cdot)$ is the conditional inverse cumulative distribution function (CDF) of average future inflation. We follow [Adrian, Boyarchenko, and Giannone \(2019\)](#)

in smoothing the quantile function using the skewed t -distribution proposed by [Azzalini and Capitanio \(2003\)](#). This flexible distribution is characterized by four parameters and given by:

$$f(\bar{\pi}_{t+1,t+4}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t) = \frac{2}{\sigma_t} \times t(z_{t,t+4}; \kappa_t) \times T\left(\eta_t z_{t,t+4} \sqrt{\frac{\kappa_t + 1}{\kappa_t + z_{t,t+4}^2}}; \kappa_t + 1\right), \quad (3)$$

where $z_{t,t+4} = \frac{\bar{\pi}_{t+1,t+4}(x_t) - \mu_t}{\sigma_t}$ and t and T respectively represent the density and cumulative distribution function of the student t -distribution. The constants $\mu_t \in \mathbb{R}$ and $\sigma_t \in \mathbb{R}^+$ are location and scale parameters, whereas the constants $\eta_t \in \mathbb{R}$ and $\kappa_t \in \mathbb{Z}^+$ control the skewness and the kurtosis of the distribution, respectively. We compute these parameters at each point in time t to minimize the squared distance between our estimated quantile function $\hat{Q}_\tau(\bar{\pi}_{t+1,t+4}|x_t)$, obtained from the quantile Phillips-curve model (2), and the quantile function of the skewed t -distribution $F_{\bar{\pi}_{t+1,t+4}}^{-1}(\tau|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t)$ to match the 10th, 25th, 75th and 90th quantiles.

2 The Time-Varying Dynamics of Inflation-at-Risk

Two distinct subsamples emerge when characterizing the determinants of the inflation distribution in the United States.⁷ The first period, running from 1973:Q1 to 1999:Q4, covers the OPEC shocks, the subsequent Volcker disinflation and the early stages of the Great-Moderation. The presence of large shocks to relative prices and the taming of inflation expectations induced large swings in the upper quantile, while changes in unemployment and past inflation affected the median. These are ubiquitous themes in the description of inflation dynamics.

The first period contrasts with the second subsample, from 2000:Q1 to 2019:Q1, characterized by large movements in credit spreads, progressively well-anchored inflation expectations but subdued inflation pressures. These patterns are studied in most of the literature discussing a favorite whipping boy – the flatness of the Phillips curve. Well-anchored long-run inflation expectations [Yellen \(2013\)](#), systematic monetary policy (e.g. [Ball and Mazumder, 2019](#), [McLeay and Tenreyro,](#)

⁷We study the role of economic and financial conditions for the risks to the United States inflation outlook using our full sample in Appendix B.

2018) and mismeasurement of labor market slack (e.g., [Stock and Watson, 2019](#)) are the usual suspects in explaining the observed muted response in average inflation, also referred to as the “missing deflation/inflation puzzle” [Williams \(2010\)](#) (more on this in the next subsection). However, the financial crisis and the period in which monetary policy has been constrained by the zero lower bound, have been followed by a period of underperformance of inflation relative to explicit or implicit inflation targets. This period has ended with reductions, of different size, of long-term inflation expectations. Some authors have pointed out that the risks of persistent below-target inflation are associated with the emergence of this phenomenon and claim that this set the seeds for further downside risks to inflation. Through this section we will show that tight credit conditions arising from financial crises also contributed to increasing odds of low inflation or even deflation, pointing to a greater role of labor market recovery and well-anchored inflation expectations in supporting average inflation. This point will be further investigated in the [Section 4](#) using contrasting evidence from the United States and the euro area.

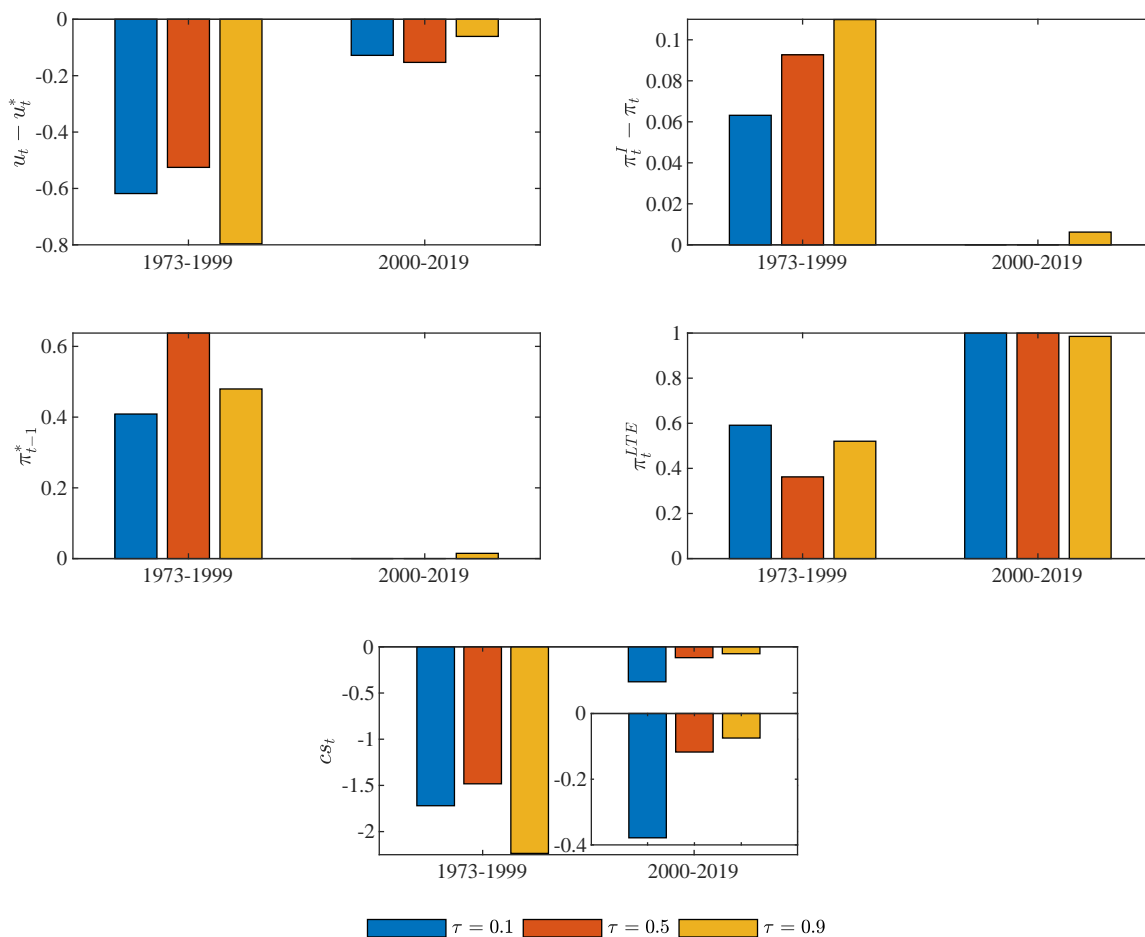
2.1 Subsample Stability and the Missing Deflation/Inflation

To investigate how the importance of risk factors changed across the two subsamples, we report their estimated quantile-specific slopes in [Figure 1](#). Three results stand out. First, relative import prices still pose threats to the upper inflation quantile, though to a lesser degree than prior to the Great Moderation. Second, inflation inertia has completely lost its grip on inflation, crowning long-run inflation expectations as the decisive inflation determinant among the variables in the modern Phillips curve, as supported by the time-varying parameter model for mean inflation in [Blanchard, Cerutti, and Summers \(2015\)](#).⁸ Third, long-run inflation expectations exert a symmetric effect on the inflation distribution. In fact, in the recent period of well-anchored long-run inflation expectations the response of average inflation to labor market slack, financial conditions and relative price

⁸Among others, we use long-term inflation expectations from Consensus Economics as it provides equivalent data for the euro area, an interesting comparison to the U.S. experience (see [Section 4](#)). Results are similar if we use long-term inflation expectations from the SPF or Michigan survey, which are respectively available from 1987Q1 and 1981Q1.

changes is dampened. However, it would be misleading to dismiss the role of these factors focusing on the conditional mean only. Instead, credit conditions and, to a lesser extent, labor market outcomes are key drivers of the asymmetry in the inflation distribution.

Figure 1: Quantile Regression Slopes Across Subsamples.



NOTE: The figure shows the estimated slopes of the quantile regression of average four-quarter-ahead core CPI inflation in the U.S., defined in expression (2). Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate the coefficients for the 10th quantile (blue), median (red) and 90th quantile (yellow).

The Role of Credit Spreads There is substantial subsample instability governing the link between the inflation tails and variations in credit spreads. The sub-period 1973-1999 is characterized by relatively small variations in credit spreads in a period of high and volatile inflation induced in part by systematic increases in energy prices – reason why the actual contribution of credit spreads to the inflation outlook in that period is smaller. From 2000 onward, low variability of inflation

around 2 percent has been a notable aspect of the stability of the macroeconomic landscape that has coexisted with substantial variation in credit spreads, a phenomenon amplified by the global financial crisis.

Focusing on the most recent subsample, a novel result is that the link between the inflation outlook and financial conditions is asymmetric, in that an increase in credit spreads is associated with a larger increase in downside risks than in upside risks. (We performed ANOVA tests on the equality of the slope coefficients across quantiles and can reject the null hypothesis of equality between the slopes on credit spreads for the 10th and 50th quantile at 1% significance level).

This result resonates with the findings by [Korobilis, Landau, Musso, and Phella \(2021\)](#) in the context of both time-varying parameter and stochastic volatility versions of both our semi-structural Phillips curve model and competitive time series models, by [Cecchetti \(2008\)](#) and [Banerjee, Contreras, Mehrotra, and Zampolli \(2020\)](#) in a cross-section of countries and is reminiscent of the results for the GDP growth outlook in [Adrian, Boyarchenko, and Giannone \(2019\)](#). Moreover, this result can be generated in nonlinear models with amplification mechanisms (e.g., financial accelerator models). Specifically, it captures the notion that when financial conditions become tighter firms cut prices disproportionately more, on average.⁹ In Appendix C, we show that our findings can for instance be replicated by applying our quantile regression framework to simulated data generated from the model by [Gertler, Kiyotaki, and Prestipino \(2019\)](#) - a fully micro-founded and nonlinear DSGE model which features the possibility of a severe financial crisis.

In Appendix D, we show that these results hold also with different measures of financial conditions and of inflation. Most importantly, the same patterns emerge when controlling for SPF forecasts, thus proving that the credit spread is informative for inflation beyond its forward-looking component and, thus, its information content for future economic activity.

⁹[Gilchrist, Schoenle, Sim, and Zakrajšek \(2017\)](#) explain why when firms experience a large drop in their liquidity that pushes them to their constraints, such as in the Great Recession, one might see less deflation than otherwise predicted from a model with homogenous firms that does not take into account liquidity constraints. Liquidity-constrained firms restrain from cutting prices below marginal costs to support their cash-flow. However, this does not mean that tighter financial conditions result in “higher” inflation. The predominant effect is still the one on the extensive margin, whereby the not liquidity-constrained firms lower their prices disproportionately more than during normal times. During regimes of low inflation, tight credit conditions thus impinge a more negative effect than usual.

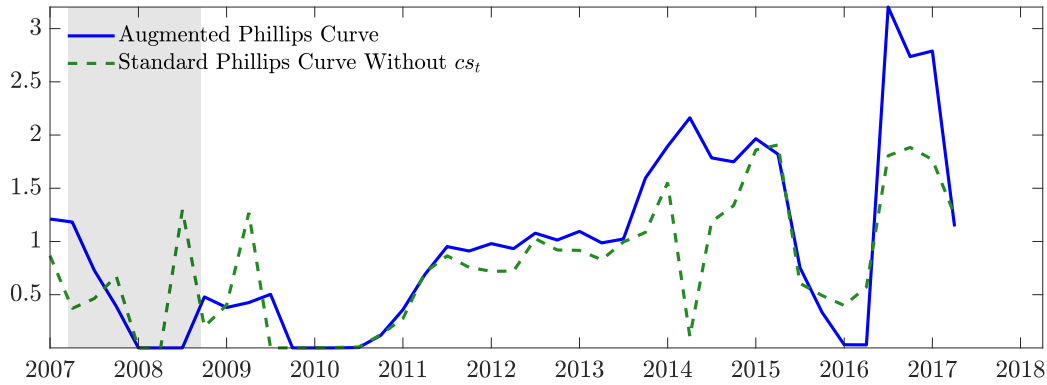
A Modern Phillips Curve The time-varying sensitivity of inflation to inflation expectations (and inflation inertia), as well as its ability to explain the “missing deflation/inflation” puzzle, had already been explored and established by [Blanchard, Cerutti, and Summers \(2015\)](#). In Appendix E we extend their findings to the entire distribution of the inflation outlook by constructing “counterfactual” inflation distributions and inflation-at-risk probabilities from the perspective of these two different subsamples. Our results confirm that the conventional wisdom from the pre-2000 economy would have suggested much larger deflation risks during the Great Recession and much lower upside risks to the inflation outlook during the recovery than the modern Phillips-curve relationship. We take this as evidence that the model estimated starting from the 2000’s does indeed better characterize the inflation dynamics in recent times.

Predictive Ability of Credit Spreads Next, we show that for the sample starting in the 2000s, the augmented Phillips curve model outperforms the standard model that ignores credit spreads in terms of predictive ability. In Section B.1 of Appendix B we find that the augmented model is also superior in terms of coverage by applying the [Rossi and Sekhposyan \(2019\)](#) test for correct calibration of the predictive density .

The reliability of the predictive distribution can be assessed by measuring the accuracy of the model’s density forecasts through its predictive scores. These are computed by evaluating the model’s predictive distribution at the realized value of the time series. A higher predictive score indicates more accurate predictions, as the model assigns a higher probability to outcomes that are closer to the realized value. We compute the predictive scores in an out-of-sample exercise where the predictive distributions are calculated using an expanding window.

Figure 2 plots the out-of-sample predictive scores for the modern Phillips curve model estimated starting in 2000:Q1 (see Section 2). The augmented Phillips curve model has a higher predictive ability, on average, than the standard Phillips curve model that ignores credit spreads. We take this as evidence of the importance of considering financial conditions in the modern Phillips curve in characterizing the inflation outlook.

Figure 2: Out-of-Sample Predictive Scores. Model Estimated from 2000:Q1.



NOTE: The figure illustrates the out-of-sample predictive scores obtained from the density constructed by fitting a skewed- t distribution on the conditional quantiles from the quantile regression model (2) that starts in 2000:Q1. The quantiles are constructed via an expanding rolling windows estimation starting in 2006:Q4. In particular, for the first out-of-sample evaluation, we use data for average core CPI inflation over the next year until 2005:Q4 to estimate the model. Then, we use data for the explanatory variables in 2006:Q4 to construct the out-of-sample predictive scores.

3 Evidence from Financial Markets and Regime-Switching Model


In this section we provide further evidence of our main finding that higher credit spreads are associated with higher downside risks to the inflation outlook. First, we look at inflation probabilities derived from inflation options and show that credit spreads are most negatively associated with low inflation probabilities and that their relationship weakens as one considers higher inflation cut-offs. We then consider predictive densities of various inflation measures derived from the Survey of Professional Forecasters and find that they are remarkably similar to those obtained from our quantile regression model. Finally, we estimate a monthly regime-switching version of our augmented Phillips curve model over the last 20 years of data and show how the main findings of our quantile regression model are directly comparable to and confirmed by this alternative approach.

3.1 Evidence from Financial Markets

Two defining features of our quantile Phillips curve model are that tight financial conditions carry substantial downside risks to inflation and that these risks diminish as one moves to the left to the upper tail of the inflation distribution. To test whether this relationship also holds in financial

markets, we run an OLS regression of options-implied inflation probabilities of one-year-ahead CPI inflation derived from inflation caps and floors contracts (as described by [Kitsul and Wright, 2013](#)) on the credit spread. In the left panel of Figure 3 we present the estimated coefficients of that regression. The slopes are rescaled so as to facilitate the comparison with those coming from our quantile Phillips curve model (the right panel reproduces the bottom-right box of Figure 8). In particular, the coefficient for the probability of one-year-ahead inflation being below 1% is rescaled to match the slope estimated on the lowest inflation quantile which arises from the quantile Phillips curve model. Further, the coefficient for the probability of inflation being below 1% is transformed from positive to negative – as a positive correlation between the credit spread and this probability is equivalent to a negative relationship between the credit spread and the lowest inflation quantile.

Figure 3: Credit Spread Slopes, Quantile Regression vs. Financial Markets.



NOTE: The left panel reports the slopes of separate regressions of inflation probabilities on the credit spread (at monthly frequency), along with their 95% confidence interval. The coefficient for the probability of future inflation being below 1% is rescaled to match the slope estimated on the lowest inflation quantile which arises from the quantile Phillips curve model (right panel, taken from Figure 8). The coefficients are transformed from positive to negative for the probability of inflation being below 1% – as a positive correlation between the credit spread and this probability is equivalent to a negative relationship between the credit spread and the lowest inflation quantile.

Despite the vast disparities in the construction of the tails of the inflation distribution, the estimated slopes are very similar to each other. Most importantly, as the inflation probability cutoffs increase, their relationship with credit spreads weakens. We thus confirm our findings from the quantile regression that credit spreads exert their biggest downward pressure on the left tail of the distribution and thus are an important factor behind its asymmetry.¹⁰

¹⁰This result is robust to the inclusion of the regressors used to purify the inflation probabilities from the oil price effects (see discussion on these effects in Appendix F).

Finally, in Appendix F we plot the time series of these financial-market-based probabilities together with the credit spread (see Figure F-3). The graphs confirm the progressive weakening in the relationship between the credit spread and inflation probabilities as one moves further up in the inflation distribution until it completely breaks down once the upper tail is reached. We also show that inflation probabilities coming from financial markets and from our quantile Phillips curve model point in the same direction.

3.2 Evidence from Surveys of Professional Forecasters

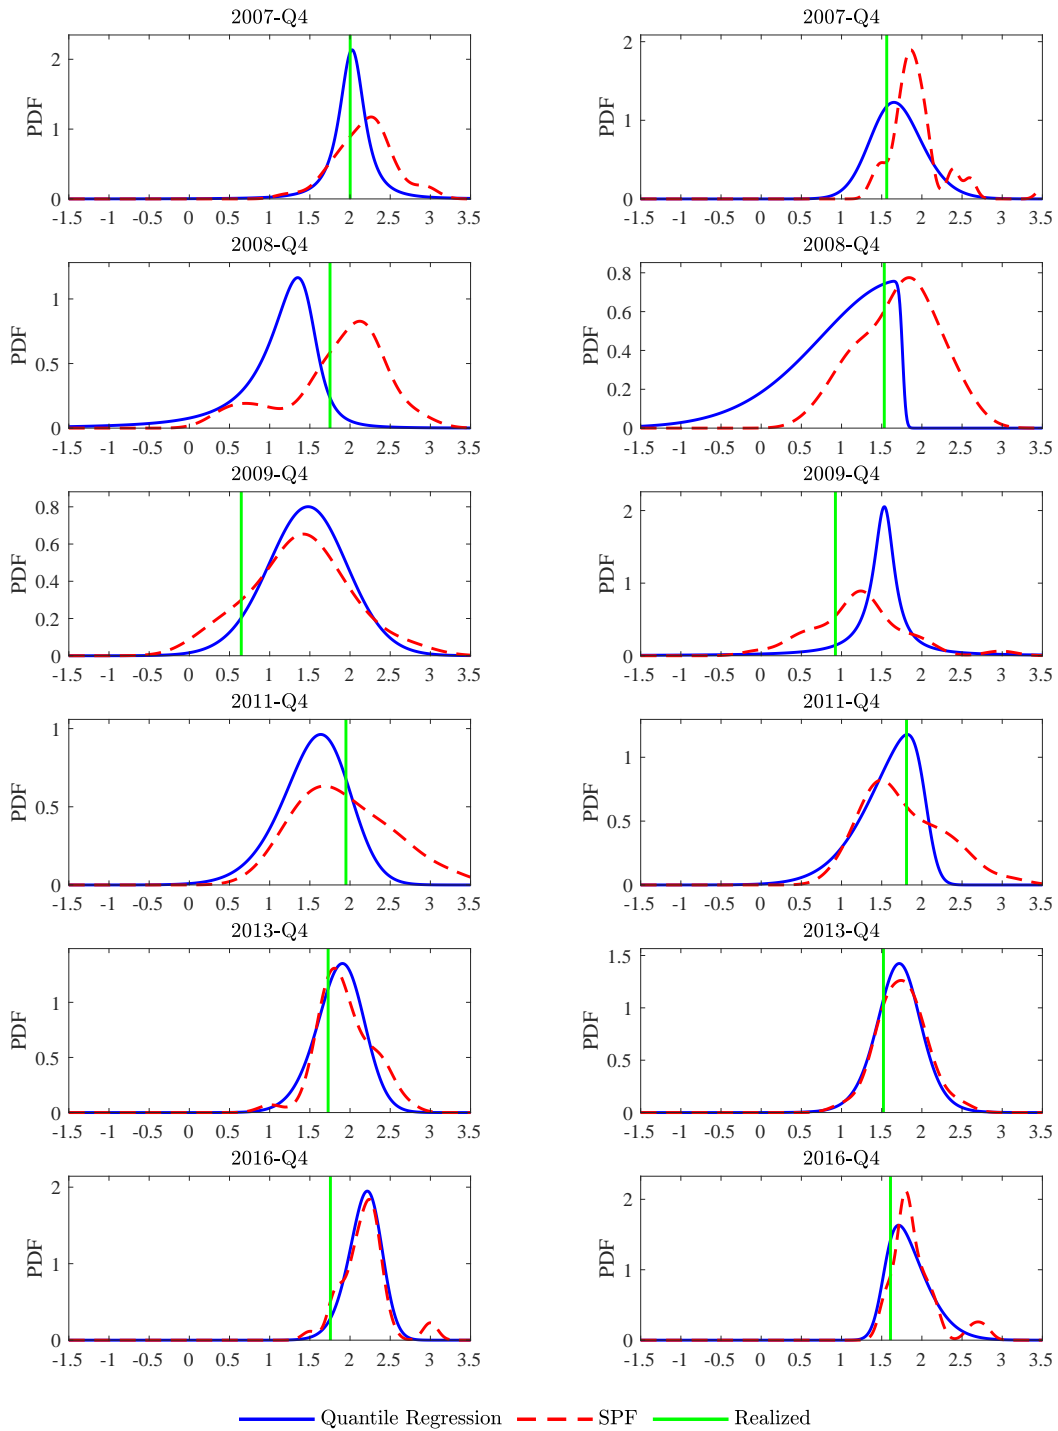
We now compare the predictive densities coming from our quantile Phillips curve model with those coming from the Surveys of Professional Forecasters (SPF). In Figure 4 we do this for both core CPI (left column) and core PCE inflation (right column), as both inflation measures are available in the SPF and applicable to our statistical model. We display the predictive densities of average one-year ahead inflation for several dates before and after the Great Recession; and always for the last quarter in a year, so that the horizon considered by the SPF and by the quantile regression model are aligned.¹¹ The SPF forecast densities (red dashed) are computed by applying a kernel smoothing function on the point estimates across survey respondents.¹² The quantile regression densities (blue solid) are computed by fitting a skewed- t distribution on the estimated quantiles, as described in Section 1.2. Finally, we report the realized value with a green solid vertical line.

The two densities are similar over time and when they do not align, the quantile Phillips curve model seems to be closer to the realized value. In 2007-Q4, for instance, the SPF densities are more optimistic than our statistical model, which is informed by the increase in credit spreads.

¹¹Ganics, Rossi, and Sekhposyan (2020) discuss the importance of predictive densities based on fixed-horizon density forecasts.

¹²Del Negro, Bassetti, and Casarin (2020) construct subjective SPF forecast distributions using non-parametric Bayesian methods.

Figure 4: Predictive Densities from Quantile Regression vs. from SPF Forecasts. Core CPI (Left) and Core PCE Inflation (Right).



NOTE: The figure shows the estimated skewed t -Student densities from the quantile Phillips curve model (blue solid) of average four-quarter-ahead core PCI inflation (left panels) and core PCE inflation (right panels). The figure also reports the realized value (green vertical line) and the density of SPF forecasts (red dashed) of the respective year-on-year inflation measures, computed by applying a kernel smoothing function on the point estimates across survey respondents.

3.3 Evidence from Regime-Switching Regression

We analyze how a nonlinear model such as a regime-switching regression compares to quantile regression estimates, as [Caldara, Cascaldi-Garcia, Cuba Borda, and Loria \(2020\)](#) in the growth-at-risk context. We do so to add further evidence that the relationships we established in our main analysis are not an artifact of the quantile regression but a genuine feature of the data that alternative nonlinear models would identify as well.

Since in the regime-switching model we estimate more parameters than in the quantile regression (in particular, the volatility of the error term and the regimes' transition probabilities), we are going to move to monthly frequency to allow for more observations and better identification of the regimes. To this end, we first introduce a monthly version of our quantile regression model and show that the results mirror the ones from the quarterly specification. Next, we compare the findings from the regime-switching model to the ones from quantile regression.¹³ As we will show, the asymmetric effect of credit spreads on the inflation distribution and the symmetric effect of inflation expectations is a robust finding across these approaches.

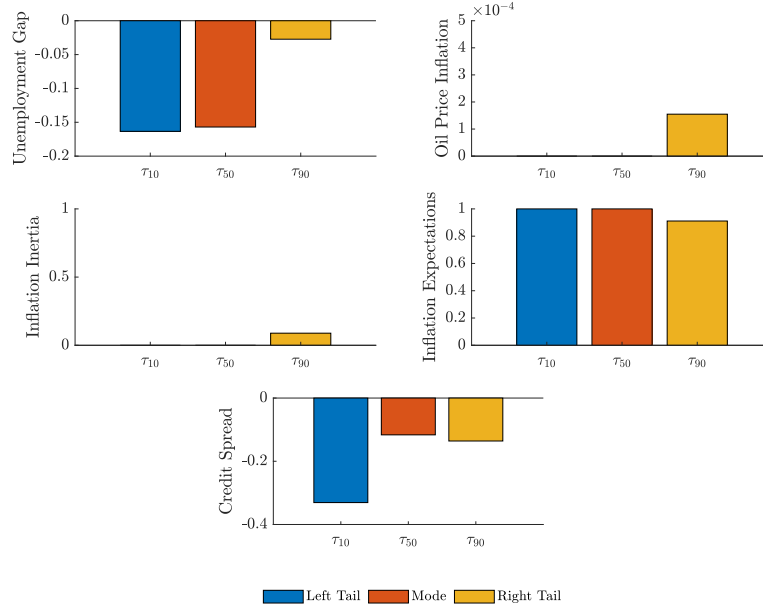
Monthly Quantile Regression We explore the monthly version of the quantile regression model described in (2) by interpolating the inflation expectations data. The sample runs from from January 2000 to April 2019, the last available date for core CPI inflation over the next 12 months. Formally, we estimate the following quantile Phillips curve

$$\widehat{Q}_\tau(\bar{\pi}_{t+1,t+12}|x_t) = (1 - \hat{\lambda}_\tau)\pi_{t-1}^* + \hat{\lambda}_\tau\pi_t^{LTE} + \hat{\theta}_\tau(u_t - u_t^*) + \hat{\gamma}_\tau(\pi_t^O - \pi_t) + \hat{\delta}_\tau cs_t, \quad (4)$$

where $\bar{\pi}_{t+1,t+12}$ is the (annualized) average core CPI inflation rate between month $t + 1$ and month $t + 12$, and where as in the quarterly model several risk factors affect the quantiles. Absent the import price index at monthly frequency, we use oil price inflation π_t^O in the relative price term. As shown in [Figure 5](#), the results resembles those found at quarterly frequency for the most recent sample (see [Figure 1](#)).

¹³In our baseline specification we stick to quarterly frequency as it allows to abstract from more volatile and temporary movements in inflation and it does not require to interpolate data.

Figure 5: Slopes of Inflation Determinants Across Quantiles, Monthly Model.



NOTE: Estimated 10th, 50th and 90th quantile regression slopes from monthly quantile regression model (4).

Monthly Regime-Switching Regression The Markov-switching regression model is given by

$$\bar{\pi}_{t+1,t+12} = \mu(s_t) + (1 - \lambda(s_t)) \pi_{t-1}^* + \lambda(s_t) \pi_t^{LTE} + \theta(s_t)(u_t - u_t^*) + \gamma(s_t)(\pi_t^O - \pi_t) + \delta(s_t)cs_t + \sigma(s_t)\varepsilon_t, \quad (5)$$

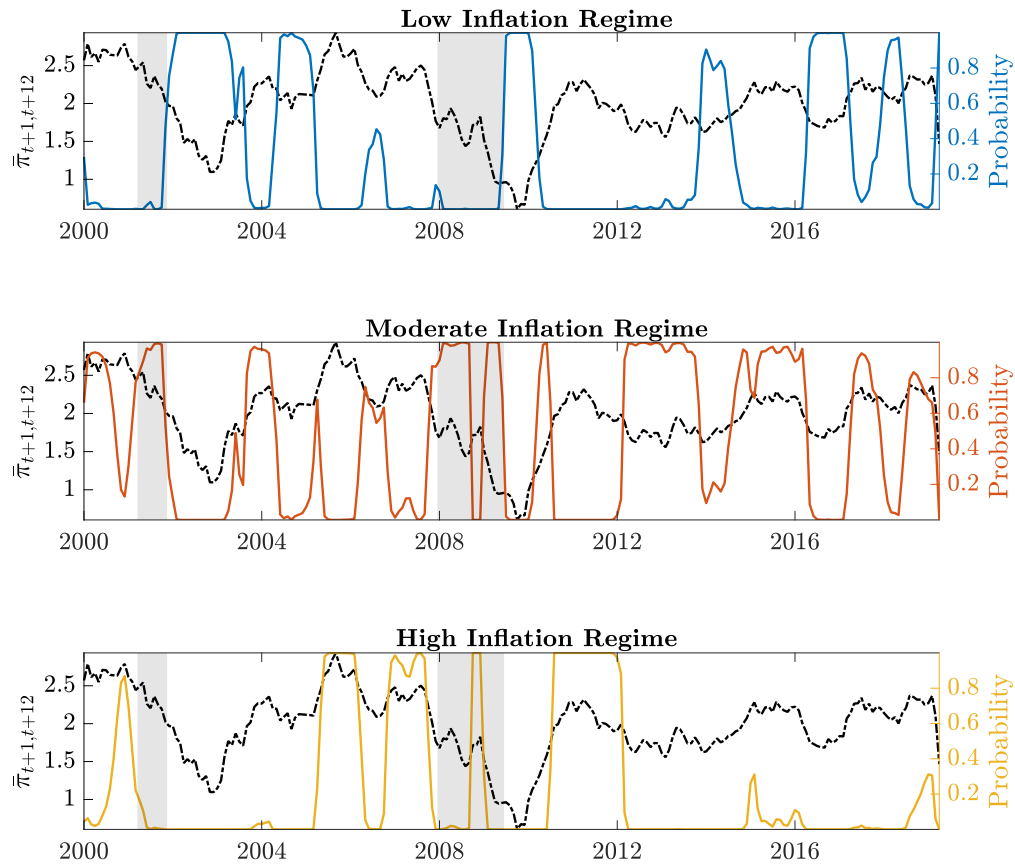
where both the coefficients $\Theta(s_t) \equiv \{\mu(s_t), \lambda(s_t), \theta(s_t), \gamma(s_t), \delta(s_t)\}$ and the standard deviation $\sigma(s_t)$ vary depending on an unobserved regime variable $s_t \in \{1, 2, 3\}$ which indicates the regime prevailing at time t . The latent variable s_t is governed by a discrete time, discrete state Markov stochastic process, defined by the transition probabilities:

$$p_{ij} = Pr(s_{t+1} = j | s_t = i), \quad \sum_{j=1}^3 p_{ij} = 1, \quad \forall i, j \in \{1, 2, 3\} \quad (6)$$

Estimation is done via Bayesian methods, details on priors and model fit are in Appendix G.

Inflation Regimes The estimated regimes broadly correspond to states where inflation is low, moderate or high. This becomes clear when comparing realized average inflation over the next year against the estimated regime probabilities, as we do in Figure 6.

Figure 6: Regime Probabilities from Markov-Switching Regression.

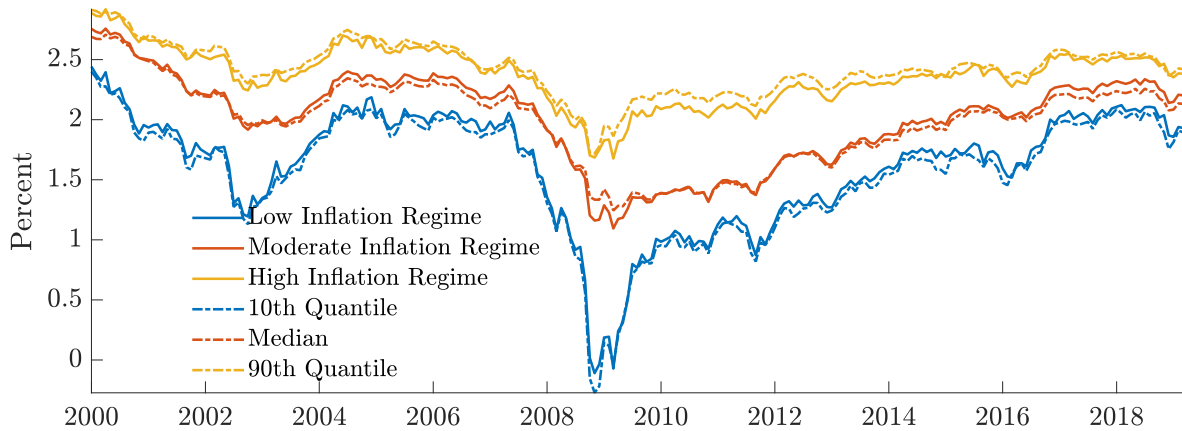


NOTE: Estimated regime probabilities from regime-switching regression (5) featuring three regimes and one Markov chain for both coefficients and volatility.

Relationship to Quantile Regression In Figure 7, we report the regime-specific fitted values along with the estimated quantiles from the monthly quantile regression model (4), which are remarkably similar. In particular, the low inflation regime corresponds to the 10th quantile, the moderate inflation regime to the median and the high inflation regime to the 90th quantile.

Notice how over the last 20 years, periods of stable inflation have been accompanied by the presence of sizeable downside risks arising from tight financial conditions. It is beyond the scope of this paper to explore which structural factors influencing our inflation determinants kept risks in check and thus stood behind the remarkable resistance of realized inflation from falling to negative territory during the financial crisis; but one can reasonably speculate that swift monetary policy communication and intervention into financial markets might have provided the needed cushion.

Figure 7: Fitted Values: Regime-Switching vs. Quantile Regression.



NOTE: The figure reports regime-specific fitted values of average core CPI inflation over the next year along with the conditional quantiles of that same variable estimated from the quantile regression. The model features one Markov chain governing both coefficients and volatility.

What explains the similarities in the assessment of risks between the regime-switching and the quantile regression models? Intuitively, this is due to the fact that the regime-switching model identifies regimes of low, moderate and high inflation that roughly correspond to the subsample of datapoints that most heavily inform the slopes on the 10th, 50th and 90th quantile in the quantile regression. This becomes most evident by comparing the estimated coefficients from the regime-switching regression (5) with the ones from the quantile regression (4), as we do in Table 1.

The regime-switching regression confirms our main findings from the quantile regression approach. Historically, credit spreads disproportionately held down inflation in regimes when inflation was low, in line with Galvão and Owyang (2018) who also found that financial conditions have stronger effect on inflation in periods of financial stress. On the contrary, inflation expectations have a symmetric effect across regimes. Moreover, in regimes of high inflation, the unemployment rate made less dent on inflation outcomes. Finally, we notice that the Markov-switching regression identifies little or no changes in the volatility of inflation across regimes.

Thus, the inflation distribution arising from the quantile regression should be interpreted as capturing the range of outcomes that inflation might take, on average, over the next year given the historical relationships with its determinants during regimes of high, moderate and low inflation.

Table 1: Estimated Coefficients Across Models

Regime-Switching Regression						
	<i>Low Inflation</i>		<i>Moderate Inflation</i>		<i>High Inflation</i>	
$\mu(s_t)$	0.30	[0.25, 0.41]	0.37	[0.18, 0.33]	0.50	[0.39, 0.63]
$\lambda(s_t)$	1.00	[0.99, 1.01]	1.00	[0.96, 1.02]	1.00	[0.85, 1.02]
$\theta(s_t)$	-0.16	[-0.19,-0.13]	-0.17	[-0.18,-0.15]	-0.07	[-0.10,-0.05]
$\gamma(s_t)$	0.00	[-0.01, 0.01]	0.00	[-0.02, 0.04]	0.00	[-0.02, 0.15]
$\delta(s_t)$	-0.33	[-0.38,-0.28]	-0.15	[-0.18,-0.13]	-0.15	[-0.18,-0.12]
$\sigma(s_t)$	0.18	[0.16, 0.21]	0.13	[0.11, 0.15]	0.14	[0.11, 0.19]

Quantile Regression						
	<i>10th Quantile</i>		<i>Median</i>		<i>90th Quantile</i>	
μ_τ	0.24	[-0.57, 1.05]	0.12	[-0.34, 0.58]	0.52	[-0.03, 1.08]
λ_τ	1.00	[0.63, 1.37]	1.00	[0.80, 1.20]	1.00	[0.77, 1.23]
θ_τ	-0.16	[-0.20,-0.13]	-0.16	[-0.18,-0.13]	-0.05	[-0.08,-0.02]
γ_τ	0.00	[0.00, 0.00]	0.00	[0.00, 0.00]	0.00	[0.00, 0.00]
δ_τ	-0.33	[-0.39,-0.27]	-0.12	[-0.15,-0.08]	-0.14	[-0.20,-0.09]

NOTE: Estimated coefficients from regime-switching regression (5) and from quantile regression (4). Both models are estimated at monthly frequency from January 2000 to April 2019. Values in brackets for regime-switching model indicate lower (5%) and upper (95%) bound of coefficients derived from their posterior distribution calculated via MCMC using 20,000 replications, a burn-in of 1,000 replications and a thinning parameter of 10 (thus for a total of 210,000 effective replications). Values in brackets for quantile regression model indicate lower (16%) and upper (84%) bound of coefficients computed via “blocks-of-blocks” bootstrap (see Appendix H) using 10,000 replications.

4 The Great Recession: United States vs. Euro Area

We now analyze the effect of inflation drivers on the evolution of the inflation distribution during the last 20 years of data, comparing the United States experience with that of the euro area. For the eurozone, the analysis focuses on euro-area-wide core HICP inflation – measured by headline HICP inflation excluding energy and unprocessed food.¹⁴ As for the U.S., the quantile regression model (2) uses the sample period available for the euro area, that is, from 1999:Q1 to 2017:Q4.¹⁵

¹⁴Busetti, Caivano, and Rodano (2015) use dynamic quantile regressions to forecast the conditional distribution of euro-area inflation. Tagliabracci (2018) estimates this distribution conditioning on Eurocoin. Busetti, Caivano, Monache, and Pacella (2019) investigate the role of domestic and global determinants of euro area core inflation by estimating a Phillips curve via expectile regression.

¹⁵The last date for which average inflation over the next four quarters is available is thus 2016:Q4. The data is described in great detail in Appendix A.

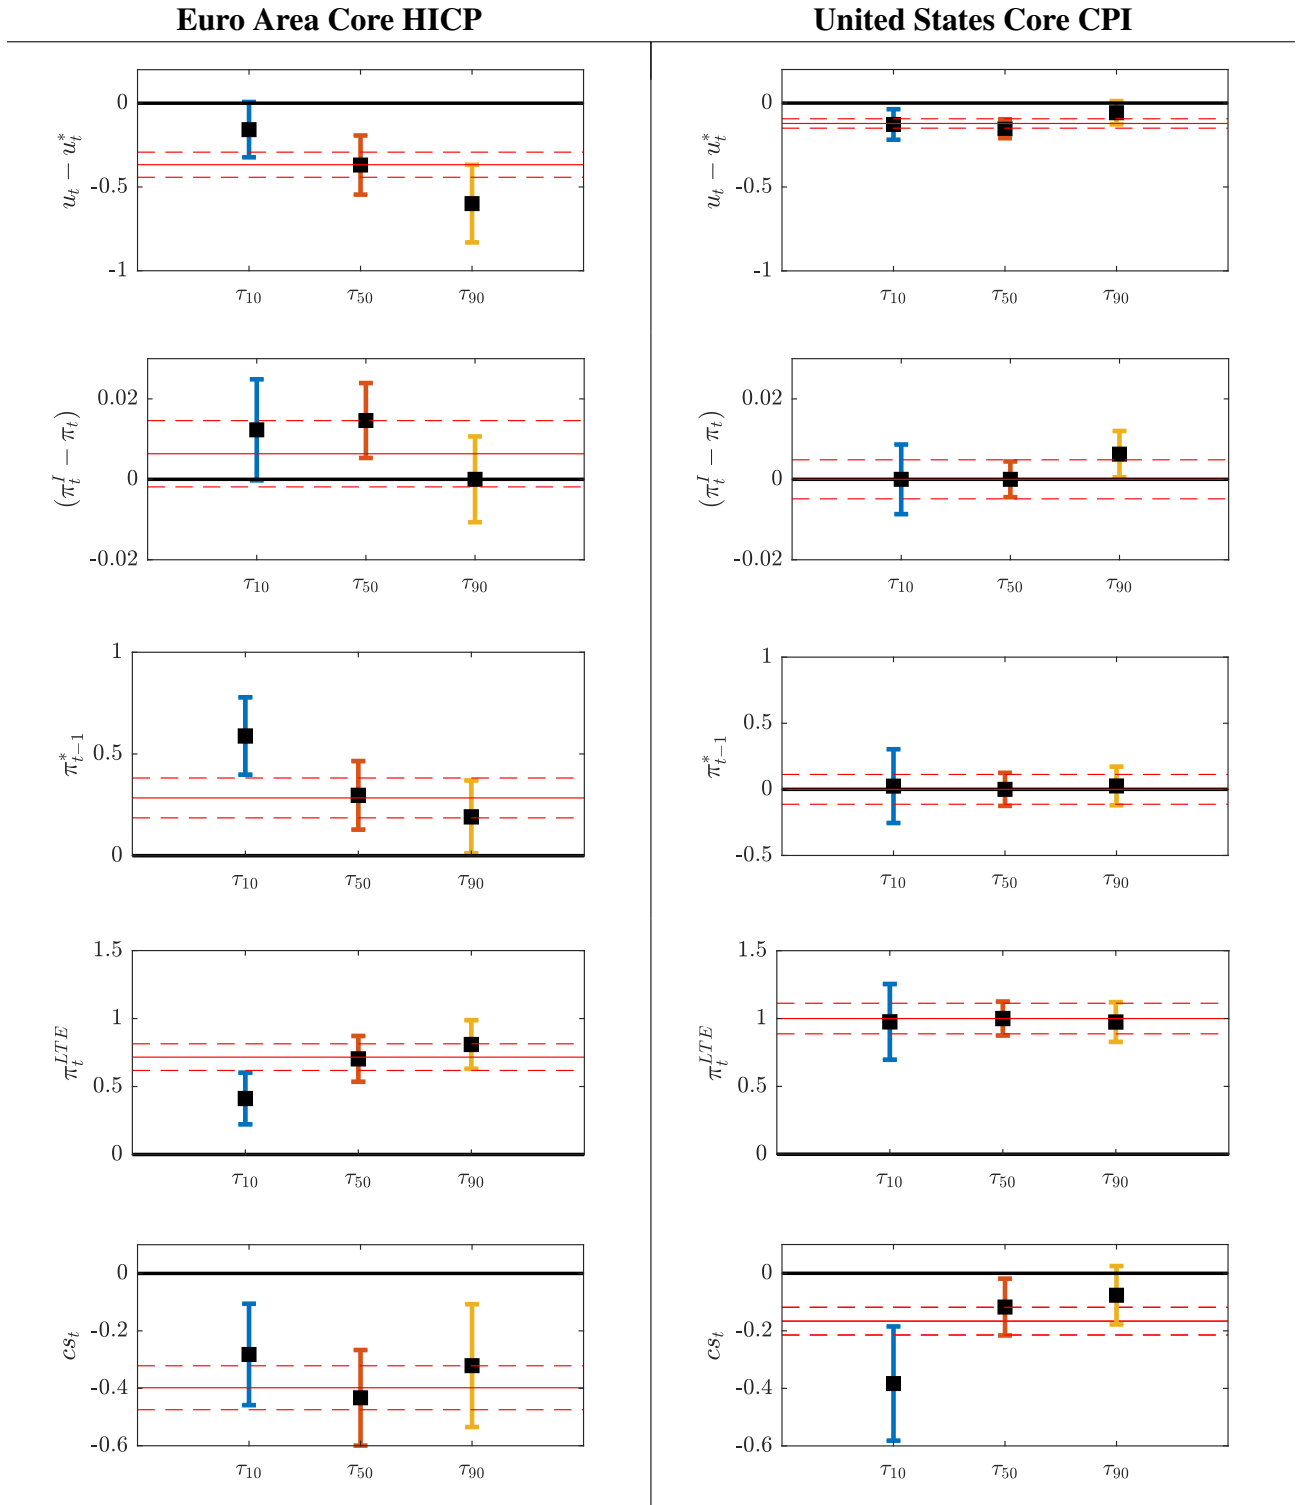
Quantile Phillips Curve Figure 8 displays the estimated slopes of the quantile regression model (2) for the euro area (left column) and the United States (right column). The information is organized as follows. Boxes in each row correspond to the covariates of the model. The black squares report the point estimates of the 10th, 50th, and 90th quantile-specific slopes. The length of the vertical lines around the point estimates corresponds to the 68 percent confidence intervals constructed by “blocks-of-blocks” bootstrap (see Appendix H). The OLS point estimates and their 95 percent confidence intervals are given by the horizontal red lines.

The unemployment gap generates fairly similar responses of median inflation in the euro area and the U.S. but important differences emerge when looking at the tails of inflation. In the euro area the upper tail is the most sensitive segment of the inflation distribution to unemployment, while the lower tail responds little. Thus, the relative odds of high inflation risks arising from a substantially tight labor market outweigh the downside risks of low inflation associated with substantial labor market slack. This pattern is reversed in the U.S., though the degree of asymmetry and the role of unemployment in general is much more muted.

During this period, changes in the relative price of imported goods played a minor role in the U.S., and a slightly larger role in the euro area. Still, there are some interesting differences across economies. In the eurozone, the median and the lower tail of the distribution seem more responsive than the upper tail of the distribution of inflation. For the U.S., these results are consistent with the previous section, in which we pointed to a greatly reduced relevance of relative prices as inflation determinants starting with the Great Moderation.

Longer-term inflation expectations and inflation inertia influence differently the overall inflation distribution in the two countries. While in the U.S. inflation expectations dominate all parts of the inflation distribution, in the euro area they only play a major role in explaining the upper tail of inflation as odds of low inflation are also driven by past inflation. In other words, in the eurozone, unmoored reductions in inflation expectations result in more persistent increases in downside risks as their negative effect is propagated over time through a higher inflation inertia.

Figure 8: Quantile Regression Slopes and Confidence Intervals.



NOTE: The figure displays the slope coefficients of the quantile regression of average four-quarter-ahead euro area core HICP (left) and United States Core CPI inflation (right) defined in (2). The black squares correspond to the point estimates whereas the vertical lines to the 68% confidence intervals computed via “blocks-of-blocks” bootstrap using 10,000 replications (see Appendix H) for the 10th quantile (blue), median (red) and 90th quantile (yellow). The estimation period is 1999:Q1 to 2017:Q4. The OLS estimates and their 95% confidence intervals are respectively represented by the solid and dashed red lines.

The last row in Figure 8 presents the role of credit spreads across inflation quantiles. In the euro area higher credit spreads (i.e., tighter credit conditions) shift the inflation distribution to the left as they have a fairly symmetric negative effect across inflation quantiles. This contrasts with the U.S. in which most of the reduction in inflation following high spreads is concentrated in the lowest tail, while the effects on the upper tail are not very significant.¹⁶ Most importantly, in the U.S. financial conditions are the only significant source of asymmetry in the inflation distribution.

We performed ANOVA tests on the equality of the slope coefficients across quantiles. The test rejected the null hypothesis of equality between the slopes on the 10th and 50th quantile at 1% significance level for average past inflation in the Euro Area and for credit spreads in the United States, thus confirming the importance of inflation inertia for the euro area and of credit conditions for the U.S. in characterizing downside risks to and left-skewness in the inflation distribution of the respective countries.¹⁷

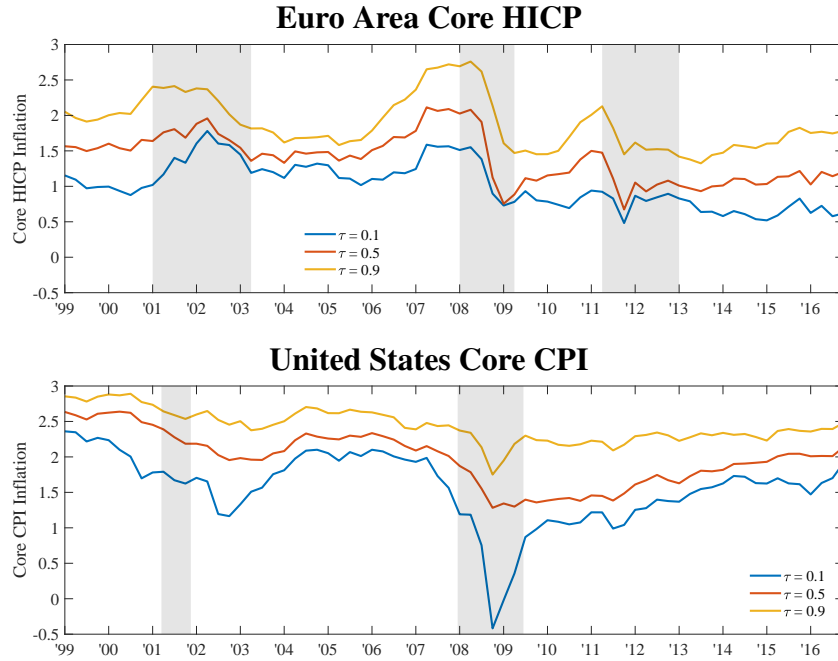
Inflation Quantiles Figure 9 highlights key aspects of the evolution of the inflation outlook by displaying the time series of the median, the 10th and the 90th inflation quantiles. The top panel shows the evolution for the euro area while the bottom panel focuses on the United States.

In the eurozone, looking at the lower tail, it appears that downside inflation risks have been important since the inception of the euro. Strikingly, the inflation distribution tends to tilt to the upside around the three recessionary episodes, while also widening up significantly. By the end of the 2001-2003 recession, downside risks to inflation started to trend down (i.e., the lower tail fell) and after a faint recovery subsequently failed to rebound to the pre-contraction level. This phenomenon is observed during the global financial crisis of 2008-2009 and repeated around the 2011-2013 recession, when downside risks increased further without recovering since then.

¹⁶This result is also robust to limiting the sample to 1995:Q1-2007:Q4 (i.e., discarding the Great Recession) as well as to using the short-term instead of the long-term CBO NAIRU measure for u_t^* (the former is higher after the Great Recession). Further, in Appendix I we extend these results to two alternative measures of inflation, core PCE and the [Stock and Watson \(2019\)](#) Cyclically Sensitive Inflation index. The effects are more symmetric in the case of core PCE, while the CSI measure exhibits a similar asymmetry as core CPI although of somewhat larger magnitude.

¹⁷The slopes for the 90th and 50th quantile are statistically different only for the U.S. unemployment rate.

Figure 9: Time Evolution of Selected Conditional Inflation Quantiles.



NOTE: The figure displays the time evolution of the conditional quantiles of euro area core HICP inflation (top) and of United States core CPI inflation (bottom) estimated from the quantile regression defined in (2). Shaded bars indicate NBER-dated recessions for the United States and OECD-based recession indicators for the euro area.

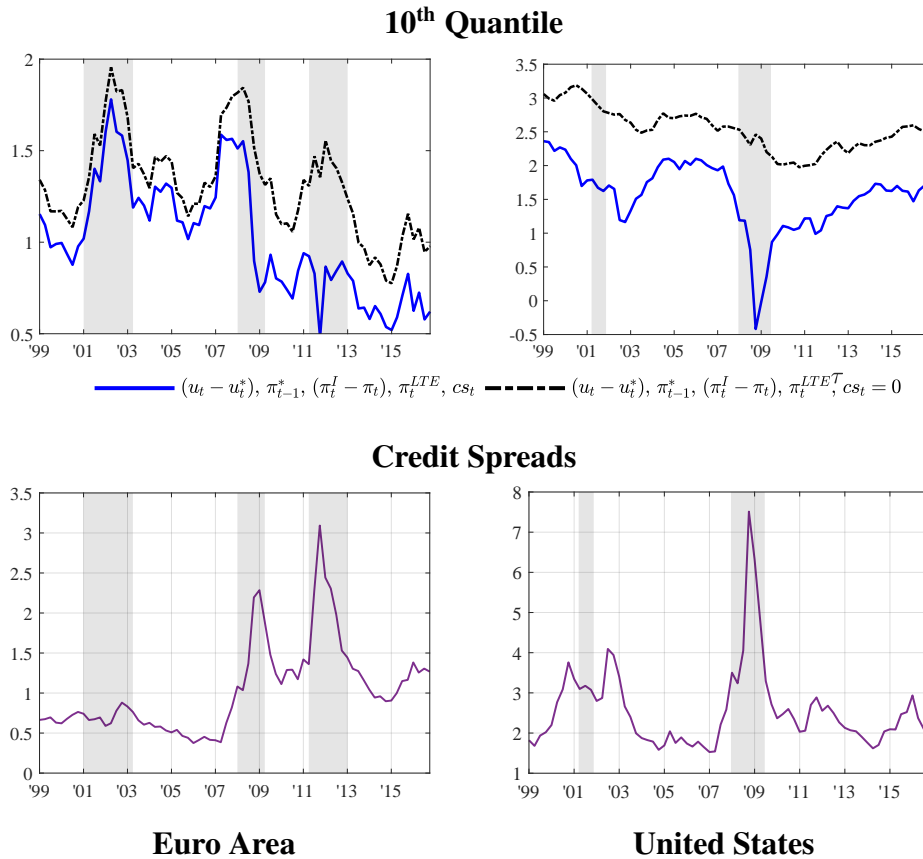
The estimated quantiles for the U.S. in the bottom panel of Figure 9 show some salient differences with the eurozone. First, the downward tilts to the distribution associated with the two recessions were primarily a result of a drop in the left tail, unlike in the euro area where the downward push was more pronounced for the median and the upper tail. This was particularly acute during the global financial crisis. However, the substantial increase in the odds of low inflation was followed by a sustained recovery until the distribution became again tightly centered slightly above 2 percent with the 10th quantile moving back to close to 2 percent. This contrasts the eurozone experience, in which the left tail failed to recover after the global financial crisis. In Appendix J we complement these results by exploring which factors contributed to the recovery of the left tail in the United States and to the failed recovery of the left tail in the euro area.

The Role of Credit Spreads We now turn our attention to the increasing role played by changes in credit conditions in influencing the downside risks of inflation throughout the recovery. In Figure

10, we illustrate the time evolution of the 10th conditional inflation quantile of euro area core HICP inflation (left) and of United States core CPI inflation (right), both for the baseline version of the model (blue solid) and for a version where the effect of credit spreads is set to zero (black dashed).

The gap between the two lines captures the partial effect of credit spreads on the 10th quantile. It is evident how tighter financial conditions exert a persistent downward pressure on downside inflation risks and more strongly so, when credit spreads are high. It is remarkable how the large spike in credit spreads observed in 2008 in the U.S. (bottom right panel of Figure 10) pushed down the lower inflation quantile, which slowly moved back to about 2 percent by the end of 2016.

Figure 10: Partial Effect of Credit Spread on 10th Inflation Quantile.



NOTE: The figure displays the evolution of the 10th conditional inflation quantile of average one-year-ahead euro area core HICP inflation (left) and of U.S. core CPI inflation (right) estimated from the quantile regressions model (2), in its baseline version (blue solid) and in its version in which the effect of credit spreads is set to zero (black dash-dotted). Shaded bars indicate NBER-dated recessions for the U.S. and OECD-based recession indicators for the euro area.

The eurozone is a slightly different story. Financial conditions, which played a more limited role in the lower tail inflation dynamics, became increasingly important after the global financial

crisis. To see this, let us focus on the bottom left panel of Figure 10. This figure clearly shows that the tightening in credit conditions occurred in two consecutive waves. The initial tightening in financial conditions happened around 2008 and 2009 and marked a sharp reduction in the lower quantile of the distribution that, even after some recovery in credit conditions during 2009, would never rebound. The second wave, linked to the European sovereign debt crisis, exacerbated this change. As of 2012, the deterioration in credit conditions lifted up substantially the odds of low inflation. It is remarkable how, early in 2013, economic conditions pointed to a recovery in the lower quantile. According to our model, however, this would have portrayed a misleading picture reflecting the lack of consideration of the substantial downward pressures in place originated by the still very tight credit conditions at that time. To see this more clearly, we translate the variation in these quantiles into changes of the entire distribution of inflation, to which we turn next.

The Distribution of Inflation At a speech in London in July 2012, Mario Draghi – President of the European Central Bank from November 2011 to October 2019 – announced the ECB’s commitment of doing “whatever it takes” to preserve the euro. The eurozone was in the throes of crisis, bond yields and credit spreads of weak euro-member governments were soaring, and financial markets doubted that European institutions could avert disaster. This is part of the historical context reflected in Figure 11, which plots the estimated euro area core HICP predictive inflation distributions (left column) and their associated quantile functions (right column) across four selected dates (for details on the construction of these objects please refer back to Section 1.2). In Appendix K we also report the inflation probabilities associated with these distributions. We start at the dawn of the global financial crisis (2007:Q4), then explore those periods in which downside risks were most acute (2008:Q4 and 2011:Q4) and finally zoom in the end of our sample (2016:Q4). The blue solid lines correspond to the baseline quantile model whereas the black dash-dotted lines parse out the effect of credit spreads.

The results reaffirm that tight financial concerns played a crucial role in shifting to the left the inflation distribution during the end of 2008 and remarkably so in the last quarter of 2011 – a few

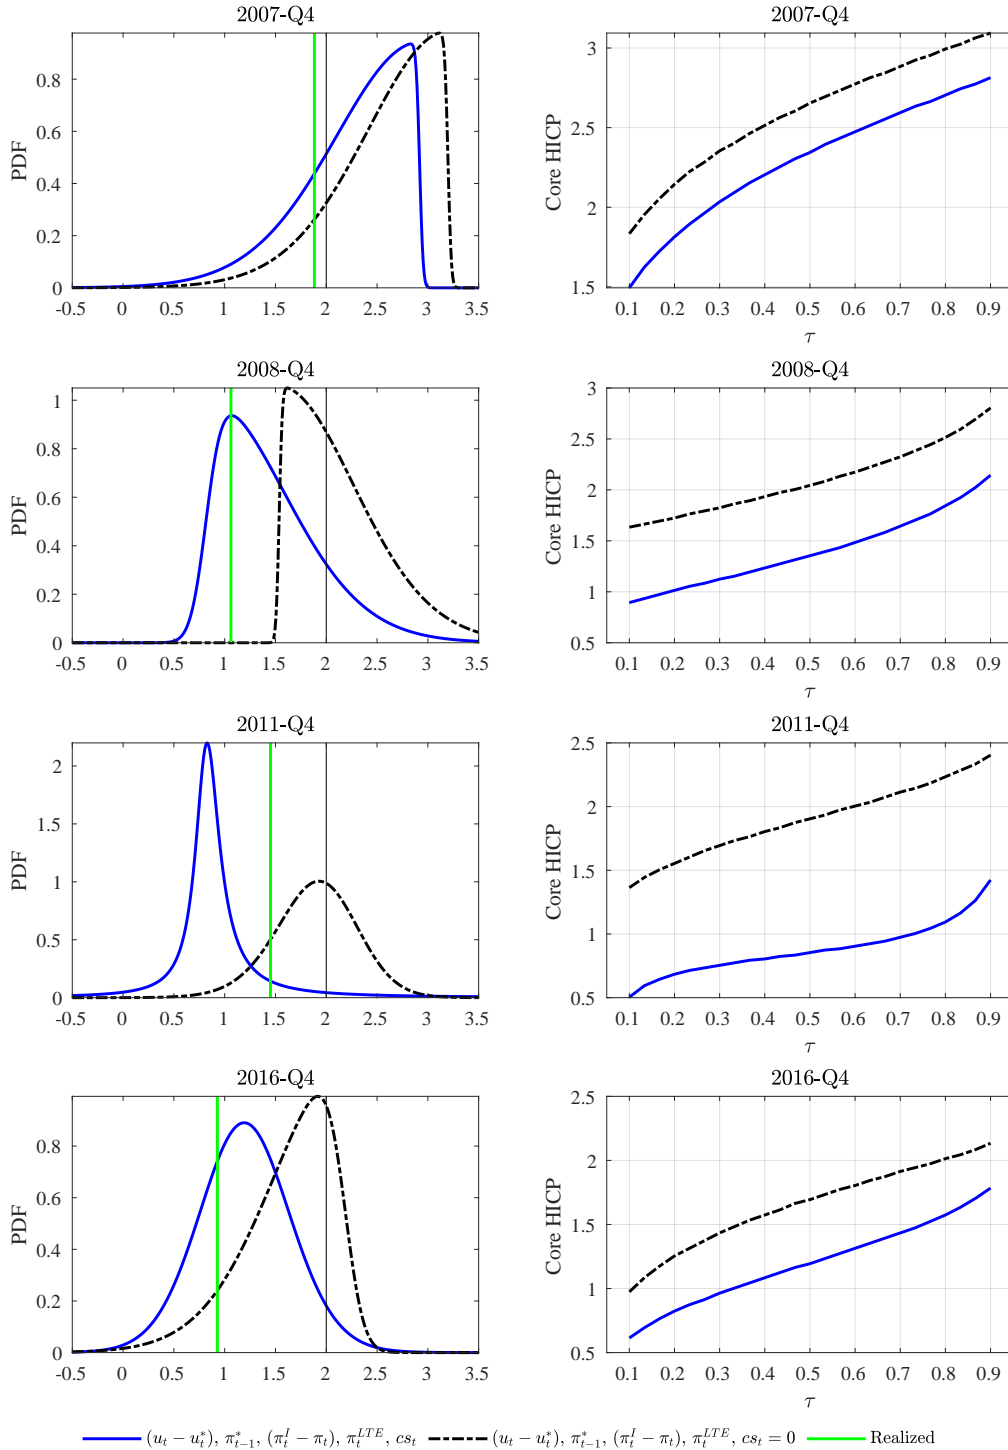
months before Draghi’s speech. It is noteworthy how by the end of 2007, inflation-at-risk above 3 percent was virtually non-existent, while the left-tail pointed to some minor downside risks of inflation running below 1 percent. Overall, credit conditions barely affected these conclusions. The two waves in which the financial crisis was reflected in tight credit conditions translated into a remarkable change of the inflation outlook. The distribution shifted to the left and concentrated around a median inflation little below 1 percent, with the odds of low inflation – or even deflation – soaring. By the end of 2016 the distribution of inflation had fatter tails, with the odds of high inflation above those observed at the peak of the crisis, but with substantial downside risks still remaining. The effects of credit conditions on inflation are also illustrated in the right column of Figure 11, which shows that the inflation quantiles which condition on credit spreads were significantly below those that solely rely on economic factors.

In Figure 12, we compare these results with the experience of the United States, for which we consider similar dates.¹⁸ We focus on 2008:Q4 as this is when downside risks were most pronounced in the U.S. following the sharp rise in credit spreads and dire economic conditions. As in the case of the eurozone, the effect of financial variables on the inflation distribution is striking during those episodes in which financial distress was most acute. In 2008:Q4, for example, tighter credit conditions contributed to pushing the entire inflation distribution to the left, while making it more dispersed and poking down substantially its left tail to the point of placing non-zero probability of deflation occurring on average within the next year (as we will show below).

Looking at the right columns of Figures 11 and 12, one important difference emerges between the euro area and the U.S. experience – a difference we had already encountered when analyzing the quantile-specific slopes of credit spreads on average future inflation in Figure 8. While in the eurozone higher credit spreads pushed down the inflation distribution symmetrically across quantiles, in the United States its effects were mostly reflected in the left tail. As we show below, the translation of these effects into the probability of low inflation (i.e., downside inflation-at-risk) is more pronounced, the more the inflation distribution is right-skewed (i.e., the fatter its left tail).

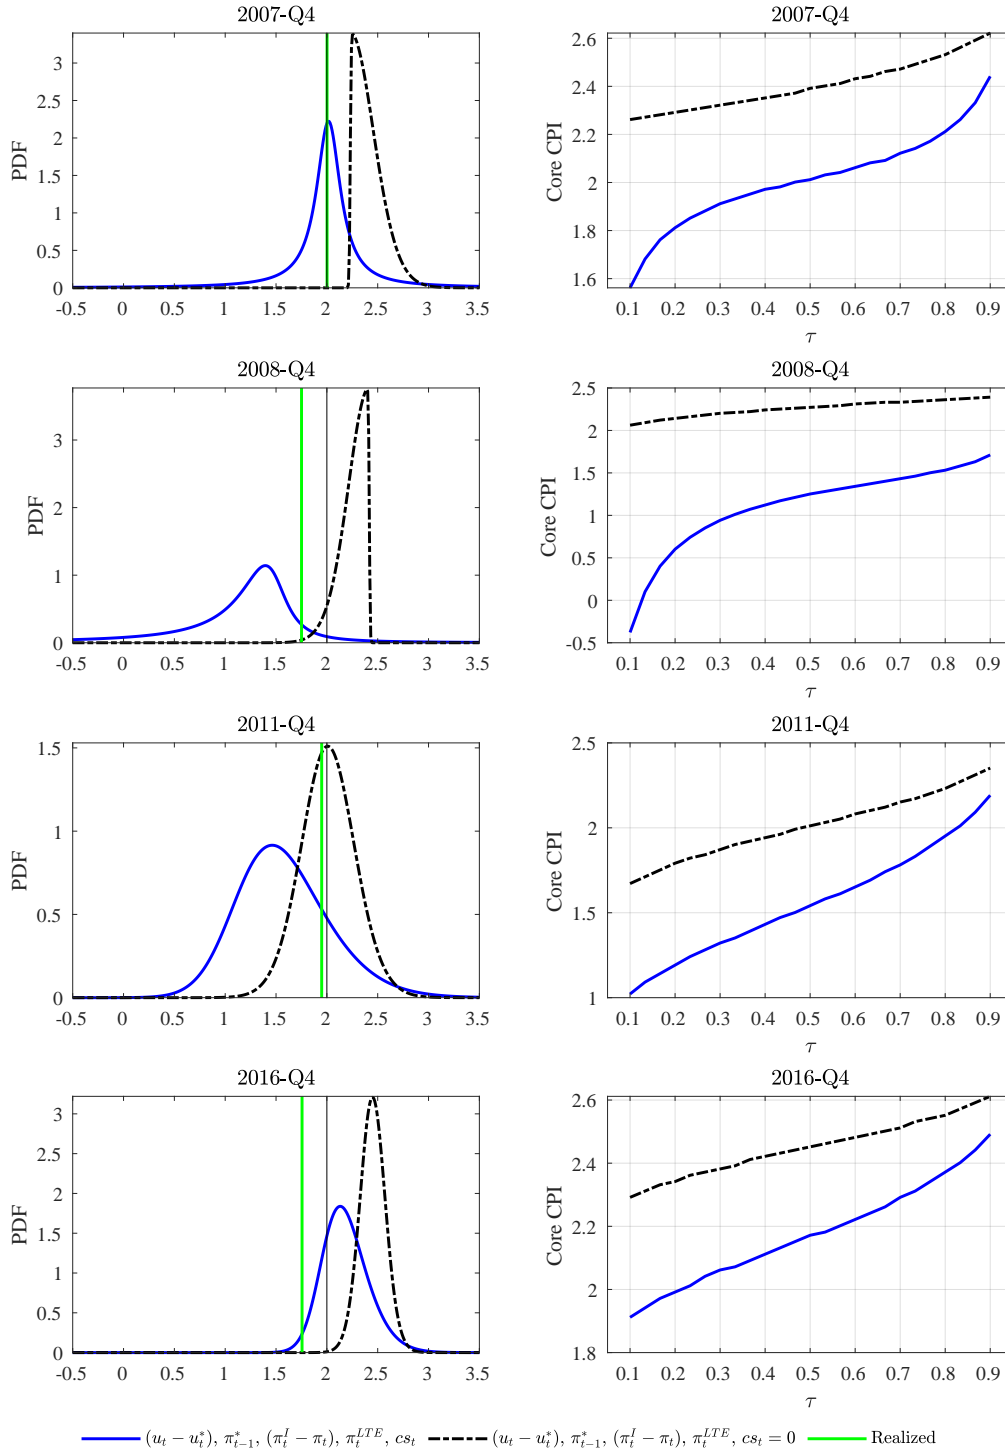
¹⁸This chart is similar to Figure 4 in Section 3.2, the only difference being that the densities have been computed using the same sample as for the Euro Area.

**Figure 11: Selected Time Episodes of Predictive Densities (Left) and Skewness (Right).
Euro Area Core HICP Inflation.**



NOTE: The left panels show the estimated skewed t -Student densities of average four-quarter-ahead euro area core HICP inflation, in its baseline version (blue solid) and in its version where the effect of credit spreads is set to zero (black dash-dotted). The same panel also reports the realized value of average four-quarter-ahead United States core PCE inflation (green vertical line) and the density of SPF forecasts of year-on-year core HICP inflation (red dotted) computed using a kernel smoothing function. The SPF density is only available in 2016-Q4 as that is the first date for which the SPF forecasts are available and the last data point in our sample. The right panels show the estimated skewed t -inverse cumulative associated with the conditional densities in the left panels.

**Figure 12: Selected Time Episodes of Predictive Densities (Left) and Skewness (Right).
United States Core CPI Inflation.**



NOTE: The left panels show the estimated skewed t -Student densities of average four-quarter-ahead United States core CPI inflation, in its baseline version (blue solid) and in its version where the effect of credit spreads is set to zero (black dash-dotted). The same panel also reports the realized value of average four-quarter-ahead United States core CPI inflation (green vertical line) and the density of SPF forecasts of year-on-year core CPI inflation (red dotted) computed using a kernel smoothing function. The right panels show the estimated skewed t -inverse cumulative associated with the conditional densities in the left panels.

5 Tracking Risks During the Covid-19 Crisis

Risks to the inflation outlook have been front and center at the peak and in the recovery from the Covid-19 crisis. We show how our model, augmented by credit spreads, allows to identify important changes in the inflation distribution during this historical episode.

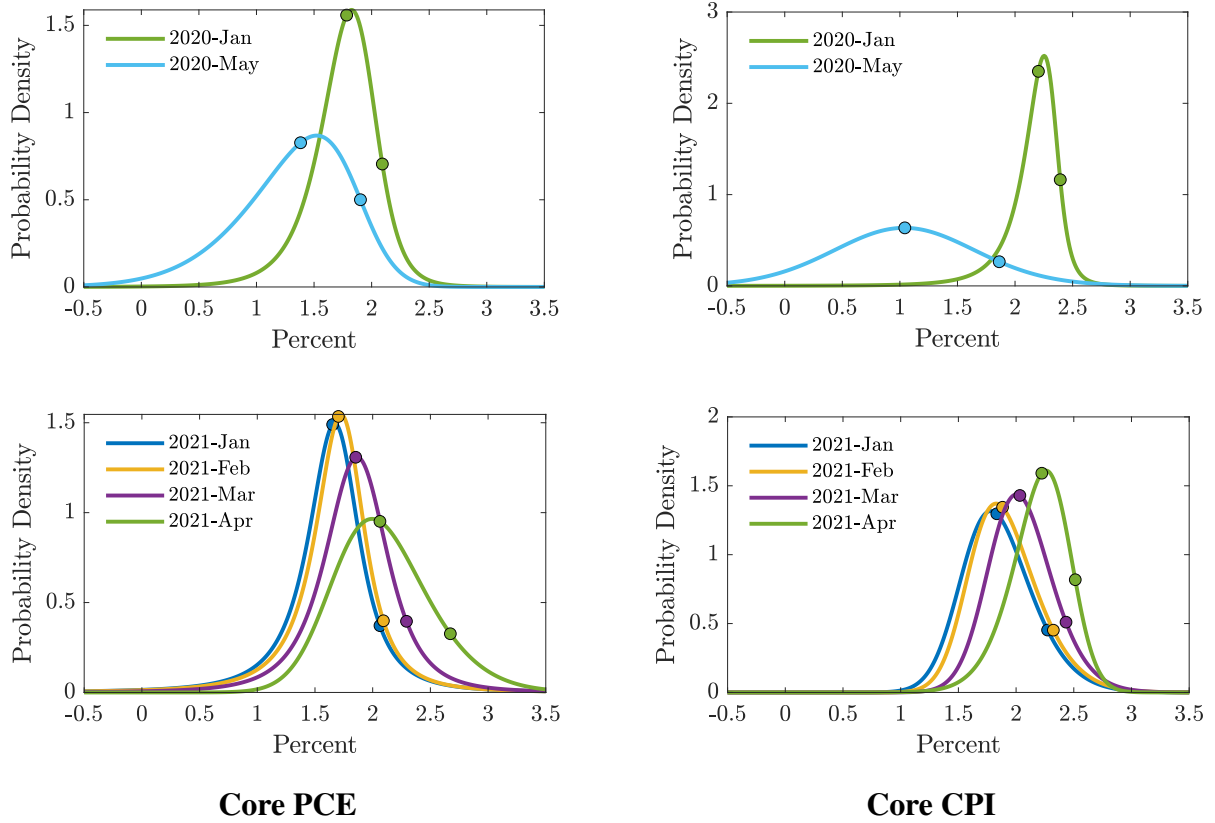
In doing this, we consider a monthly version of our model in which we use for past inflation π_t^* the average inflation rate between t and $t - 11$ when forecasting one-year-ahead inflation $\bar{\pi}_{t+1,t+12}$ as of time t . This is to allow to a more “real-time” assessment of inflation risks that incorporates information on the current period inflation rate.¹⁹ The model is estimated from January 2000 to April 2020, using as a dependent variable one-year-ahead inflation up to April 2021.

Figure 13 shows results for average inflation over the next 12 months, with core PCE inflation on the left and core CPI inflation on the right. Specifically, it shows the distributions since January 2020, with markers given at the median and at the 90th quantile. Both core PCE and core CPI inflation distributions have shifted to the right since the beginning of this year. Indeed, the selected months show the quick build-up of downside risks to inflation during the harsh initial months of the global pandemic (January to May 2020, top panels) which were then followed by the recent increase in upside risks to inflation of the recent months (bottom panels).

We now ask what would have been the predictive distributions over the next 12 months had the financial conditions deterioration of March 2020 persisted in May 2021. Figure 14 presents distributions of average inflation over the next 12 months in May 2021 for our baseline (blue solid lines) and in a counterfactual scenario in which credit conditions are those prevailing in March 2020 (black dashed lines). The easing of financial conditions since the onset of the pandemic has moved the distributions of both core PCE and core CPI inflation rates to the right reducing downside risks of low inflation. Indeed, had the financial conditions not eased during the period from their May 2020 state, the probability of inflation running at or above 2.5 percent would be about 7 percentage points lower for both core PCE and core CPI inflation rates as of May 2021.

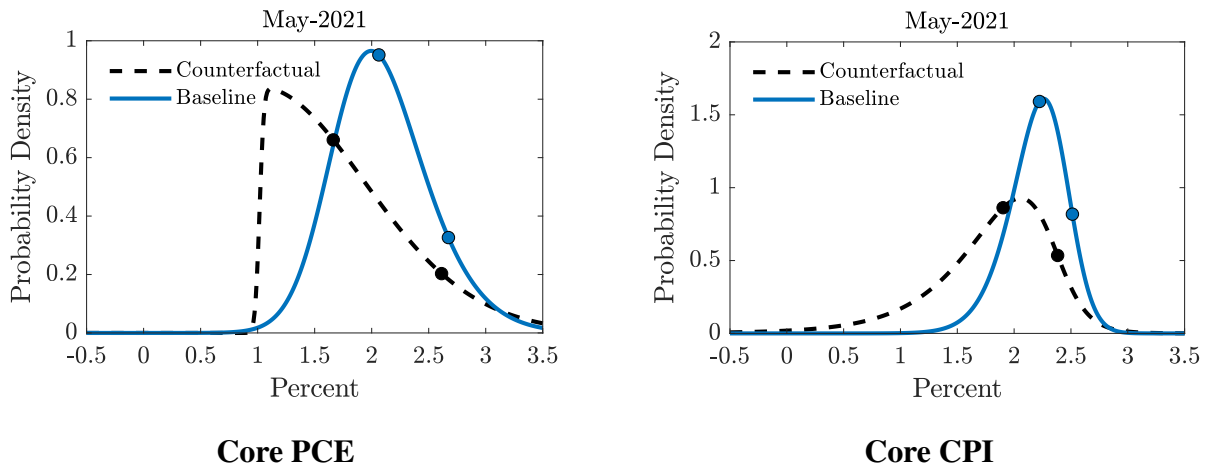
¹⁹In our baseline model, we do not assume that the researcher knows inflation as of time t when forecasting future inflation, due to lags in the releases of inflation data.

Figure 13: Predictive Densities of One-Year-Ahead Inflation Measures During Covid-19.



NOTE: Predictive densities of inflation over the next 12 months in selected episodes. Onset of the COVID-19 pandemic (top panel) and 2021 (bottom panel), for core PCE inflation (left) and core CPI inflation (right). Markers appear at the median and at the 90th quantile.

Figure 14: Predictive Densities of One-Year-Ahead Inflation Measures in April 2021. The Role of Credit Spreads.



NOTE: Predictive densities of inflation over the next 12 months in May 2021. The baseline case is in solid blue lines and the counterfactual case (in which credit spreads are at March 2020) levels in dotted black lines, for core PCE inflation (left) and core CPI inflation (right). Markers are at the median and at the 90th quantile.

6 Conclusion

In this paper we show that the recent muted response of the *conditional mean* of inflation to economic conditions does not necessarily convey an adequate representation of inflation dynamics. Indeed, we find ample variability in the tail risks to inflation, even when focusing on the post-2000 period of stable and low *mean* inflation. In particular, we document that tight financial conditions generated times of substantial and persistent downside risks to inflation. We also show that evidence from financial market quotes, from SPF forecasts and from a regime-switching model of inflation is consistent with the previous findings.

Our paper offers a new empirical perspective to macroeconomic modelers and to policymakers, showing that changes in credit conditions are key to understand tail-risk dynamics of inflation. Our results provide empirical guidance and suggest more efforts in modeling the linkages between the entire inflation distribution and financial markets in the context of nonlinear models.

As a future research avenue, we encourage the exploration of whether credit spreads in specific sectors make the inflation outlook particularly vulnerable and whether certain inflation components are particularly sensitive to movements in financial markets.

References

- Adams, P. A., T. Adrian, N. Boyarchenko, and D. Giannone (2021). Forecasting Macroeconomic Risks. *International Journal of Forecasting* 37(3), 1173–1191.
- Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable Growth. *American Economic Review* 109(4), 1263–89.
- Andrade, P., E. Ghysels, and J. Idier (2012). Tails of Inflation Forecasts and Tales of Monetary Policy. Working Papers, Banque de France.
- Azzalini, A. and A. Capitanio (2003). Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew t -distribution. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65, 367–389.
- Ball, L. and S. Mazumder (2019). A Phillips Curve with Anchored Expectations and Short-Term Unemployment. *Journal of Money, Credit and Banking* 51(1), 111–137.
- Banerjee, R. N., J. Contreras, A. Mehrotra, and F. Zampolli (2020). Inflation at Risk in Advanced and Emerging Economies. BIS Working Papers 883, Bank for International Settlements.
- Baumeister, C. and L. Kilian (2016). A General Approach to Recovering Market Expectations from Futures Prices with an Application to Crude Oil.
- Berkowitz, J., I. Biegean, and L. Kilian (1999). On the Finite-Sample Accuracy of Nonparametric Resampling Algorithms for Economic Time Series.
- Blanchard, O., E. Cerutti, and L. H. Summers (2015). Inflation and Activity: Two Explorations and Their Monetary Policy Implications. Working Paper Series WP15-19, Peterson Institute for International Economics.
- Boons, M., F. Duarte, F. de Roon, and M. Szymanowska (2020). Time-Varying Inflation Risk and Stock Returns. *Journal of Financial Economics* 136(2), 444–470.

- Busetti, F., M. Caivano, D. D. Monache, and C. Pacella (2019). The Time-Varying Risk of Italian GDP. Mimeo, Bank of Italy.
- Busetti, F., M. Caivano, and L. Rodano (2015). On the Conditional Distribution of Euro Area Inflation Forecast.
- Caldara, D., D. Cascaldi-Garcia, P. Cuba Borda, and F. Loria (2020). Understanding Growth-at-Risk: A Markov-Switching Approach. *mimeo*.
- Cecchetti, S. G. (2008). Measuring the Macroeconomic Risks Posed by Asset Price Booms. *Asset Prices and Monetary Policy*, 9–43.
- Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2015). Understanding the Great Recession. *American Economic Journal: Macroeconomics* 7(1).
- Christiano, L. J., R. Motto, and M. Rostagno (2014). Risk Shocks. *American Economic Review* 104(1), 27–65.
- Del Negro, M., F. Bassetti, and R. Casarin (2020). A Bayesian Approach for Inference on Probabilistic Surveys. *mimeo*.
- Del Negro, M., M. P. Giannoni, and F. Schorfheide (2015). Inflation in the Great Recession and New Keynesian Models. *American Economic Journal: Macroeconomics* 7(1), 168–96.
- Favara, G., S. Gilchrist, K. F. Lewis, and E. Zakrajšek (2016). Updating the Recession Risk and the Excess Bond Premium. Working paper series, FEDS Notes. Washington: Board of Governors of the Federal Reserve System.
- Galvão, A. B. and M. T. Owyang (2018). Financial Stress Regimes and the Macroeconomy. *Journal of Money, Credit and Banking* 50(7), 1479–1505.
- Ganics, G., B. Rossi, and T. Sekhposyan (2020). From Fixed-Event to Fixed-Horizon Density Forecasts: Obtaining Measures of Multi-horizon Uncertainty from Survey Density Forecasts. CEPR Discussion Papers 14267, C.E.P.R. Discussion Papers.

- Gertler, M., N. Kiyotaki, and A. Prestipino (2019). A Macroeconomic Model with Financial Panics. *The Review of Economic Studies* 87(1), 240–288.
- Ghysels, E., L. Iania, and J. Striaukas (2018). Quantile-Based Inflation Risk Models.
- Gilchrist, S. and B. Mojon (2018). Credit Risk in the Euro Area. *Economic Journal* 128(608), 118–158.
- Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek (2017). Inflation Dynamics during the Financial Crisis. *American Economic Review* 107(3), 785–823.
- Gilchrist, S. and E. Zakrajšek (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review* 102(4).
- Kiley, M. T. (2018). Unemployment Risk. Technical report.
- Kilian, L. and H. Lütkepohl (2018). *Structural Vector Autoregressive Analysis*. Number 9781107196575 in Cambridge Books. Cambridge University Press.
- Kilian, L. and S. Manganelli (2007). Quantifying the Risk of Deflation. *Journal of Money, Credit and Banking* 39(2-3), 561–590.
- Kilian, L. and S. Manganelli (2008). The Central Banker as a Risk Manager: Estimating the Federal Reserve’s Preferences under Greenspan. *Journal of Money, Credit and Banking* 40(6), 1103–1129.
- Kitsul, Y. and J. H. Wright (2013). The Economics of Options-Implied Inflation Probability Density Functions. *Journal of Financial Economics* 110(3), 696–711.
- Koenker, R. (2005). *Quantile Regression*. Econometric Society Monographs. Cambridge University Press.
- Koenker, R. and P. Ng (2005). Inequality Constrained Quantile Regression. *Sankhyā: The Indian Journal of Statistics (2003-2007)* 67(2), 418–440.

- Korobilis, D. (2017). Quantile Regression Forecasts of Inflation Under Model Uncertainty. *International Journal of Forecasting* 33(1), 11 – 20.
- Korobilis, D., B. Landau, A. Musso, and A. Phella (2021). The Time-Varying Evolution of Inflation Risks. *mimeo*.
- Loria, F., C. Matthes, and D. Zhang (2019). Assessing Macroeconomic Tail Risk. Finance and Economics Discussion Series 2019-026, Board of Governors of the Federal Reserve System (US).
- Maih, J. (2015). Efficient Perturbation Methods for Solving Regime-Switching DSGE Models. Working Paper 2015/01, Norges Bank.
- Manzan, S. and D. Zerom (2013). Are Macroeconomic Variables Useful for Forecasting the Distribution of U.S. Inflation? *International Journal of Forecasting* 29(3), 469 – 478.
- McLeay, M. and S. Tenreyro (2018). Optimal Inflation and the Identification of the Phillips Curve.
- Rossi, B. (2014). Density Forecasts in Economics and Policymaking. *Els Opuscles del CREI* (37).
- Rossi, B. and T. Sekhposyan (2019). Alternative Tests for Correct Specification of Conditional Predictive Densities. *Journal of Econometrics* 208(2), 638 – 657.
- Stock, J. H. and M. W. Watson (2019). Slack and Cyclically Sensitive Inflation. Working Paper 25987, National Bureau of Economic Research.
- Tagliabracci, A. (2018). Essays on Macroeconomic Forecasting. Doctoral thesis, Universitat Autònoma de Barcelona.
- Villemot, S. . and J. Pfeifer (2017). BVAR Models “à la SIMS” in Dynare. *mimeo*.
- Williams, J. (2010). Sailing Into Headwinds: The Uncertain Outlook for the U.S. Economy. Speech 85, Federal Reserve Bank of San Francisco.

Wolters, M. H. and P. Tillmann (2015). The Changing Dynamics of US Inflation Persistence: A Quantile Regression Approach. *Studies in Nonlinear Dynamics & Econometrics* 19(2), 161–182.

Yellen, J. (2013). Panel Discussion on “Monetary Policy: Many Targets, Many Instruments. Where Do We Stand?”. Speech 622, Federal Reserve Board.

A Data Appendix

In this section we provide details on the data for the United States and the euro area.

A.1 United States

- Core Consumer Price Index
 - Source: FRED.
 - Details: CPILFESLPCA, Consumer Price Index for All Urban Consumers: All Items Less Food and Energy, Compounded Annual Rate of Change, Quarterly, Seasonally Adjusted.
- [Stock and Watson \(2019\)](#) Cyclically Sensitive Inflation
 - Source: Replication Material of [Stock and Watson \(2019\)](#) on [Professor Mark Watson's Website](#).
 - Details: Quarterly CSI inflation rates.
- Core Personal Consumption Expenditures
 - Source: FRED.
 - Details: PCEPILFELPCA, Personal Consumption Expenditures Excluding Food and Energy (Chain-Type Price Index), Compounded Annual Rate of Change, Quarterly, Seasonally Adjusted.
- Long-Term Inflation Expectations
 - Source: [Blanchard, Cerutti, and Summers \(2015\)](#) and Consensus Economics (provided by the Prices and Wages section of the Research & Statistics Division at the Federal Reserve Board).

- Details: From 1990:Q4 onwards, we use six- to-ten-year-ahead mean CPI inflation forecasts from semiannual surveys from Consensus Economics. The inflation expectations data from Consensus Economics are available semi-annually in April and October. We impute the same inflation expectations value as the last available observation for the missing months in between the data releases (e.g., if in 1991 October the value is 2.5 and in 1991 April the value is 2, then we impute a value of 2.5 between 1991 May and 1991 October and a value of 2 between November 1990 and April 1991). Before 1990:Q4 that date we use the series from [Blanchard, Cerutti, and Summers \(2015\)](#).
 - Transformations: In the quarterly version of the model we average the monthly observations within a quarter. In the monthly version of the model we interpolate the quarterly series.
- Unemployment Rate
 - Source: FRED.
 - Details: UNRATE, Civilian Unemployment Rate, Percent, Quarterly, Seasonally Adjusted.
 - Natural Rate of Unemployment
 - Source: FRED.
 - Details: NROU, Natural Rate of Unemployment (Long-Term), Percent, Quarterly, Not Seasonally Adjusted.
 - Import Price Index
 - Source: FRED.
 - Details: B021RG3Q086SBEA_CCA, Imports of goods and services (chain-type price index), Solidly Compounded Annual Rate of Change, Quarterly, Seasonally Adjusted.
 - Oil Price

- Source: FRED.
- Details: WTISPLC_CCA, Spot Crude Oil Price: West Texas Intermediate (WTI), Solidly Compounded Annual Rate of Change, Quarterly, Not Seasonally Adjusted.
- [Gilchrist and Zakrajšek \(2012\)](#) Credit Spread and Excess Bond Premium
 - Source: Data regularly updated in [FEDS Note](#) by [Favara, Gilchrist, Lewis, and Zakrajšek \(2016\)](#).
 - Details: Credit spread and excess bond premium as constructed by [Gilchrist and Zakrajšek \(2012\)](#).
- Corporate Bond Spread
 - Source: FRB/US model package available at this [Federal Reserve Board website](#).
 - Details: RBBB minus RG10. RBBB, yield on BBB-rated corporate bonds. RG10, yield on 10-year Treasury security.
- National Financial Conditions Index
 - Source: FRED.
 - Details: NFCCI, Chicago Fed National Financial Conditions Index, Index, Quarterly, Not Seasonally Adjusted.
- Inflation Probabilities from Financial Markets
 - Source: Provided by the Monetary and Financial Markets Analysis Section of the Monetary Affairs Division of the Federal Reserve Board.
 - Details: Probabilities are inferred from inflation caps and floors contracts as in [Kitsul and Wright \(2013\)](#).
- Survey of Professional Forecasters (SPF) Core CPI and Core PCE Inflation Forecasts

- Source: Philadelphia FED.
 - Details: [Philadelphia FED Website](#) (Mnemonic: CORECPI and COREPCE). Q/Q Rate of Change in the Quarterly-Average Core CPI Level (annualized percentage points).
- Oil Price Surprises
 - Source: Downloaded from [Professor Christiane Baumeister's Website](#).
 - Details: Probabilities are inferred from inflation caps and floors contracts as in [Baumeister and Kilian \(2016\)](#).

Figure A-1: Inflation Measures, United States.





Figure A-2: Regressors, United States.



A.2 Euro Area

- Harmonized Index of Consumer Prices
 - Source: Statistical Office of the European Communities and Haver Analytics (provided by the Advanced Foreign Economies section of the International Finance Division at the Federal Reserve Board).
 - Details: EA19, Total excluding energy, food, alcohol and tobacco. Quarter-over-quarter annualized growth rates, seasonally adjusted.

- Long-Term Inflation Expectations
 - Source: Consensus Economics (provided by the Advanced Foreign Economies section of the International Finance Division at the Federal Reserve Board).
 - Details: Six- to-ten-year-ahead mean CPI inflation forecasts from semiannual surveys from Consensus Economics. France and Germany.

- Unemployment Rate
 - Source: The Area-wide Model (AWM) database.
 - Details: Unemployment Rate, Percentage of civilian work force, Total (all ages), Total (male and female), Seasonally adjusted, but not work ing day adjusted data.

- Natural Rate of Unemployment
 - Source: Authors' Calculations.
 - Details: HP-filtered trend (with smoothing parameter $\lambda = 1600$) of unemployment rate.

- Import Prices
 - Source: The Area-wide Model (AWM) database.

- Details: Imports of Goods and Services Deflator, Index, Index base year 1995 (1995 = 1). Defined as the ratio of nominal, and real imports of goods and services. Based on the gross concept, i.e. both extra- and intra- area trade flows are accounted for.

- Oil Prices
 - Source: The Area-wide Model (AWM) database.
 - Details: Oil Prices, United Kingdom, Petroleum: UK Brent, US dollars per barrel.

- Credit Spread
 - Source: Data regularly updated at this [Banque de France website](#).
 - Details: euro area bank credit spreads from [Gilchrist and Mojon \(2018\)](#).

- OECD Recession Dates
 - Source: FRED.
 - Details: EUROREC, OECD based Recession Indicators for euro area from the Period following the Peak through the Trough, +1 or 0, Quarterly, Not Seasonally Adjusted.

Figure A-3: Inflation Measures, Euro Area.





Figure A-4: Regressors, Euro Area.



B Inflation-at-Risk using Full-Sample Estimates

In this section, we first present full-sample estimates of the Phillips-curve quantile model to gauge the influence of the inflation drivers on the tails of the conditional inflation distribution. This naturally leads us to consider the presence (or not) of non-linear inflation dynamics, where the non-linearity is intended as arising from the asymmetry in the importance of inflation determinants across quantiles. We close the section by providing estimates of the conditional distribution of average future inflation.

Our measure of inflation is “core inflation”. This measure provides information about the rate toward which headline inflation will converge in the medium term if present patterns continue; as volatile transient shocks will fade over time, the core rate is intended to be a reliable predictor of future headline inflation. We focus on core CPI inflation, where inflation is measured as quarter-over-quarter annualized growth rates in the underlying price index. In particular, our working measure of inflation is the average inflation rate between t and $t + 4$ quarters. Our sample spans the period from 1973:Q1 to 2019:Q1, as the [Gilchrist and Zakrajšek \(2012\)](#) credit spread is only available starting in the early 70’s.

The four top panels of [Figure B-5](#) report the estimated slope coefficients $\hat{\theta}_\tau$, $(1 - \hat{\lambda}_\tau)$, $\hat{\lambda}_\tau$ and $\hat{\gamma}_\tau$ of the quantile regression model (2).²⁰ They also visualize the partial fitted regression lines along with scatterplots of one-year-ahead average inflation against the relevant inflation determinant. In all figures we focus on three partial fitted regression lines, corresponding to the 10th, 50th and 90th quantiles. We also include the partial fitted OLS regression line, which is obtained from the commonly estimated Phillips curve. These slopes are informative about whether economic and financial conditions affect the tails of the inflation distribution differently than the median, which indicates the presence of non-linearities in inflation dynamics.²¹

The top-left panel of [Figure B-5](#) presents the quantile-specific Phillips curve coefficients associated with variations in the unemployment gap. The results are in line with the recent evidence

²⁰The quantile slopes, OLS estimates and their confidence intervals can be found in [Figure H-1](#).

²¹In [Section 4](#), we complement the information in these figures by showing the confidence bands of the estimated slopes constructed by “blocks-of-blocks” bootstrapping (see [Appendix H](#)).

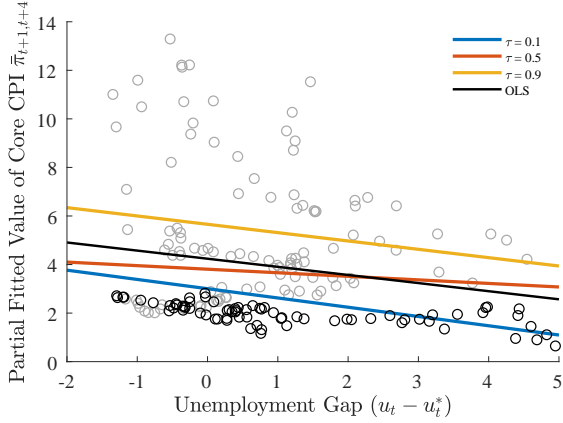
suggesting a substantial flatness in the Phillips curve, as the conditional median of inflation remains relatively muted in its response to changes in the unemployment gap. This pattern carries over to the tails, albeit to a lesser extent. Indeed, the lower tail is somewhat more responsive to the unemployment gap than the median. These results point to a mildly asymmetric response of inflation to changes in the unemployment gap.

As the top-right panel of Figure B-5 reveals, changes in relative import price inflation most strongly affect the upper tail of inflation. Increases in relative import prices tilt the inflation distribution to the upside, hence substantially increasing the odds of upside inflation risks. However, reductions in relative prices make the distribution tighter around the median, a consequence of the less significant response of the lower tail.²²


The second row of panels in Figure B-5 shows how the inflation quantiles respond to average past inflation and to inflation expectations. Here, we uncover yet another interesting asymmetry: While movements in the median and in the upper tail are mostly dominated by average past inflation, the lower tail of the distribution shows the largest response to changes in inflation expectations. That is, persistently high past inflation experiences tend to tilt the distribution to the upside, hence creating upside risks to the inflation outlook (and barely affecting the lower tail). In contrast, the modest effect of past inflation on the lower tail of the distribution implies that persistently low inflation experiences do not generate significant downside risks to the inflation outlook as the distribution does not shift to the left, but rather gets more compressed around the median. Conversely, changes in long-run inflation expectations translate one-for-one to the left tail, while the effects on the median and the upper tail are smaller. In other words, a sustained decline in longer-run inflation expectations poses serious downside inflation risks, while the effects of such a decline on upside risk are much more muted.

²²Results are similar using relative oil price inflation ($\pi_t^O - \pi_t$) (see Appendix B.3.1), share-weighted core import prices and the real exchange rate (results available upon request).


Figure B-5: Quantile Regression Slopes.




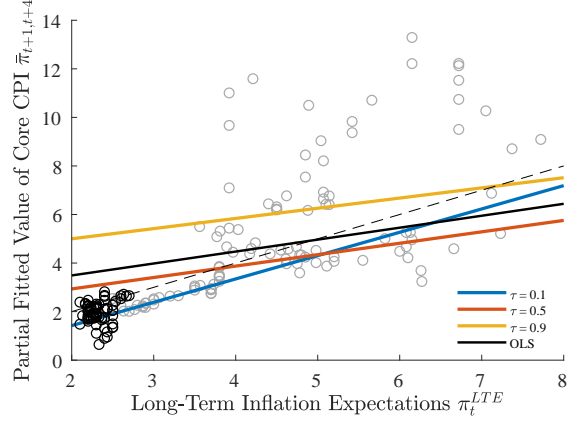
$$\hat{\theta}_\tau = \{\hat{\theta}_{0.1} = -0.38, \hat{\theta}_{0.5} = -0.15, \hat{\theta}_{0.9} = -0.34\}$$



$$\hat{\gamma}_\tau = \{\hat{\gamma}_{0.1} = 0.04, \hat{\gamma}_{0.5} = 0.04, \hat{\gamma}_{0.9} = 0.09\}$$



$$\hat{\lambda}_\tau = \{\hat{\lambda}_{0.1} = 0.96, \hat{\lambda}_{0.5} = 0.47, \hat{\lambda}_{0.9} = 0.42\}, \text{ where } \lambda_\tau \text{ is coefficient on } \pi_t^{LTE} \text{ and } (1 - \lambda_\tau) \text{ on } \pi_{t-1}^*$$



$$\hat{\delta}_\tau = \{\hat{\delta}_{0.1} = -0.19, \hat{\delta}_{0.5} = -0.02, \hat{\delta}_{0.9} = -0.19\}$$

NOTE: The figure displays the slope coefficients of the quantile regression of average four-quarter-ahead core CPI inflation defined in expression (2). The lines illustrate the slopes associated with the median (red), the 10th (blue) and the 90th (yellow) inflation quantile. The black lines are the OLS estimates. Circles indicate scatterplots of average future inflation against a given inflation determinant. Grey circles indicate scatterplots of average future inflation against a given financial variable prior to 1999:Q4 whereas black circles indicate the scatterplot for the period starting in 2000:Q1.

For each economic factor, we highlight its relationship with the inflation outlook during the most recent period with the black cloud of points which focuses on observations from the year 2000 onwards. As we will show in Section 2, the roles of the unemployment gap and of relative prices in accounting for variations in average future inflation are considerably dampened. At the same time, we find that the ability of inflation inertia to move the inflation distribution is dramatically reduced, bestowing its predominant role to long-run inflation expectations.

The lowest panel of Figure B-5 shows the effects of changes in credit spreads on the inflation quantiles.²³ Overall, the negative sign suggests that high credit spreads (i.e., tight financial conditions) generate downside inflation risks. Interestingly, credit spreads affect both tails of the inflation distribution. However, as the figure shows, there is substantial subsample instability governing the link between the tails and variations in credit spreads. The sub-period 1973-1999 is characterized by relatively small variations in credit spreads in a period of high and volatile inflation induced in part by systematic increases in energy prices. This is captured by the light-grey cloud of points. From 2000 onward, low variability of inflation around 2 percent has been a notable aspect of the stability of the macroeconomic landscape that has coexisted with substantial variation in credit spreads, a phenomenon amplified by the global financial crisis. These combinations correspond to the black cloud of points. As we show in Section 2, this more recent period helps in correctly identifying the relationship between the tails of the distribution and the credit spread, which is confounded by the time aggregation. We find that in the post-2000 period, most of the reduction in inflation following high credit spreads is concentrated in the lower tail of the distribution, while the effects on the upper tail are poorly estimated (i.e., the point estimates are associated with high levels of uncertainty). Thus, the results point to a close relationship between a tightening of financial conditions and risks of “low inflation”, while periods of “frothiness” and bully financial markets have little effects on the upper tail of the inflation distribution; instead, they make the distribution of inflation more concentrated around the median.

Finally, we performed an ANOVA test on the equality of the slope coefficients across quantiles.

²³In Appendix B.3.2 we show that these results are robust to the choice of other financial variables.

The test rejected the null hypothesis of equality between the slopes on the 10th and 50th quantile at 5% significance level for all variables except for relative import prices, thus highlighting the importance of both labor market, inflation and financial determinants in shaping downside risk and left-skewness of the inflation distribution. On the other hand, relative prices is the only variable whose slope for the 90th and 50th quantile are statistically different, confirming its importance for upside risks to and right-skewness in the inflation outlook.

B.1 The Role of Financial Conditions

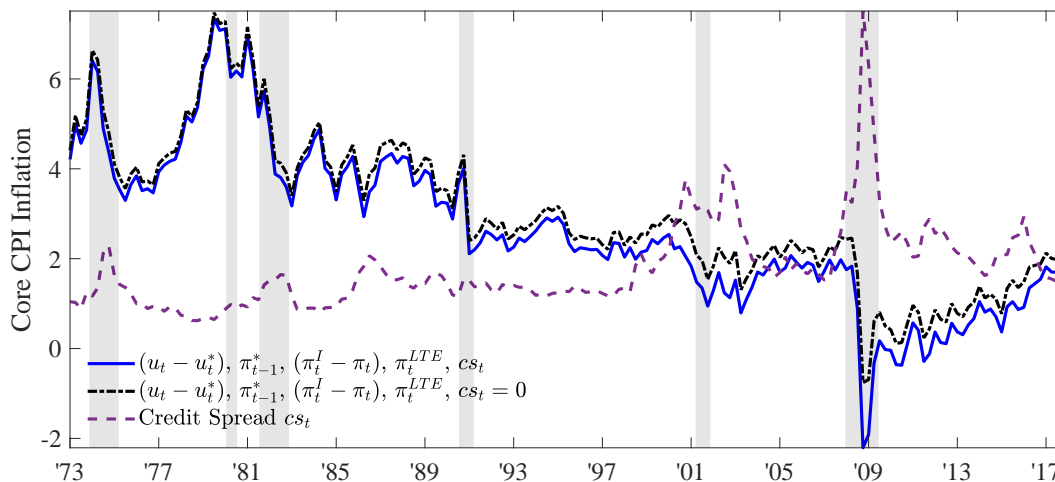
We now illustrate the influence of credit spreads on downside risks to inflation and their variations over time (later on, we will focus on the last subsample starting in 2000 by comparing the United States with the experience in the euro area). To do so we construct the 10th quantile of inflation arising from the quantile model in its baseline version and in a version in which ignores the role of financial variables.

Figure B-6 displays the evolution over time of the 10th inflation quantile in the baseline model – which includes the effects of credit spreads (solid blue line) – and the 10th quantile constructed by shutting down the effects of this financial variable (black dash-dotted line). The graph also includes the time series of the credit spread (purple dashed line). It is evident that the quantile model in which the role of financial variables is disregarded can be a misleading measure of downside inflation risk if there are significant changes in credit spreads. As credit spreads have been growing over time, so does this model's miss. Indeed, earlier in the sample the 10th quantile is barely affected by credit conditions, while starting in the early 2000s – once the model accounts for more pronounced variations in credit spreads – headwinds coming from financial conditions substantially increase the odds of low inflation.

During the 1990s, there is a progressive reduction in the lower tail of the distribution that remained fairly insensitive to financial developments. Starting in the 2000s, the 10th quantile showed a remarkable resistance to go well below 2 percent. This phenomenon ended at the onset of the global financial crisis and the subsequent zero lower bound episode. The lower tail of the distribu-

tion was such that downside inflation risks materialized, with non-zero deflation probabilities. The aftermath of the global financial crisis shows that the lower tail of the distribution exhibits substantial persistence. That is, the tightening in credit conditions tilted the distribution to the downside for a prolonged period. The reduction in downside risks was enabled by improvements in the labor market and sustained by inflation expectations.

Figure B-6: Time Evolution of 10th Inflation Quantile Across Models.



NOTE: The figure displays the time evolution of the 10th inflation quantile estimated from the quantile regressions model (2), in its baseline version (blue solid) and in its version where the effect of credit spreads is set to zero (black dash-dotted). The credit spread (purple dashed) is also reported. Shaded bars indicate NBER-dated recessions.

Calibration of Predictive Density We now formally assess how financial variables influence the accuracy with which the quantile model characterizes the actual distribution of average future inflation. In particular, we test for correct calibration of the conditional predictive distributions implied by the baseline model in one case and by the model which does not condition on financial variables in the other. To do so we use the test of Rossi and Sekhposyan (2019), which evaluates the absolute predictive ability of a model at its estimated parameter values and, thus, in finite samples.²⁴ In this sense, both the parametric model and the estimation technique employed are being evaluated.

To run the test, we first define the probability integral transform (PIT), i.e., the conditional

²⁴See Rossi (2014) for an excellent summary of density forecast evaluations.

quantile z_t that corresponds to the realized observation $\bar{\pi}_{t,t+4}^*$:

$$z_t \equiv F_{\bar{\pi}_{t+1,t+4}^*}^{-1}(\tau|x_t) = Prob(\bar{\pi}_{t+1,t+4} < \bar{\pi}_{t+1,t+4}^* | x_t), \quad (\text{B-1})$$

where $F_{\bar{\pi}_{t+1,t+4}^*}^{-1}(\tau|x_t)$ refers to the inverse of the conditional CDF or, equivalently, to the conditional quantile function evaluated at the realized value $\bar{\pi}_{t,t+4}^*$. In a perfectly calibrated model, the predictive density should feature a CDF which is uniform, i.e., equal to the 45° line. This property implies that the probability that the realized value is above or below the predicted value is the same (on average, across time) irrespectively of whether high or low realizations of the predicted variable are considered. Following this logic, if the empirical CDF of the PITs lies outside of the 5% critical values, then the [Rossi and Sekhposyan \(2019\)](#) rejects the null hypothesis of correct calibration.

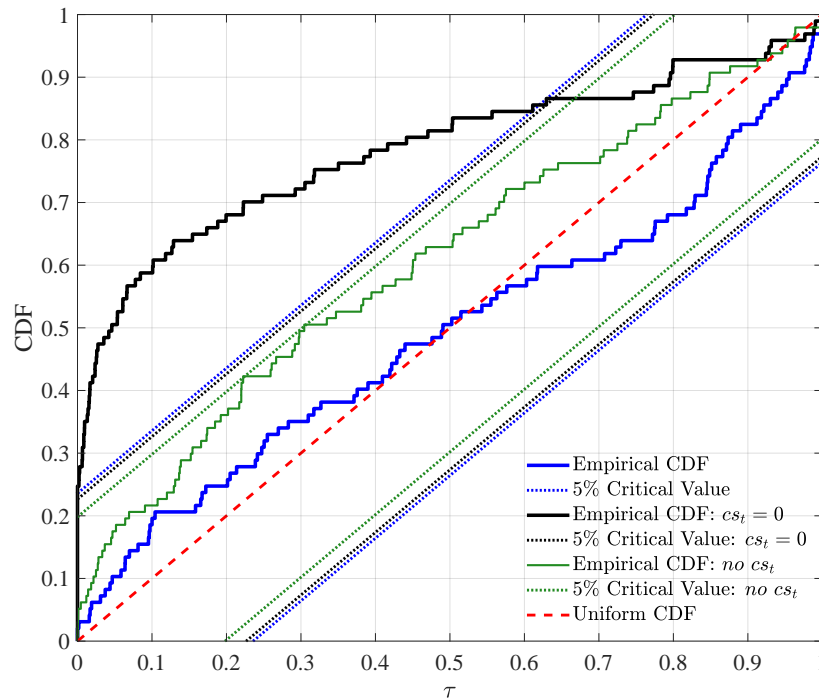
In [Figure B-7](#) we plot the CDF of a uniform distribution (red, dashed) as well as the empirical CDFs of the PITs obtained from the baseline model (blue line) and from two versions which either do not condition on the credit spread in the estimation (green line) or in the construction of the inflation quantiles (black line), along with their 5% critical values (dash-dotted lines). The model which nets out credit spreads in the construction of the quantiles (black line) can be interpreted as quantifying the partial effect of credit conditions in the augmented Phillips curve model.

The critical values are bootstrapped following the [Rossi and Sekhposyan \(2019\)](#) procedure for multi-step-ahead forecasts. As in [Adrian, Boyarchenko, and Giannone \(2019\)](#), the PITs are constructed via an expanding rolling windows estimation initially using 20 years of data. Confidence bands should thus be taken as general guidance since [Rossi and Sekhposyan \(2019\)](#) derive them for PITs computed using a fixed rolling window scheme.²⁵

Unlike the baseline model, the model that disregards the role of financial variables in the construction of the quantiles (black test) does not pass the test for correct calibration – as it poorly specifies the predictive inflation distribution by placing too much mass on its lower tail. The CDF

²⁵Using a rolling window scheme with 20 years of data we confirm that the model which does not use credit spreads in the construction of the quantiles fails to pass the test because of poor calibration of the left tail. Also, we still can't reject correct specification of the predictive density of our baseline model (results are available upon request).

Figure B-7: Rossi and Sekhposyan (2019) Test for Correct Calibration of Predictive Density.



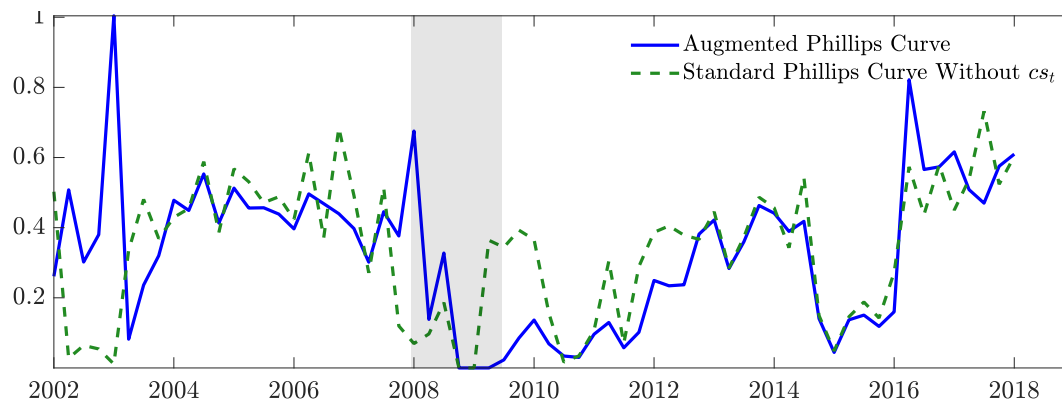
NOTE: The figure illustrates the CDF of a uniform distribution along with the empirical CDFs of out-of-sample PITs obtained from the quantile regressions model (2), in its baseline version (blue) and in two versions which either do not condition on financial variables in estimation (green) or in the construction of the inflation quantiles (black). The 5% critical values for each model (dashed-dotted), are bootstrapped following the Rossi and Sekhposyan (2019) procedure for multi-step-ahead forecasts. As in Adrian, Boyarchenko, and Giannone (2019), the PITs are constructed via an expanding rolling windows estimation initially using 20 years of data. Confidence bands should thus be taken as general guidance since Rossi and Sekhposyan (2019) derive them for PITs computed using a fixed rolling window scheme.

of the PITs from the model neglecting financial variables in estimation (green line) is within the critical values but performs worse than the baseline along the entire inflation distribution except on the upper tail. We take this as evidence that our model calibration is at least as good, if not better, than the model which does not condition on credit spreads.

Predictive Scores We further evaluate the reliability of the predictive distribution by measuring the accuracy of the model’s density forecasts through its predictive scores. These are computed by evaluating the model’s predictive distribution at the realized value of the time series. A higher predictive score indicates more accurate predictions, as the model assigns a higher probability to outcomes that are closer to the realized value. We compute the predictive scores in an out-of-

sample exercise where the predictive distributions are calculated using an expanding window.

Figure B-8: Out-of-Sample Predictive Scores. Model Estimated from 1973:Q1.



NOTE: The figure illustrates the out-of-sample predictive scores obtained from the density constructed by fitting a skewed- t distribution on the conditional quantiles from the quantile regression model (2) that starts in 1973:Q1. The quantiles are constructed via an expanding rolling windows estimation starting in 2001:Q1. In particular, for the first out-of-sample evaluation, we use data for average core CPI inflation over the next year until 2000:Q1 to estimate the model. Then, we use data for the explanatory variables in 2001:Q1 to construct the out-of-sample predictive scores.

Figure B-8 plots the scores of the predictive distribution from the baseline model starting in 1973:Q1 of a Phillips curve augmented with credit spreads, together with the scores of the predictive distribution from the standard Phillips curve model that does not condition on credit spreads. The model with and without credit scores performs about equally well, with the augmented model and the standard model having higher predictive ability respectively at the onset and toward the end of the global financial crisis.

B.2 The Predictive Distribution of Inflation

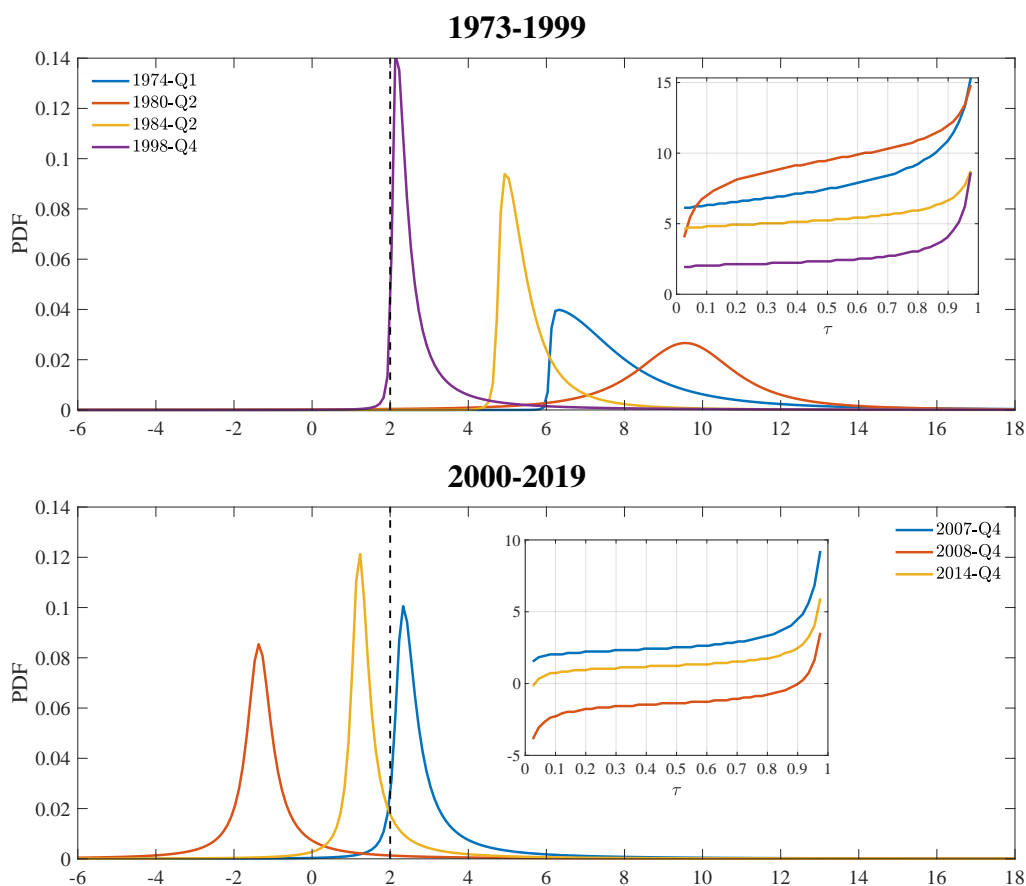
Figure B-9 displays, for selected dates, the estimated conditional predictive densities of average one-year-ahead inflation and their associated fitted inverse cumulative distribution functions – shown in the inset boxes.²⁶ The top and the bottom panels illustrate the contrast between the odds of high inflation, which characterized the inflation distribution during the first part of the sample, and the progressive switch toward downside risks to the inflation outlook which built up

²⁶As formally described in Section 1, we construct the skewed t -Student probability density function of inflation using the quantiles estimated using the regression model (2).

at the onset of the global financial crisis.

As shown in the top panel, during the first subsample we select four dates. We start our time travel in early 1974:Q1, around the recession triggered by the first wave of oil shocks and the easing cycle of the Federal Reserve. The second quarter of 1980 is chosen to capture the effects of the second OPEC shock. We pick these two dates as representative of the Great Inflation period. Then, we look at the distribution in the mid-eighties, more precisely in 1984:Q2, to capture the effects of the Volcker disinflation – a disinflationary transition period that led the U.S. economy into the so-called Great Moderation. This last period is represented by showing the estimated conditional inflation distribution in 1998:Q4.

Figure B-9: Conditional Predictive Inflation Densities at Selected Time Episodes.



NOTE: The figures shows for selected time episodes the estimated skewed- t conditional densities of average four-quarter-ahead core CPI inflation associated with the quantile regressions model (2). The inset box reports the values of average future inflation across quantiles at the same selected periods. More formally, it depicts the estimated skewed- t inverse CDF associated with the conditional densities in the main panel.

Overall, the estimated quantile models capture how the inflation distribution moved from the right - with significant upside inflation risks associated with the persistent effects of the oil shocks in the mid-70s and early 80s – to the left, with almost negligible upside risks of inflation falling above 4 percent at the eve of the 2000s. Beyond these general changes in the distribution, it is worth noting how the shape of the distribution substantially changed over time. The first OPEC shock led to an asymmetric inflation distribution, with almost negligible odds of inflation falling below 6 percent but a long right tail creating very large upside risks to the inflation outlook. These risks materialized after the second OPEC shock. The distribution shifted further to the right, with upside risks becoming more balanced around a much higher average future inflation rate. Chairman Volcker’s reaction to the great concern about the rise in long-run inflation expectations led to the aggressive monetary policy reaction designed to curb inflationary pressures and progressively hamper inflation expectations. The effects of this policy are reflected in the noticeable shift-to-the-left in the estimated inflation distribution, with upside risks substantially reduced by the mid-80s. During those years, the inflation distribution became more symmetric and substantially more concentrated around the median. This disinflationary process continued during the 90s, and by the end of the millennium the distribution concentrated around 2.5 percent, with the lower tail remaining quite insensitive to economic or financial developments and showing a remarkable resistance to go below 2 percent, a feature we analyze in Section 2.

The bottom panel of Figure B-9 selects a few dates in the evolution of the inflation distribution during the last 20 years, and it depicts a completely different story from the first part of our sample – although there is remarkable similarity between the inflation distribution at the eve of the Great Recession, the blue line in the bottom panel that corresponds to 2007:Q4, with the one shown for the last quarter of 1998 in the top panel. During this more recent period the reasons for concern move from upside inflation risks to low-inflation or even deflation risks – with the ghost of the Great Depression frightening central banks during the aftermath of the global financial crisis. Although we devote Section 4 to develop this issue in depth, the three dates chosen in the bottom panel of Figure B-9 serve as a useful preamble to that discussion.

The global financial crisis and the dramatic increase in credit spreads translated into a right-skewed (i.e., fatter left-tailed) inflation distribution, with the median moving progressively closer to the lower tail. This phenomenon was exacerbated during the subsequent zero lower bound episode. The lower tail of the distribution was such that downside inflation risks materialized, with non-zero probabilities of deflation (see the red line that displays the distribution in 2008:Q4).

This emergence of substantial downside risks to inflation has been the main source of increasing concern among researchers and policymakers. Monetary policy provided accommodation to support a strong job market, to abate the lingering headwinds from the financial crisis, and to keep inflation expectations well-anchored. These effects translated into a substantial shift to the right in the inflation distribution, curtailing the odds of deflation by the end of 2014 (yellow distribution shown in the bottom panel of Figure B-9).

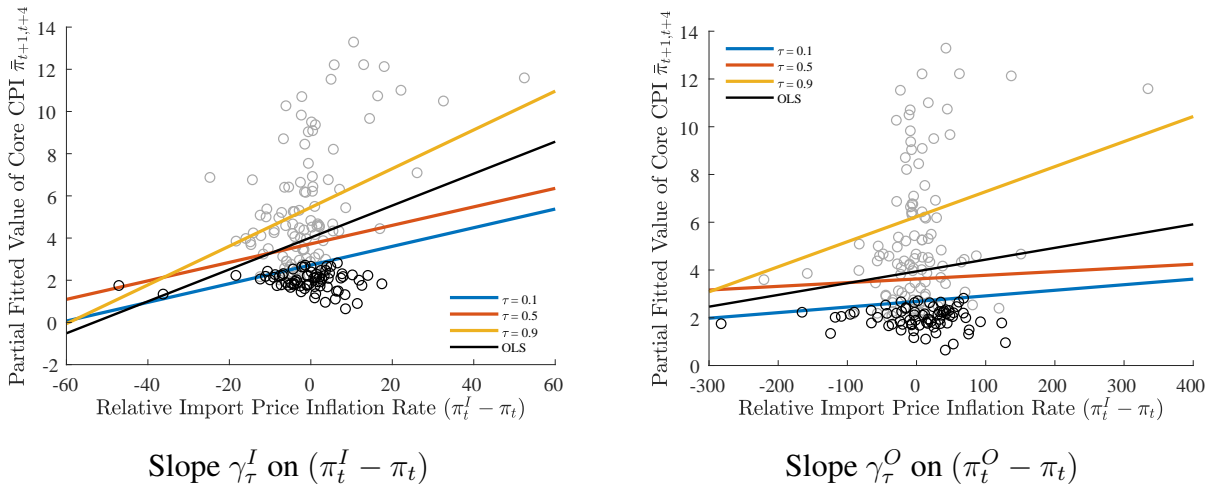
Finally, it is interesting to notice how the inflation distributions implied by the model estimated over the full sample suggested bigger risks to the inflation outlook than those that actually materialized. As we will show in the next section, this is a result of the fact that the Phillips-curve relationships evolved over time and are substantially different in the two subsamples just considered. Indeed, the inflation risks during the Great Recession which are informed by full-sample estimates mostly reflect conventional Phillips-curve relationships—which predominate the first subsample and which imply a higher sensitivity of inflation to macroeconomic conditions—rather than the modern Phillips-curve relationships—which are informed by the last subsample where stable inflation expectations take the center stage and, thus, inflation is less sensitive to developments in the economy. In Section 2, we investigate the relationships in these two subsamples separately so as to uncover a more accurate picture of the relative importance of risk factors over time.

B.3 Robustness

We now present additional robustness results for the full sample estimation. The data are described in Appendix A. First, we show that conditioning on energy prices instead of imported goods yields very similar results (see Figure B-10). Figure B-11 displays how changes in the financial variable affect the estimated slopes in the baseline quantile regression model – for core CPI inflation – over the full sample period. We consider three alternative financial variables: corporate bond spreads (top-left panel), excess bond premium as constructed by [Gilchrist and Zakrajšek \(2012\)](#) (top-right panel), and the national financial conditions index (bottom-center panel). The results are striking. The lower tail of the distribution of inflation is highly negatively responsive to changes in financial conditions, but the upper tail of distribution is not.

B.3.1 Oil vs. Import Relative Price Inflation

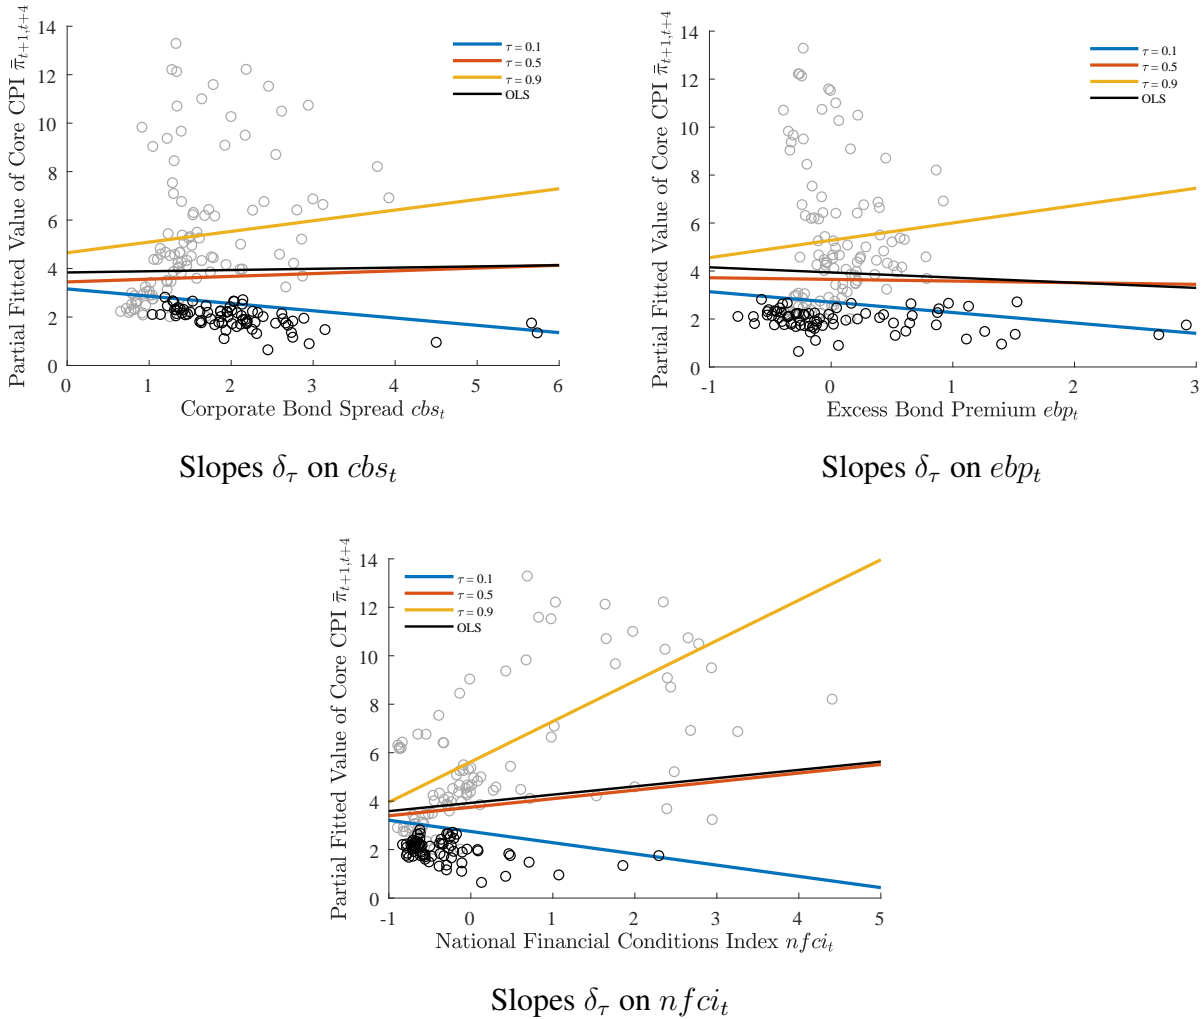
Figure B-10: Quantile Regressions Slopes Across Relative Price Measures.



NOTE: The figure displays the slope coefficients on relative prices of the quantile regression of average four-quarter-ahead core CPI inflation defined in expression(2). The lines illustrate the slopes associated with the median (red), the 10th (blue) and the 90th (yellow) inflation quantile. The black lines are the OLS estimates. Circles indicate scatterplots of average future inflation against a given inflation determinant. The left panel corresponds to the model where the relative price measure is relative import price inflation, whereas the right panel considers the model where the relative price measure is relative oil price inflation. Grey circles indicate scatterplots of average future inflation against a given financial variable prior to 1999:Q4 whereas black circles indicate the scatterplot for the period starting in 2000:Q1.

B.3.2 Different Financial Variables

Figure B-11: Quantile Regressions Slopes Across Financial Variables.



NOTE: The figure displays the estimated coefficients of the quantile regression of average four-quarter-ahead core CPI inflation defined in (2), using corporate bond spreads (top, left), the [Gilchrist and Zakrajšek \(2012\)](#) excess bond premium (top, right) and the National Financial Conditions Index (center, bottom) and the same sample period as the baseline. The lines illustrate the slopes associated with the median (red), the 10th (blue) and the 90th (yellow) inflation quantile. The black lines are the OLS estimates. Grey circles indicate scatterplots of average future inflation against a given financial variable prior to 1999:Q4 whereas black circles indicate the scatterplot for the period starting in 2000:Q1.

C The Relationship between Credit Spreads and Inflation Quantiles in a Nonlinear DGSE Model with Financial Panics

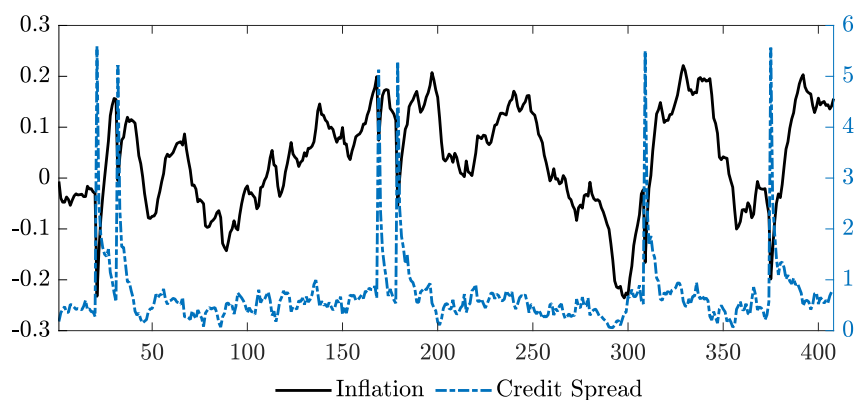
The Model We consider the fully micro-founded nonlinear DSGE model of [Gertler, Kiyotaki, and Prestipino \(2019\)](#), which features the possibility of a severe financial crisis through a bank run. There are two equilibria in their model, one with and one without a financial panic. When shocks are small, the economy fluctuates around a standard equilibrium.²⁷ In contrast, a big negative shock pushes the economy into a bank run equilibrium. Combined with a sunspot shock, it triggers a financial panic and bank net worth collapses. Banks are forced to sell assets, which ultimately disrupts firms' borrowing. Consequently, economic activity drops substantially more than in the equilibrium without a bank run.

Parameterization As in [Loria, Matthes, and Zhang \(2019\)](#), we simulate the model using the original calibration of the deep parameters and of the capital quality shock process (the only fundamental shock in the model). In order to generate rare financial crisis, we calibrate the process for the sunspot shock such that a bank run equilibrium arises after a big negative shock (above two standard deviations).²⁸

Simulation We simulate this model 1000 times for 413 periods and store the inflation rate, the credit spread, and the capital quality shock. Results are robust to the use of alternative measures of financial conditions in the model. The simulated data shown in the top panel of [Figure C-1](#) indicate that also in this model there is a non-linear relationship between inflation and financial conditions. Indeed, large credit spreads are associated with extremely negative inflation realizations.

²⁷We follow the original paper in focusing on a capital quality shock as a representative structural shock. However, other shocks would give rise to qualitatively very similar results because the mechanism creating asymmetry is not specifically tied to one structural shock.

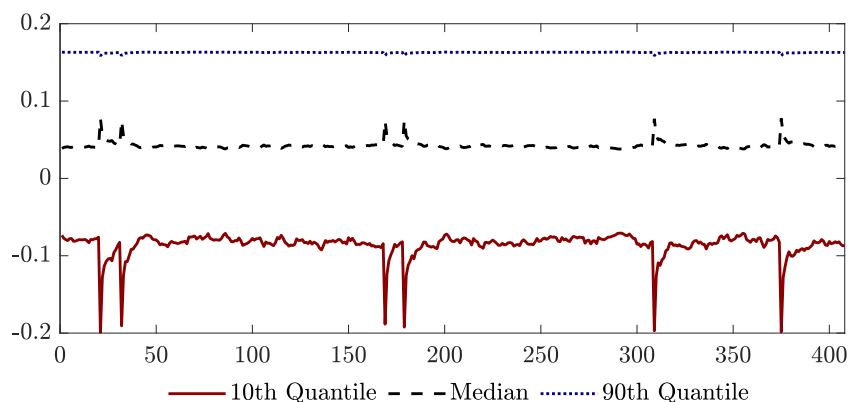
²⁸The calibration intends to make the bank run event (see [Figure 2](#) of the original paper) re-occurring in the simulated sample. In their event study, the authors feed in two consecutive negative capital quality shocks of roughly one standard deviation, together with a sun-spot shock, to generate a bank-run. We deviate slightly from their event study by having sun-spot shocks occurring concurrently with a negative two standard deviations capital quality shocks to ensure the probability of bank run of 2.5% across simulated samples.



NOTE: Example from one simulation. Inflation(left axis) and Credit Spread (right axis).

Figure C-1: Simulated Data from Gertler, Kiyotaki, and Prestipino (2019) Model.


Quantile Regression We focus on current inflation as the shocks are transitory and only create a sharp but short-lived recession. Further, we only consider the credit spread as conditioning variable since the nonlinearity in the model is coming from financial conditions. The quantile regression picks up the nonlinear relationship between financial conditions and inflation, suggesting that financial conditions have a more negative effect on the left tail than on the median. Thus, also in this example, the left tail of the distribution falls substantially during times of extreme financial distress, characterizing a vulnerable inflation outlook. This is further illustrated in the bottom panel of Figure C-2 which plots the quantiles of inflation associated with the simulated data.



NOTE: Example from one simulation.

Figure C-2: Quantiles from Gertler, Kiyotaki, and Prestipino (2019) Model.

The Role of Credit Spreads As shown in Figure C-3, the quantile regression slopes echo our empirical findings that higher credit spreads are associated with an increase in downside risks to inflation, as signaled by the negative slope coefficient for the 10th quantile. In terms of identification, what delivers this result is that during bad times, those featuring a bank run, the conventional channel whereby lower demand results in subdued price pressures is strongest. The reason why the 90th quantile indicates a positive relationship is that in good times, without a bank run, a capital quality shock reduces capital, and thus results in an increase in the rental rate of capital and in marginal costs. The contrasting effects of the demand and cost of capital channels become evident in the bottom-right panel of Figure 2 in [Gertler et al. \(2019\)](#), where in response to a capital quality shock inflation increases in a no bank run equilibrium and drops in a bank run equilibrium. The median quantile captures the tension between these two effects. Indeed, it is around zero as in normal times these two effects almost offset each other.




NOTE: Black squares are medians across simulations. Shaded areas are 68% confidence bands.

Figure C-3: Quantile Slopes from [Gertler, Kiyotaki, and Prestipino \(2019\)](#) Model.

D Robustness of Subsample Results

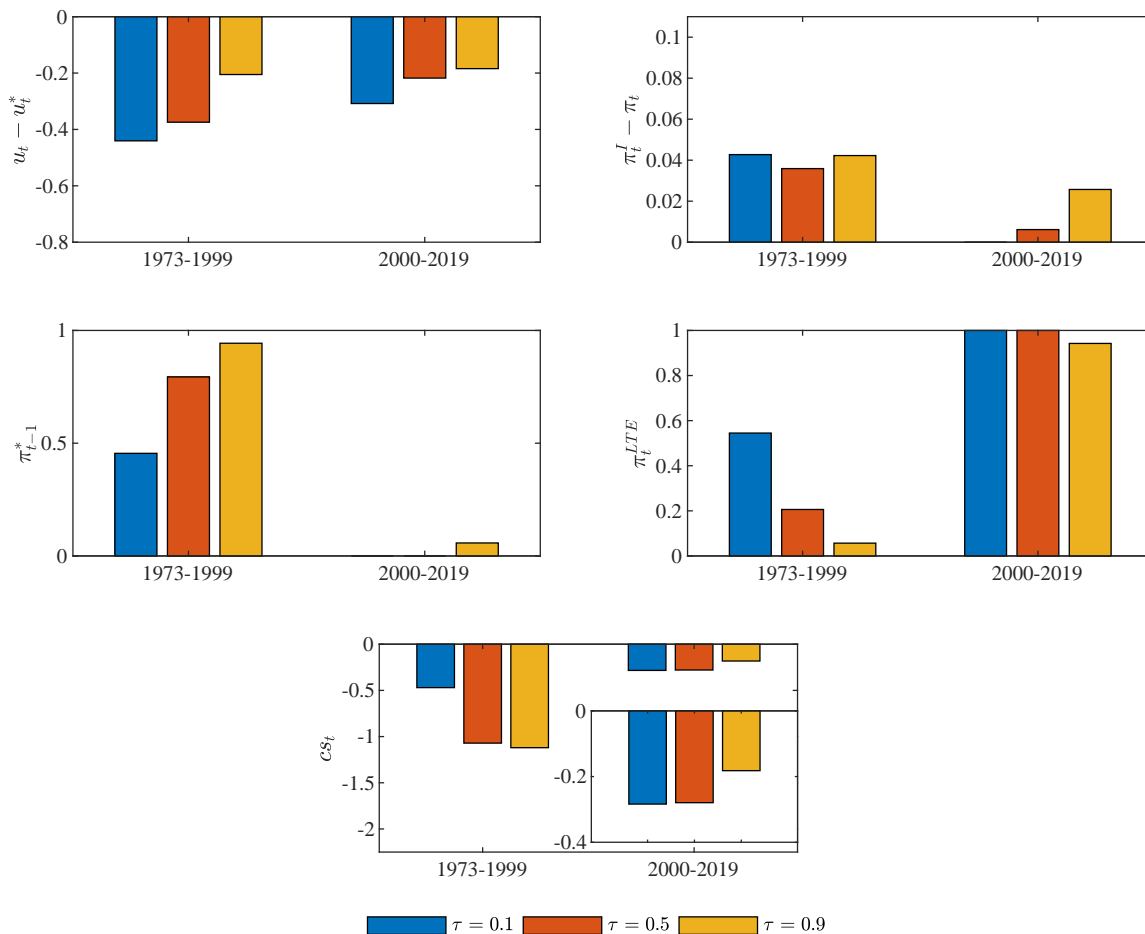
As noted in the main text, there is substantial subsample instability in the relationship between credit spreads and inflation. This is confirmed by the subsample results shown in Figure D-1 and Figure D-2 which thus reiterate how the importance of risk factors changed across the two subsamples and mimic the one in the main text in which we report the estimated quantile-specific slopes for core CPI (i.e., Figure 1). Importantly, once we control for subsample stability, the results are extremely similar across these different inflation measures. We further confirm our findings by substituting the credit spread as a measure of credit conditions with corporate bond spreads in Figure D-3 and with the excess bond premium in Figure D-4.

Figure D-1: Quantile Regression Slopes Across Subsamples, Core PCE.



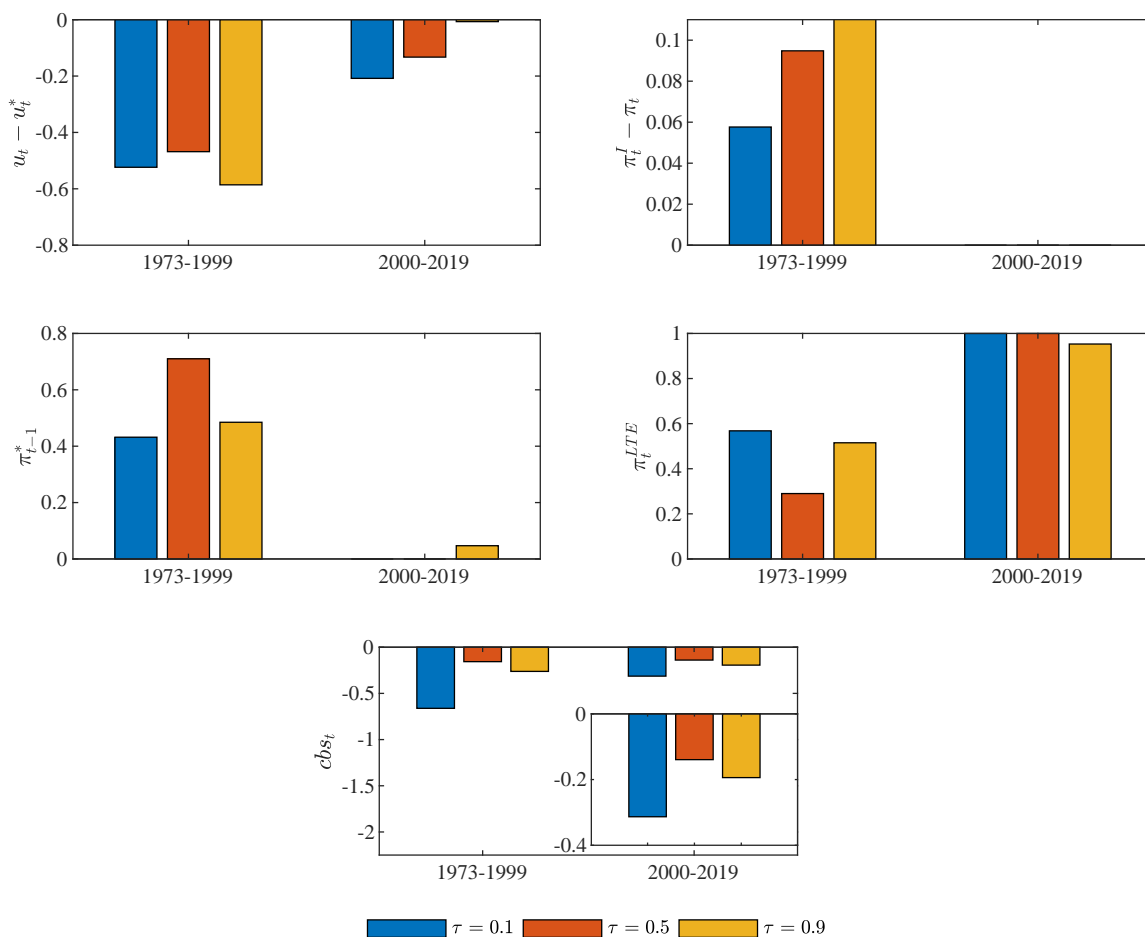
NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead core PCE inflation defined in expression (2). Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate the coefficients associated with the 10th quantile (blue), median (red) and 90th quantile (yellow).

Figure D-2: Quantile Regression Slopes Across Subsamples, [Stock and Watson \(2019\)](#) CSI.



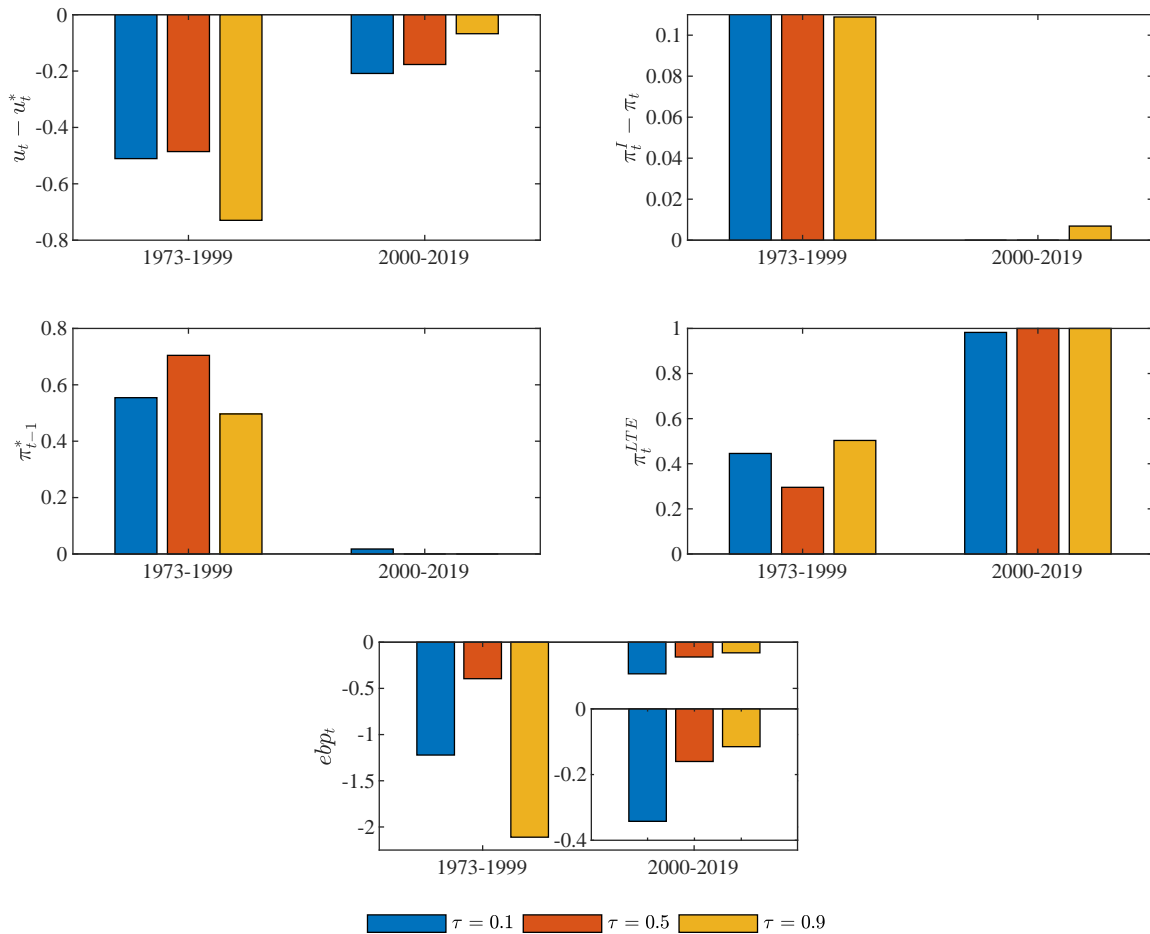
NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead [Stock and Watson \(2019\)](#) Cyclically Sensitive Inflation defined in expression (2). Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate the coefficients associated with the 10th quantile (blue), median (red) and 90th quantile (yellow).

Figure D-3: Quantile Regression Slopes Across Subsamples, Corporate Bond Spreads



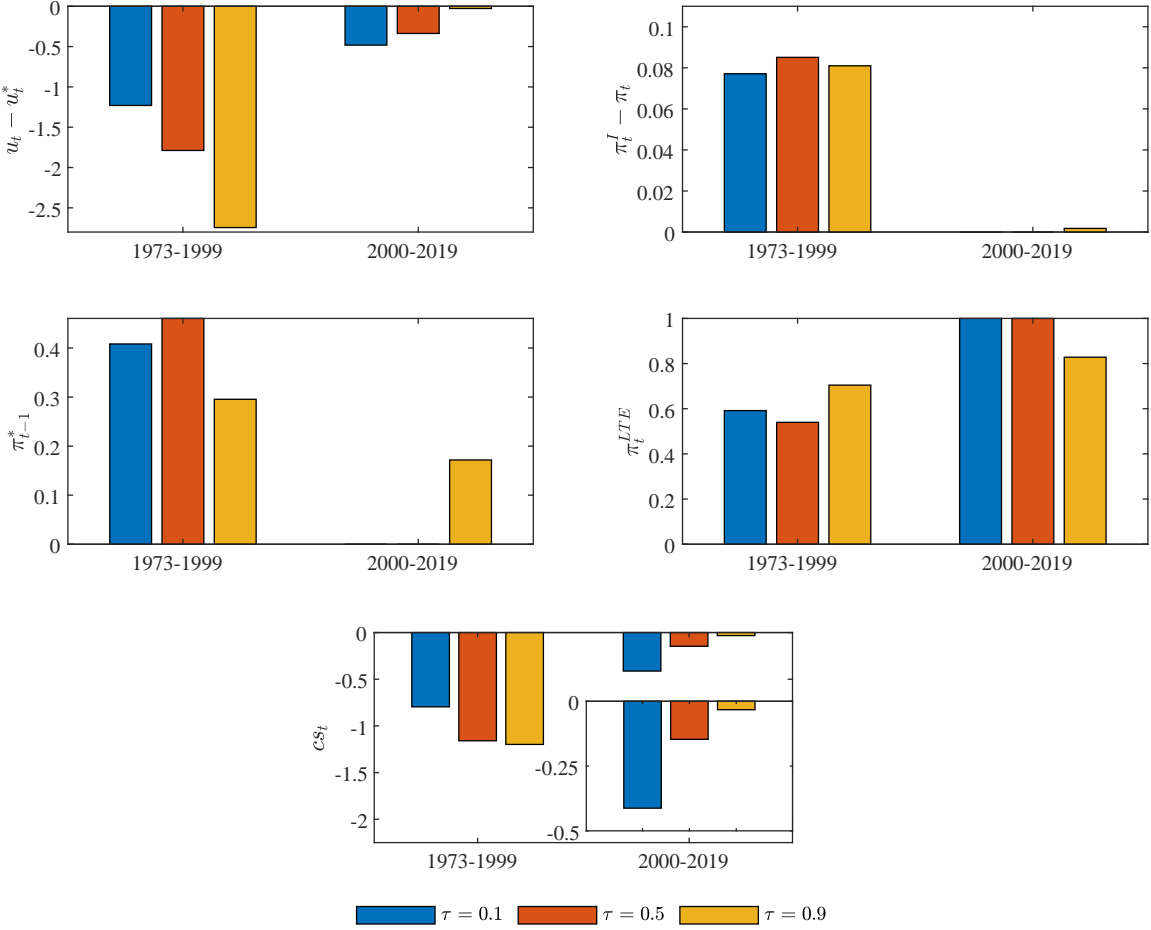
NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead core CPI inflation defined in expression (2), using corporate bond spreads and the same sample period as the baseline. Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate the coefficients associated with the 10th quantile (blue), median (red) and 90th quantile (yellow).

Figure D-4: Quantile Regression Slopes Across Subsamples, Excess Bond Premium.



NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead core CPI inflation defined in expression (2), using the [Gilchrist and Zakrajšek \(2012\)](#) excess bond premium and the same sample period as the baseline. Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate the coefficients associated with the 10th quantile (blue), median (red) and 90th quantile (yellow).

Figure D-5: Quantile Regression Slopes Across Subsamples, Controlling for SPF Forecasts.



NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead core CPI inflation defined in expression (2), controlling for SPF forecasts of output growth and unemployment rate (averaging across horizons) and using the same sample period as the baseline. Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate the coefficients associated with the 10th quantile (blue), median (red) and 90th quantile (yellow).

E The Missing Deflation and Inflation Debate

In this section we ask how the inflation distribution and the inflation-at-risk probabilities would have looked like from the perspective of these two different subsamples. In particular, we perform the following experiment. We first consider the subsample running from 1973 to 1999 and use the estimated relationship between inflation and its economic and financial determinants in that period to compute the “counterfactual” inflation quantiles in the post-2000 era. These quantiles are then used to construct the inflation distribution and the associated inflation-at-risk probabilities – as if the conditions characterizing inflation in the first subsample had prevailed in the last part of the sample. We then run the opposite experiment. That is, we use the quantile regression model estimated over the sample ranging from 2000 to 2019 to construct the inflation distribution and inflation-at-risk probabilities prevailing over that same sample. The assumption of stability in the underlying relationships characterizing the inflation distribution in the first sample leads to the appearance of “missing” inflationary and deflationary episodes.

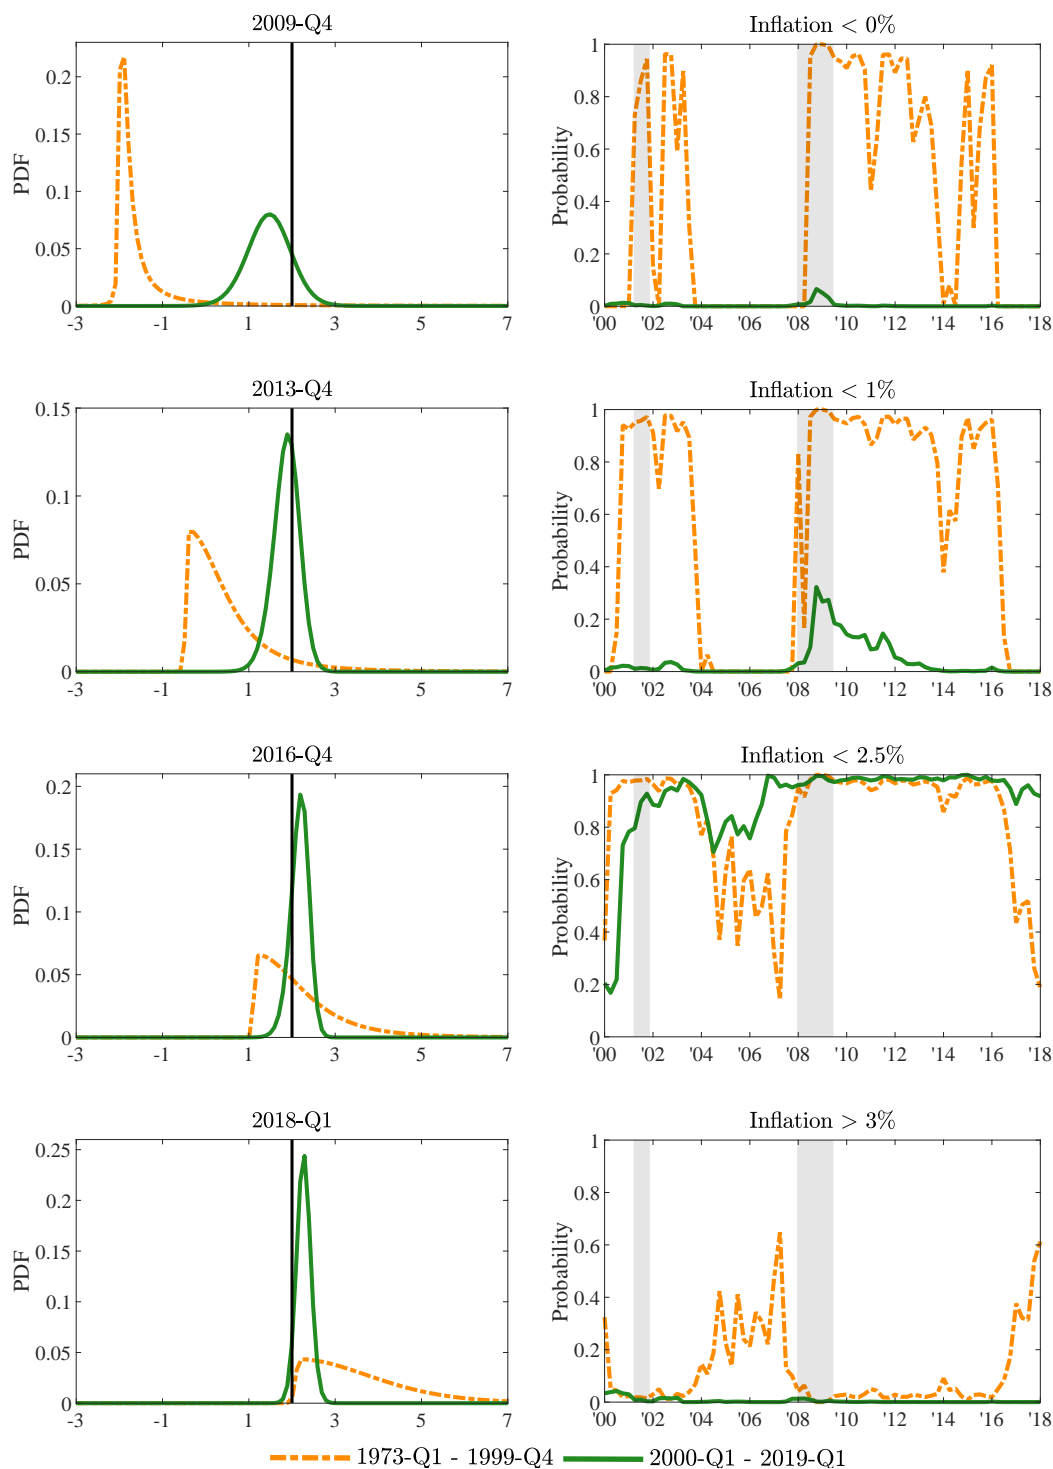
In Figure E-1 we present the conditional predictive densities and inflation probabilities associated with these two “counterfactual” economies. Two striking results stand out. First, from the point of view of the pre-2000 economy (yellow dash-dotted line), during the zero lower bound episode we should have observed disinflation, and even deflation, with probability one as well as right-skewed distributions. In contrast, the post-2000 economy suggests only tiny deflation and disinflation probabilities and more concentrated inflation distributions. This result directly speaks to the debate on the “missing deflation”, i.e., the observed discrepancy between the deflation/disinflation predicted by conventional economic wisdom given the weak inflation fundamentals and the observed resistance of actual inflation falling to negative territory during the zero lower bound period. Second, the pre-2000 economy (green solid line) supports a rise in inflation with the recovery and subsequent expansion of the U.S. economy, as reflected by the increase in the probability of inflation being above 3 percent for the most recent years and by the right-skewed distributions. Conversely, the post-2000 economy wouldn’t have suggested any change in the inflation odds. This tension is related to debate on the “missing inflation”, the mirror image of the

“missing deflation” conundrum.

These results can be easily rationalized by recalling the different role of inflation expectations in the two “counterfactual” economies. While in the pre-2000 economy, the inflation distribution is equally responsive to inflation inertia and inflation expectations as well as very sensitive to changes in the unemployment gap and in relative import prices, in the post-2000 economy inflation dynamics are mainly driven by inflation expectations. As the latter have been extremely well-anchored around 2 percent since the early 2000s (as opposed to the period prior to that, see Appendix A), the post-2000 economy would have predicted average future inflation and its tails to stay in check. Notice that the time-varying sensitivity of inflation to inflation expectations (and inflation inertia), as well as its ability to explain the “missing deflation/inflation” puzzle, had already been explored and established by [Blanchard, Cerutti, and Summers \(2015\)](#). Our analysis thus extends their findings to the entire distribution of the inflation outlook.

While well-anchored inflation expectations go a long-way in accounting for the stability of the inflation distribution throughout the Great Recession and the subsequent recovery period, it would be misleading to think that it was its sole driver. In fact, as we pointed out before, downside risks were also sensitive to changes in the labor market and financial conditions. Thus, monetary policy not only ensured price stability on average by keeping inflation expectations in check but also avoided inflation risks by supporting the job market and easing credit conditions. We discuss the relative role of these risk factors for the inflation outlook next.

Figure E-1: “Counterfactual” Predictive Densities (left) and Inflation Probabilities (right).



NOTE: The left panels show “counterfactual” skewed t -Student conditional densities of average four-quarter-ahead core CPI inflation computed using “counterfactual” conditional quantiles over 2000-2019 which were obtained using different subsample estimates of the quantile regressions model (2). The right panels show “counterfactual” inflation probabilities for different cutoffs. These probabilities are computed from the “counterfactual” skewed t -Student conditional densities shown in the left panels. Shaded bars indicate NBER-dated recessions.

F Financial Market Probabilities

This section is devoted to find an external validation of the previous results for the United States. We frame the analysis on whether the distribution of inflation embodied in financial options supports some of our conclusions about inflation risks derived from our quantile regression analysis. Specifically, to put the emphasis on the lower tail of the inflation distribution, we first compare the inflation probabilities implied by the “quantile Phillips curve model” with the one-year-ahead CPI inflation probabilities derived from inflation caps and floors contracts (as described by [Kitsul and Wright, 2013](#)). We focus on the probability of future inflation being below 1 percent.²⁹

The black line shown in the top panel of Figure [F-1](#) displays the options-implied probability of *headline* CPI inflation next year being below 1% – from mid-2009 until the end of 2016. A solid reading from this market portrays a quite striking picture. Market participants have systematically been pricing a probability of inflation below 1% of around 40 percent up until 2016, and only after that date this probability has been moving down to levels slightly below 20 percent – ten years after the global financial crisis. However, many analyses have attributed the level and evolution of this probability to the high correlation between market-based measures of inflation expectations and oil-price related shocks – especially since the onset of the global financial crisis.

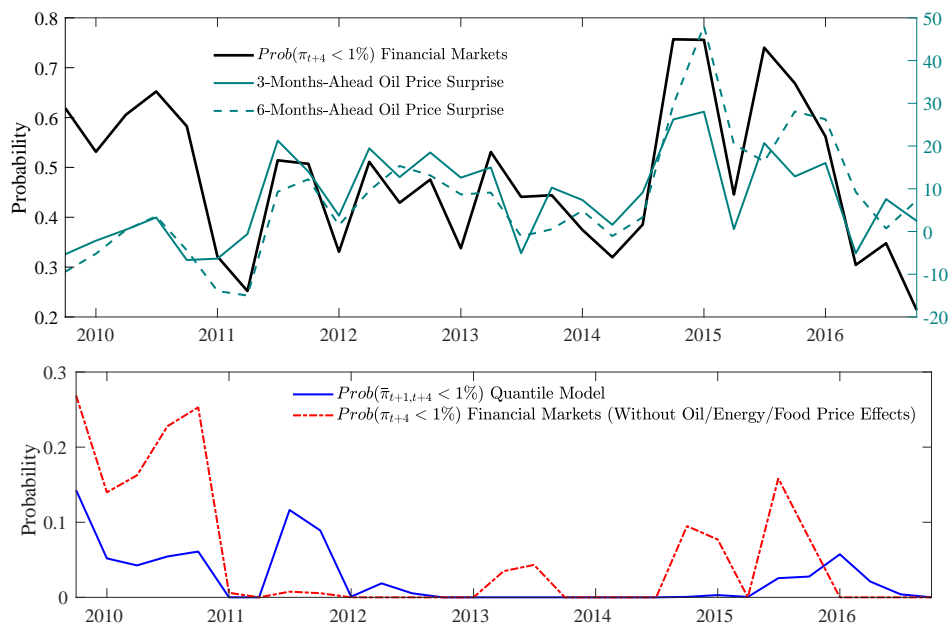
The top panel of Figure [F-1](#) substantiates this claim by displaying the inflation probabilities along with 3-months- and 6-months-ahead oil price surprises – computed using the oil market price expectations that [Baumeister and Kilian \(2016\)](#) recovered from oil futures prices and after controlling for changes in the risk premium.³⁰ As the figure shows, the options-implied inflation probabilities exhibits a high correlation with the oil-price surprises. Indeed, concerns about low inflation associated with the rise in the probability around mid 2014 to late 2015 coincide with a period in which financial markets have been steadily surprised to the downside in their oil price expectations.

²⁹The inflation probabilities are virtually identical if we consider one-year-ahead inflation instead of average one-year-ahead inflation, as in our quantile regression model.

³⁰The oil price surprises are computed as the difference between the market expectation of oil prices x -months ahead and the realized price of the West Texas Intermediate. While these surprises are not i.i.d. but rather feature some persistence, they still portray the actual surprise in oil price expectations of financial markets participants.

To improve comparability of the options-based *headline* CPI inflation probability with our measure of the tail of the *core* CPI distribution, we purify the financial markets' inflation probability from effects of changes in oil, energy and food prices. In particular, we regress it on the two oil price surprises as well as on energy and food price inflation, which also correlate with the options-based headline CPI inflation probability (see Figure F-2).³¹

Figure F-1: Inflation Probabilities, Quantile Model vs. Financial Markets.



NOTE: The top panel shows the options-implied inflation probabilities of United States headline CPI inflation next year being above 1% against the 3 and 6-months-ahead oil price surprises computed using the oil market price expectations that Baumeister and Kilian (2016) recovered from oil futures prices (top panel). The bottom panel shows the probability of average one-year-ahead core CPI inflation being below 1% coming from the quantile regression model as well as the probability of one-year-ahead headline CPI being below 1% implied from inflation caps and floors contracts as in Kitsul and Wright (2013), purified from oil, energy and food price effects and transformed to quarterly.

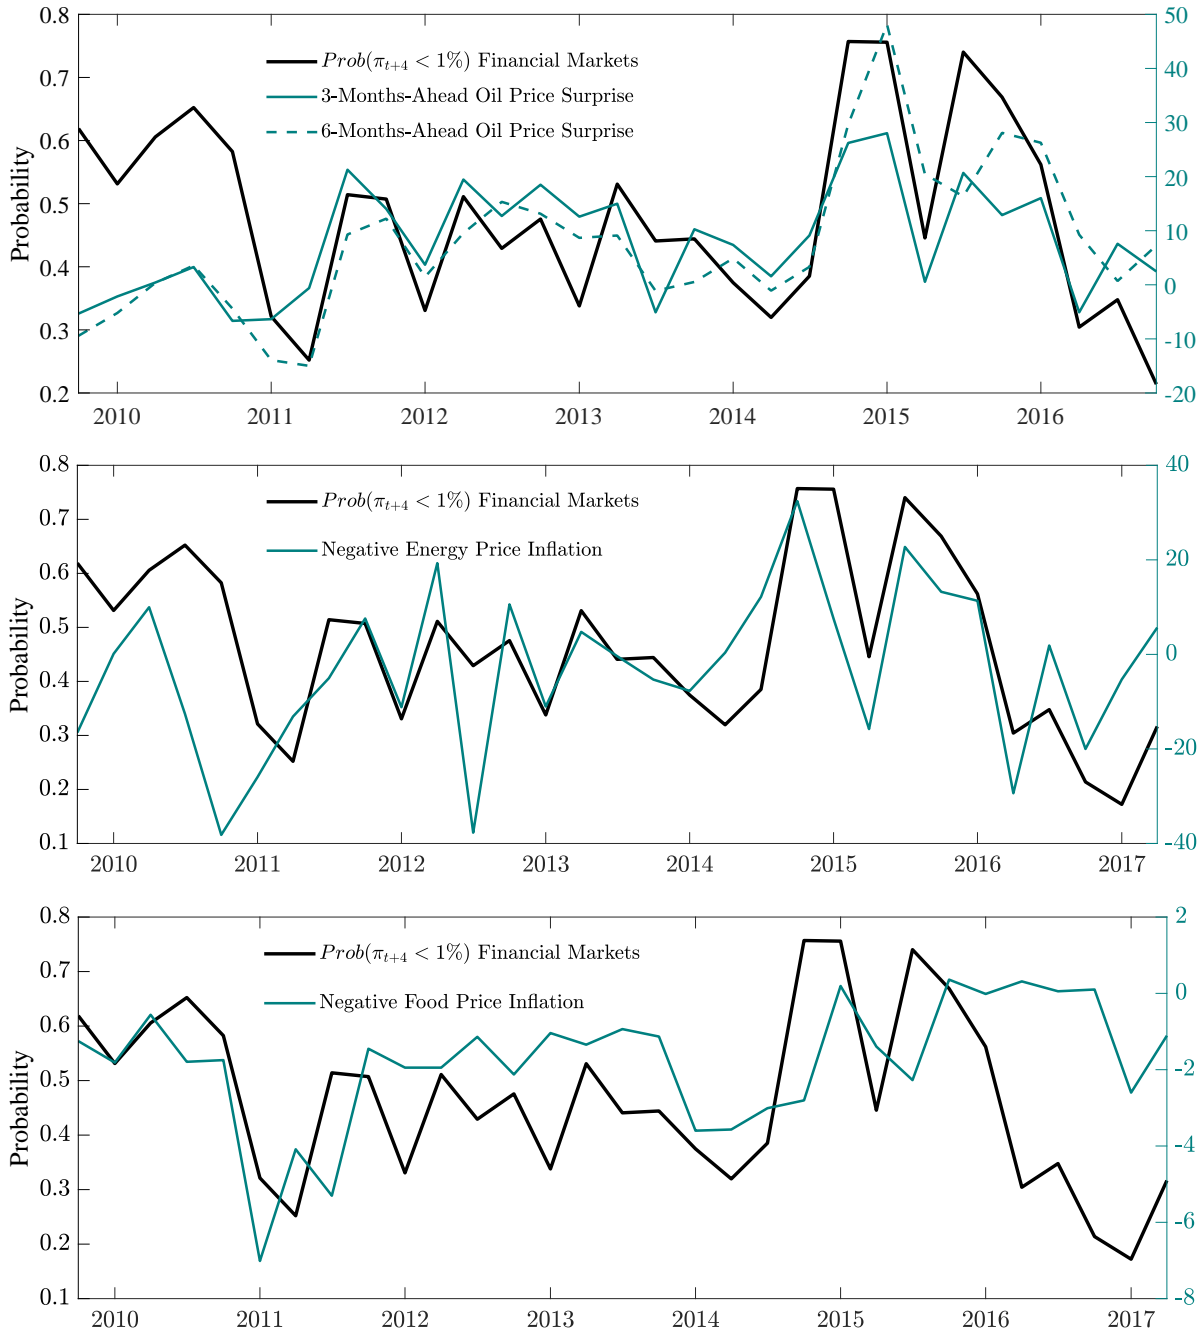
The dashed red line displayed at the bottom panel of Figure F-1 corresponds to the residual of this regression (where negative values have been set to zero). The bottom panel compares this purified financial-market-based probability with the probability of average future U.S. core CPI inflation being below 1% which arise from the quantile Phillips curve model (displayed in the top right panel of Figure K-2). The figure is very suggestive as it shows how both measures point

³¹Since the dependent variable of the regression is a probability which falls between zero and one, we estimate a generalized linear model with a logit link and the binomial family to ensure that the predicted values are between zero and one. A standard OLS regressions delivers virtually identical results.

in the same direction during most of the sample period. The probability of low inflation in the U.S. increased immediately after the global financial crisis and it subsequently falls to almost zero – remaining close to zero until the last quarter of 2016, with the exception of 2014/2015 when market participants consistently expected higher oil prices and when energy prices fell considerably. Accordingly, financial markets' expectations of headline CPI next year being below 1% rose accordingly during that time (see Figure F-2).

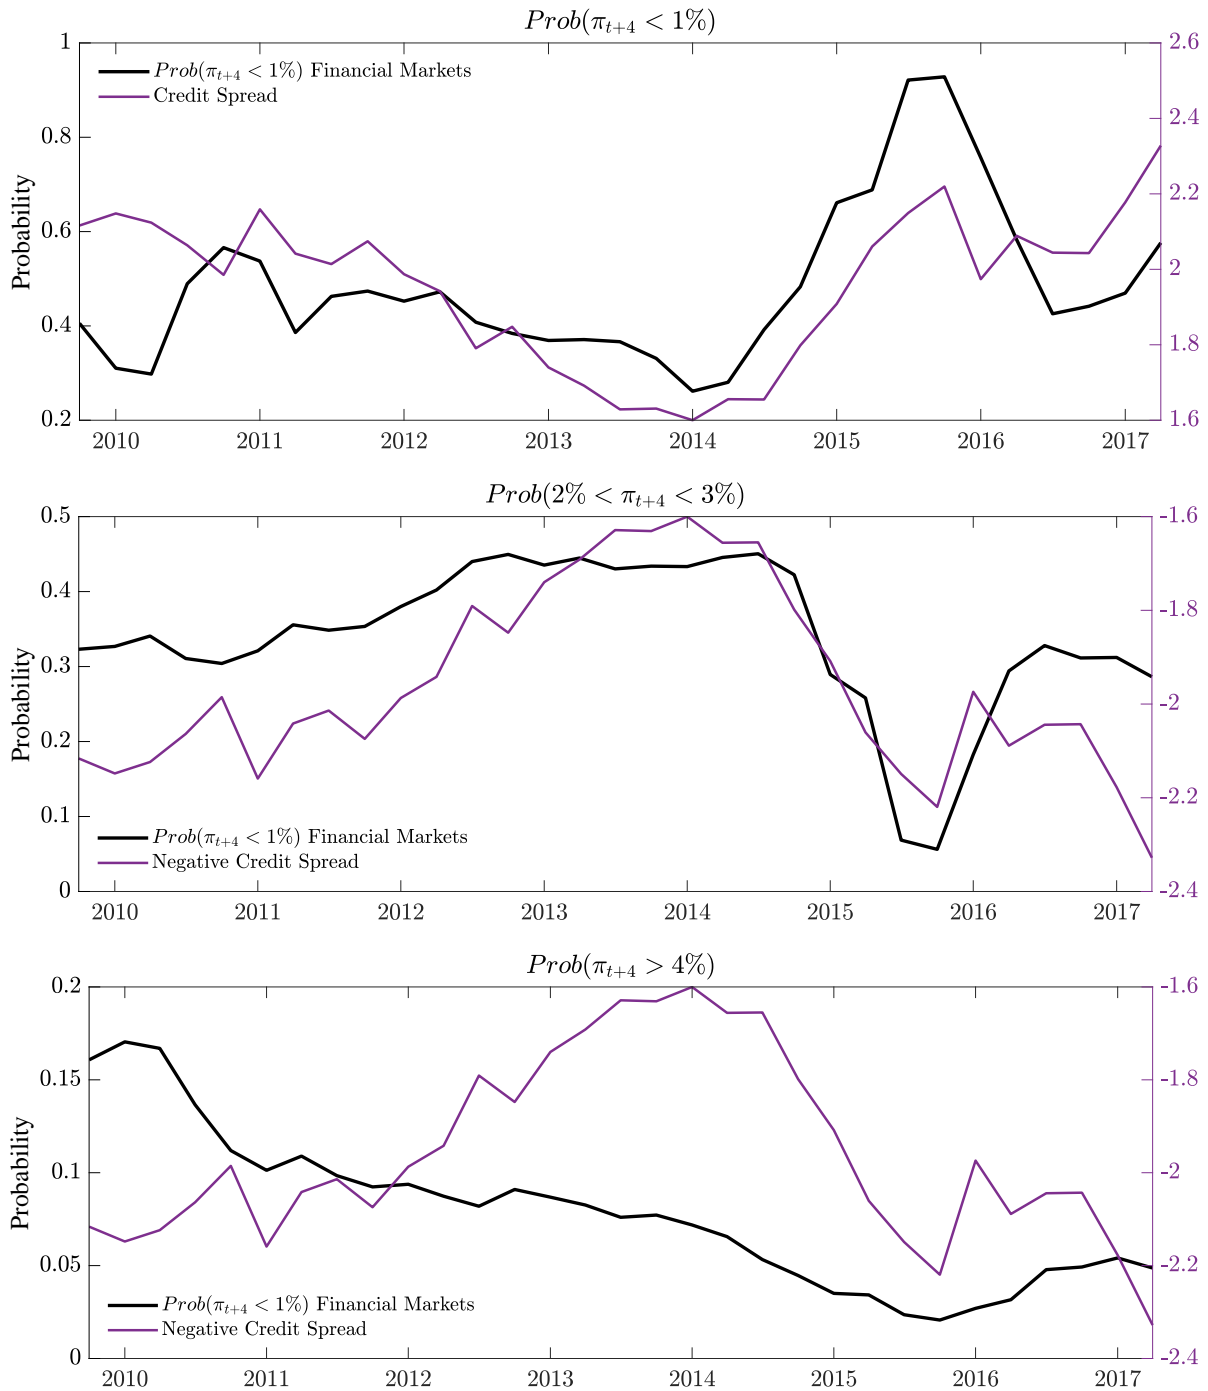
The small remaining differences can be explained by several factors. First, quantile-regression-based inflation probabilities come from a statistical model in which relative prices are estimated to play no role for the lower inflation tail whereas market participants seem to pay attention to the latter. More importantly, our regression purifies the financial markets' headline CPI inflation probabilities only from their average relationship with oil, energy and food prices – failing to fully capture times in which market participants strongly extracted information from these prices such that they comoved perfectly with the inflation probabilities (as in 2014/2015).

Figure F-2: Financial Markets' Inflation Probabilities vs. Oil, Energy and Food Price Measures.



NOTE: The figure shows the monthly options-implied inflation probabilities of headline CPI inflation next year being above 1% along with the 3-months- and 6-months-ahead oil price surprises computed using the oil market price expectations that [Baumeister and Kilian \(2016\)](#) recovered from oil futures prices (top panel), the negative of energy price inflation (mid panel) and the negative of food price inflation (bottom panel).

Figure F-3: Inflation Probabilities from Financial Markets vs. Credit Spread.



NOTE: The figures shows the credit spread against quarterly options-implied inflation probabilities of headline CPI inflation next year being above 1% (top panel), between 2% and 3% (mid panel) and above 4% (top panel).

G Regime-Switching Regression

Estimation Procedure We solve the model in the RISE toolbox³² using the perturbation methods developed by [Maih \(2015\)](#). The model is then estimated using Bayesian methods with prior hyperparameters specified in [Table G-1](#). We choose a Minnesota-type prior used in Dynare (see [Villemot and Pfeifer, 2017](#)) for the coefficients and a Dirichlet prior for the transition probabilities.

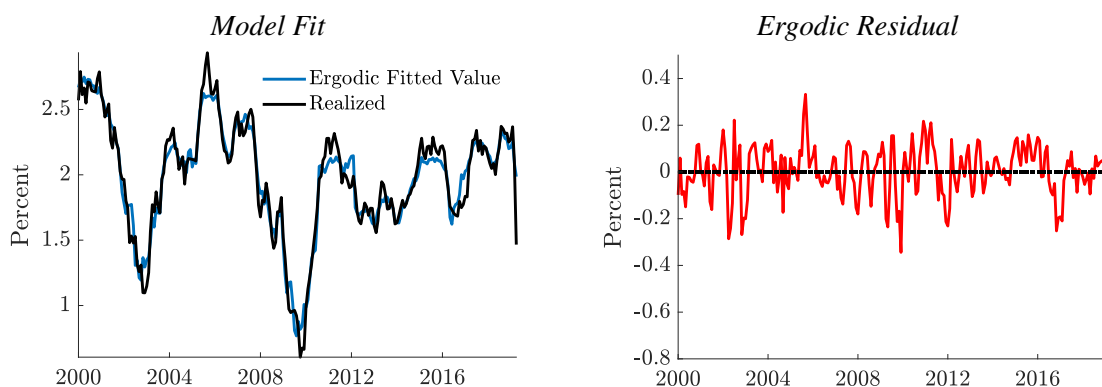
Table G-1: Prior Hyperparameters.

Coefficients: Minnesota-Type Villemot and Pfeifer (2017) Prior		
<i>Parameter</i>	<i>Description</i>	<i>Chosen Value</i>
τ	Overall tightness	1
d	Speed at which lags greater than 1 converge to zero	Not applicable since no lags
ω	Covariance dummies	3
λ	Co-persistence	No sum-of-coefficients
μ	Own-persistence	No dummy initial observations

Transition Probabilities: Dirichlet Prior		
<i>Parameter</i>	<i>Prior Mean</i>	<i>Prior Standard Deviation</i>
p_{12}, p_{21}	0.2	0.1
p_{13}, p_{31}	0.1	0.1
p_{23}, p_{32}	0.1	0.2

Model Fit We notice that this simple model is able to fit the data well, as shown in [Figure G-1](#), where we report the ergodic (average across regime) fitted values and residuals.

Figure G-1: Model Fit of Regime-Switching Regression.



NOTE: Ergodic fitted values, realized values and residuals of average core CPI inflation over the next year from regime-switching model (5).

³²The toolbox was developed by Junior Maih and is freely available at https://github.com/jmaih/RISE_toolbox.

H Bootstrap Method

To compute confidence bands for the quantile regression model we revert to “blocks-of-blocks” bootstrap. While more details on this methodology can be found in [Kilian and Lütkepohl \(2018\)](#) (see Chapter 12 therein), we here provide a brief summary of the bootstrap procedure.

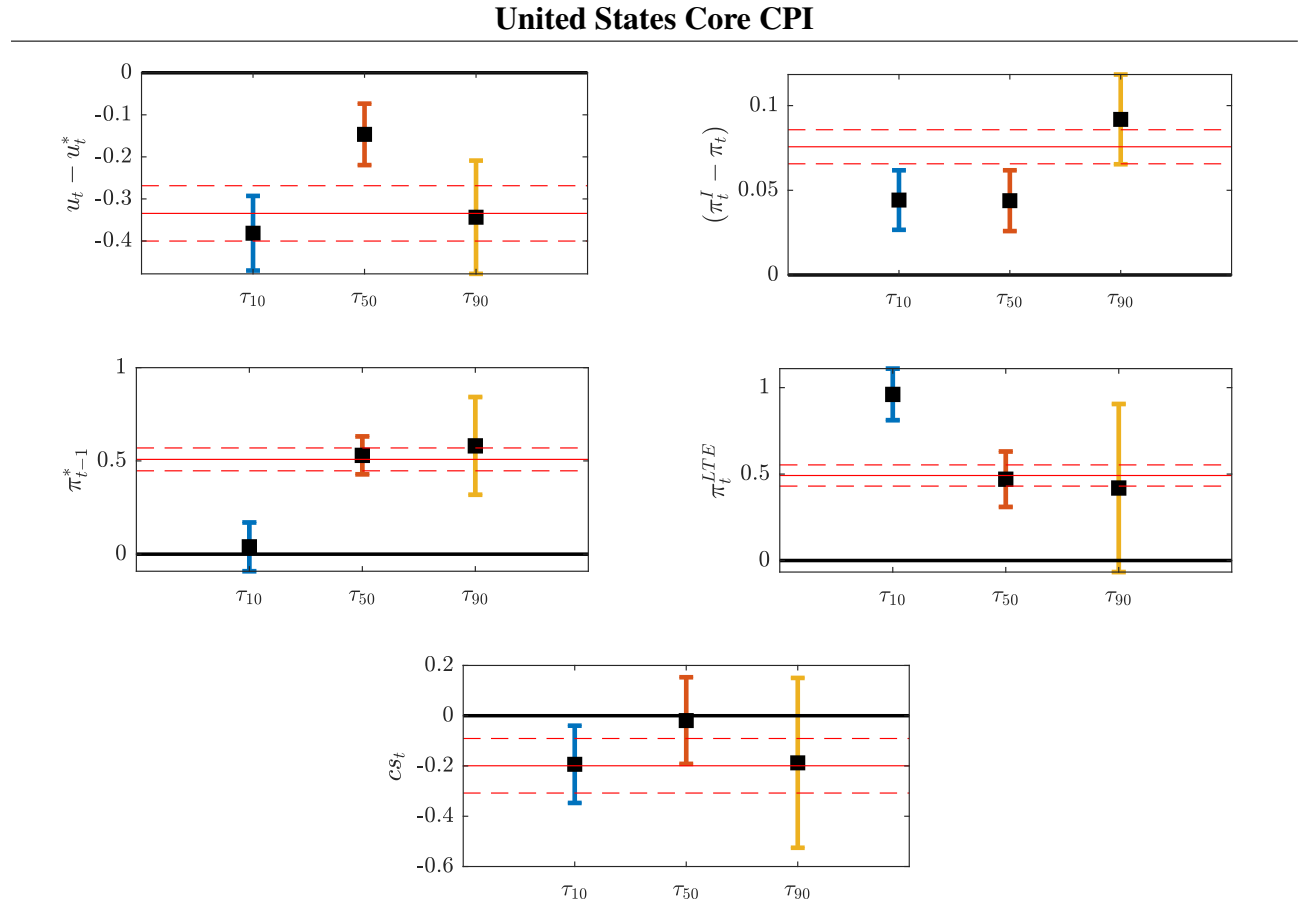
“Blocks-of-blocks” bootstrap is used in cases where a researcher is interested in computing confidence intervals around nonsymmetric statistics of the underlying data (e.g., autocorrelations or estimators of autoregressive slope coefficients in a time-series context). This is relevant in our case since not only the quantile regression slopes are non-linear functions of the data but also, we are de facto running a h -step predictive regression of inflation on its (past) determinants. The “blocks-of-blocks” bootstrap procedure allows to preserve the (time-series) dependency in the data, which would in most cases be destroyed by a naive bootstrap.

More specifically, the “blocks-of-blocks” bootstrap procedure relies on first dividing the dependent variable y and the regressors X into consecutive blocks of all possible m -tuples. At each bootstrap replication, blocks of data are randomly drawn to form a new sample of the same size as the original data. Importantly, the blocks are resampled in the same order for both the dependent variable y and the regressors X , a key step which preserves the time-dependency in the data. In our particular application, we run the quantile regression (2) and store the estimates corresponding to each bootstrap replication. From the distribution of these estimates, 68 percent confidence intervals are constructed and centered around the point estimate obtained with the original sample. The procedure is asymptotically valid for stationary processes if the block size l increases at a suitable rate as $T \rightarrow \infty$. Following [Berkowitz, Biegean, and Kilian \(1999\)](#) we set $m = \sqrt[3]{T}$, where T is the sample size. Finally, this bootstrap procedure preserves the quantile regression feature of being agnostic about the underlying distribution of the error terms, as this is not a residual-based procedure.

Figure [H-1](#) displays the slope coefficients of the quantile regression of average four-quarter-ahead United States Core CPI inflation defined in (2). The black squares correspond to the point estimates whereas the vertical lines to the 68% confidence intervals computed via “blocks-of-

blocks” bootstrap using 10,000 replications for the 10th quantile (blue), median (red) and 90th quantile (yellow). The estimation period is 1973:Q1 to 2019:Q1. The OLS estimates and their 95% confidence intervals are respectively represented by the solid and dashed red lines.

Figure H-1: Quantile Regression Slopes and Confidence Intervals.



NOTE: The figure displays the slope coefficients of the quantile regression of average four-quarter-ahead United States Core CPI inflation defined in (2). The black squares correspond to the point estimates whereas the vertical lines to the 68% confidence intervals computed via “blocks-of-blocks” bootstrap using 10,000 replications for the 10th quantile (blue), median (red) and 90th quantile (yellow). The estimation period is 1973:Q1 to 2019:Q1. The OLS estimates and their 95% confidence intervals are respectively represented by the solid and dashed red lines.

I Inspecting Other Inflation Measures

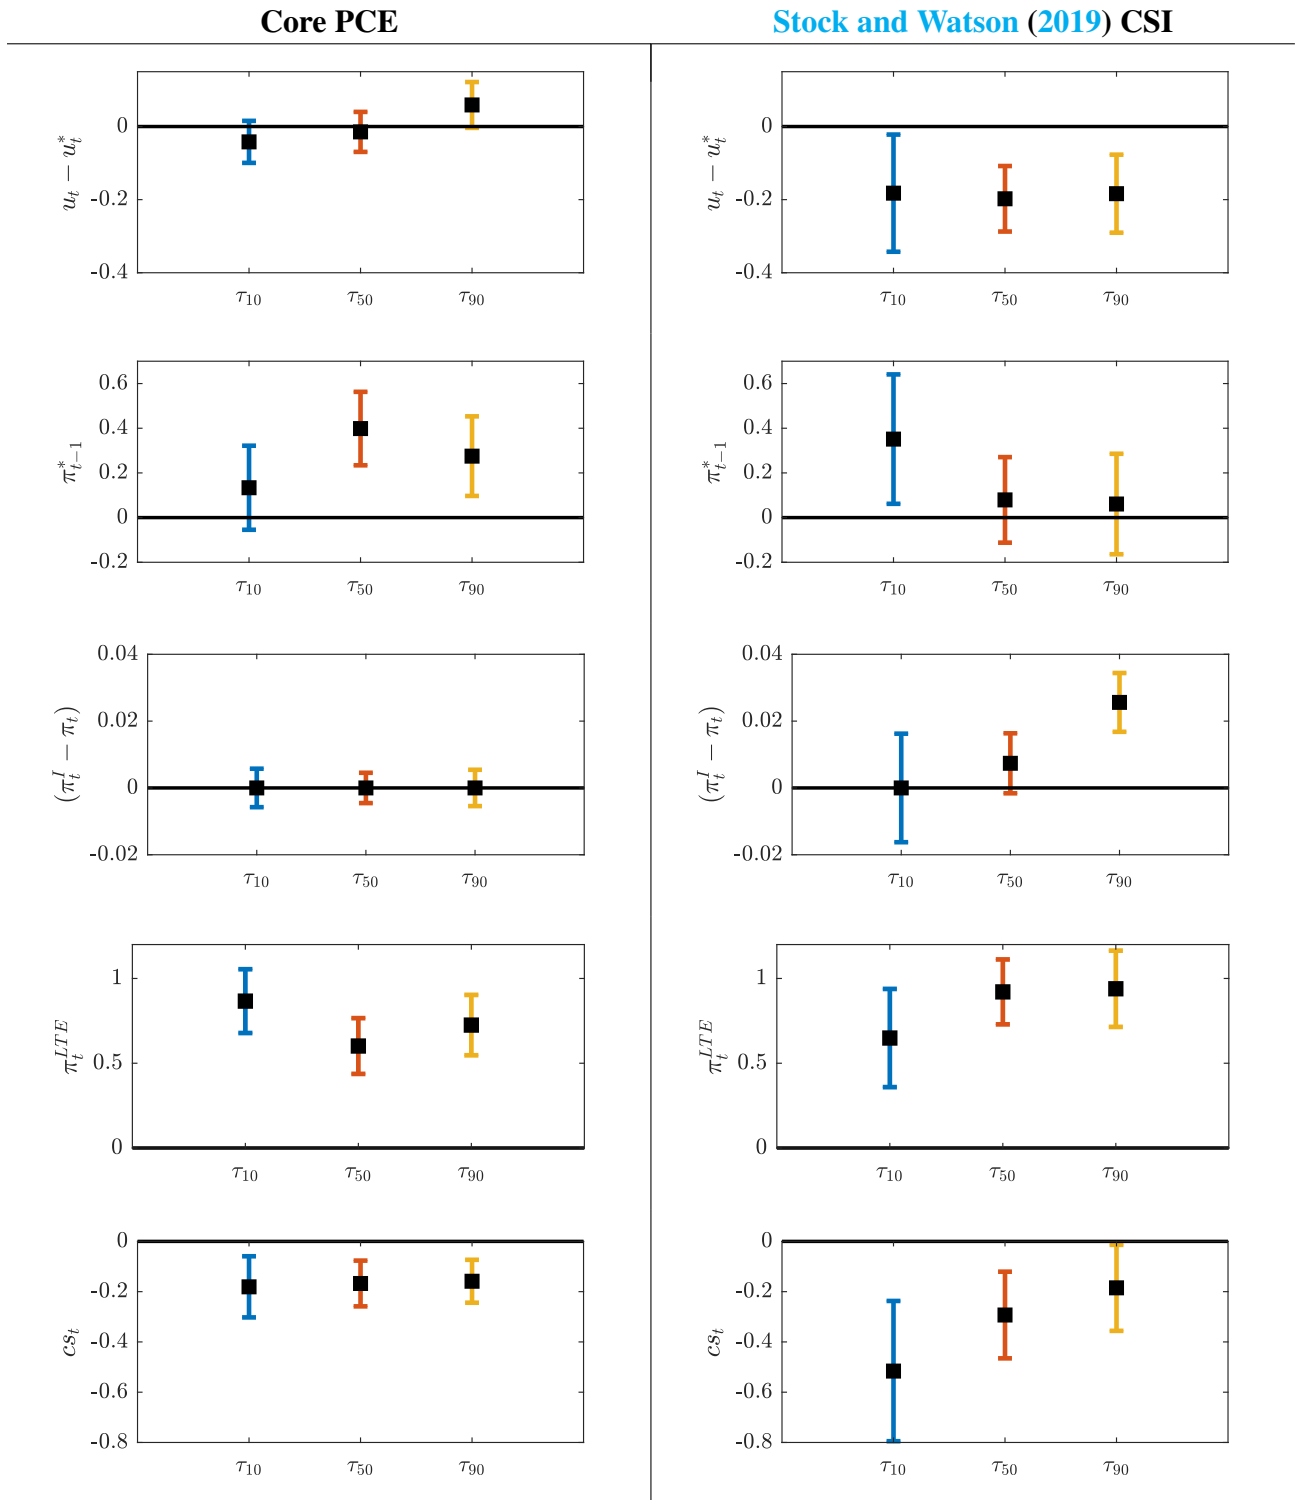
In this appendix we reestimate the quantile regression (2) replacing core CPI with two alternative measures of inflation: core PCE and [Stock and Watson \(2019\)](#) “Cyclically Sensitive Inflation”, CSI. As in the baseline analysis, the dependent variable is the average inflation rate over the period t and $t + 4$ quarters ahead. The CSI weights 17 core PCE components by their cyclical covariation with real activity. More specifically, the weights are computed so as to maximize the correlation between a composite index of cyclical activity (developed in the same paper) and the year-over-year change in the Cyclically Sensitive Inflation index. The CSI is thus meant to provide a real-time measure of cyclical fluctuations in inflation (see [Stock and Watson, 2019](#) for details).

Figure I-1 mirrors Figure 8 of the main text by displaying the estimated slopes of the quantile regression model (2) for two measures of inflation: core PCE (left column) and CSI (right column), along with their bootstrapped confidence intervals constructed as described in Appendix H. First, and not surprisingly, CSI is clearly more responsive to changes in unemployment, while core PCE is barely sensitive to labor market slack. The last row presents the role of credit spreads across inflation quantiles and inflation measures. The effects are more symmetric in the case of core PCE, while the CSI measure exhibits a similar asymmetry as core CPI although of somewhat larger magnitude.

Figure I-2 confirms the important influence of credit spreads on the 10th quantile of the distribution both for core PCE and for CSI inflation. This figure mimics the top-right panel of Figure 1 in the main text.

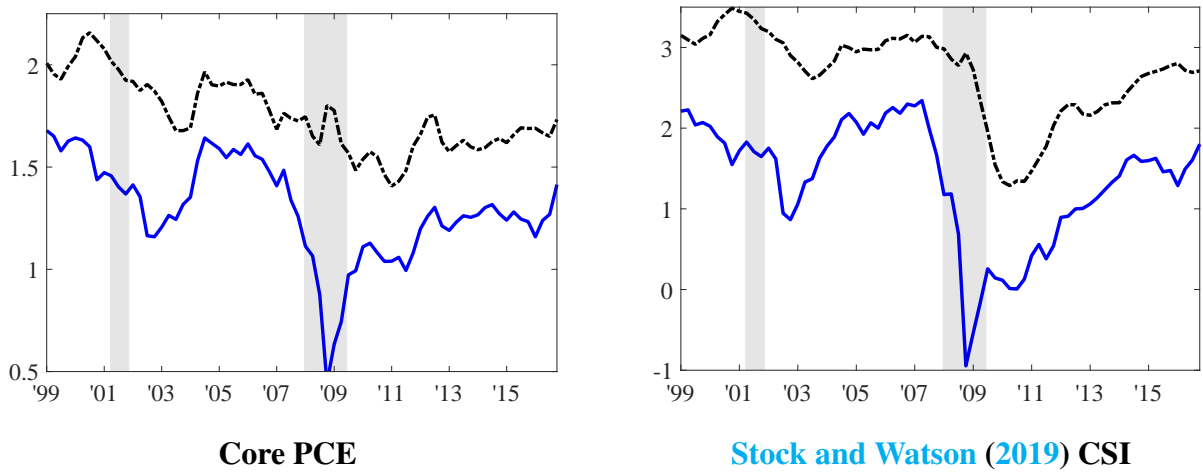
Figure I-3 displays similar exercises to those presented in the main text for these two alternative measures of inflation, core PCE and CSI, respectively.

Figure I-1: Quantile Regression Slopes and Confidence Intervals.



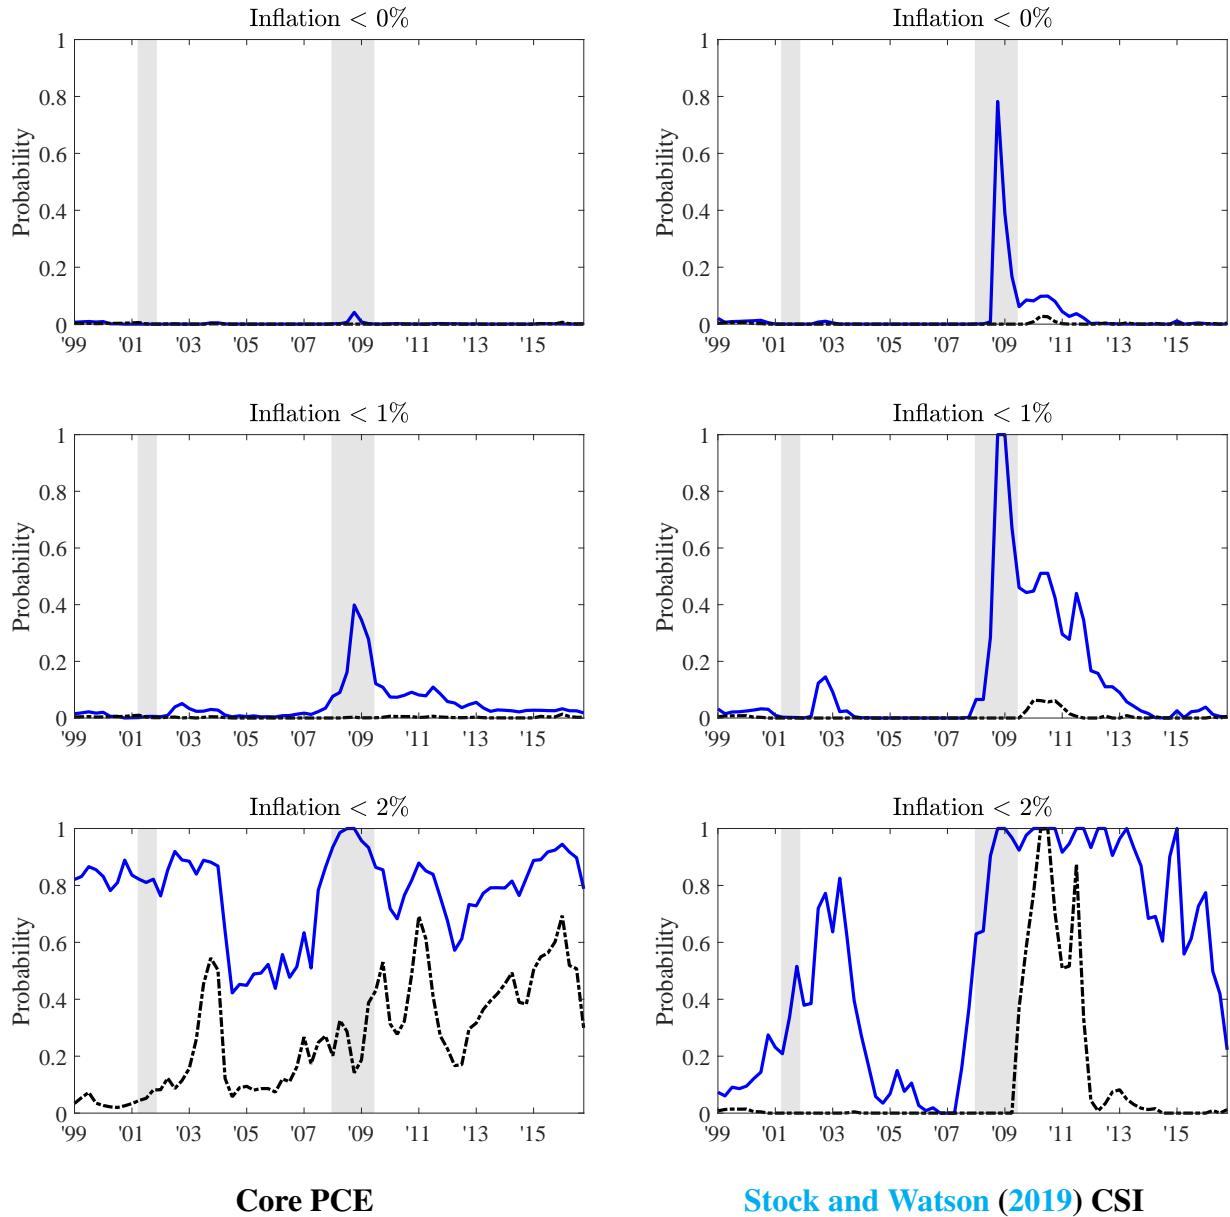
NOTE: The figure displays the slope coefficients of the quantile regression of average four-quarter-ahead of core PCE inflation (left) and Stock and Watson (2019) Cyclically Sensitive Inflation (right) defined in (2). The black squares correspond to the point estimates whereas the vertical lines to the 68% confidence intervals computed via “blocks-of-blocks” bootstrap (see Appendix H) using 10,000 replications for the 10th quantile (blue), median (red) and 90th quantile (yellow). The estimation period is 1999:Q1 to 2017:Q4.

**Figure I-2: Partial Effect of Credit Spread on 10th Inflation Quantiles.
Core PCE and [Stock and Watson \(2019\)](#) CSI Inflation.**



NOTE: The figure displays the time evolution of the 10th conditional inflation quantile of core PCE inflation (left) and [Stock and Watson \(2019\)](#) Cyclically Sensitive Inflation (right) estimated from the quantile regressions model (2), in its baseline version (blue solid) and in its version where the effect of credit spreads is set to zero (black dash-dotted). Shaded bars indicate NBER-dated recessions.

**Figure I-3: Inflation Probabilities for Alternative Cutoff Values.
Core PCE and [Stock and Watson \(2019\)](#) CSI Inflation.**



NOTE: The figure shows the time evolution of inflation probabilities of core PCE inflation (left) and [Stock and Watson \(2019\)](#) Cyclically Sensitive Inflation (right) for different cutoffs. These probabilities are computed from the skewed t -Student conditional densities of the average four-quarter-ahead inflation measures which were fitted on the estimated conditional quantiles for alternative specifications of the quantile regression model (2). Both panels are reported for the specification with and without the credit spread (in blue solid and black dash-dotted lines, respectively). Shaded bars indicate NBER-dated recessions.

J Historical Contributions

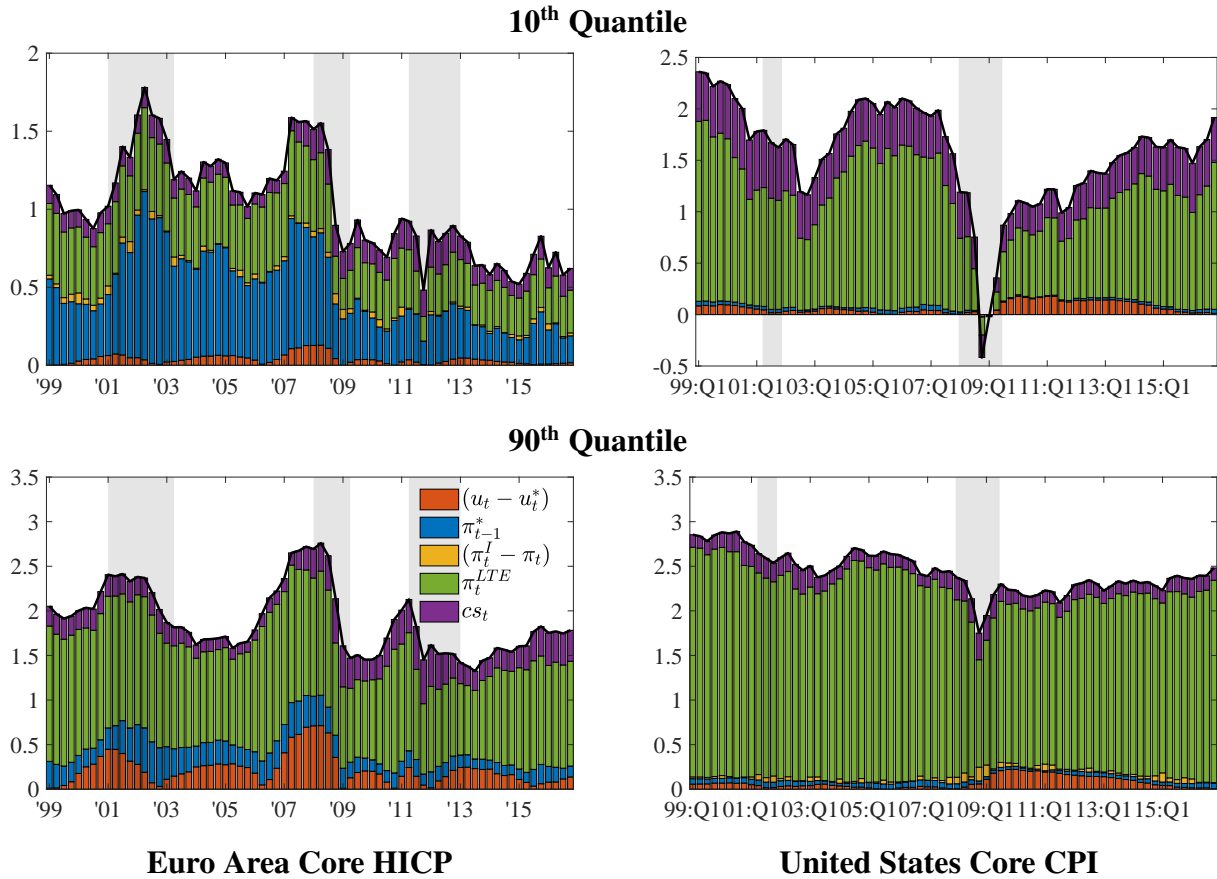
We explore which factors contributed to the recovery of the left tail in the case of the United States and to the failed recovery of the left tail in the case of the euro area. We thus next investigate the role of the inflation determinants from the Phillips-curve quantile model in influencing the inflation tails. In this regard, Figure J-4 complements the results in Figure 9 by presenting the contribution of economic and financial factors to changes in the lower and upper quantiles of the inflation distribution.

Focusing on the United States (the two charts in the right column of Figure J-4) it is striking how long-term inflation expectations have played a predominant role in sustaining the recovery of the left tail, supported to some extent by improvements in the labor market and more importantly by the easing in credit conditions. On average, across time, 66 percent of the variation in the upper quantile of the distribution is explained by changes in long-term inflation expectations, with the residual difference explained by financial conditions (27 percent), the unemployment gap (5 percent) and relative import price inflation (2 percent).

In the euro area, on the other hand, inflation expectations and labor market conditions had much less grip on downside inflation risks, with an average share of 37 percent and 4 percent respectively. Rather, average past inflation had the predominant role in holding down the lower inflation tail, its average share amounting to 42 percent. As in the U.S., financial factors played an important role also in the eurozone (15 percent share) and relative prices had no meaningful implications on the inflation outlook (2 percent share). A striking difference to the U.S. is how the lack of recovery in inflation expectations has driven most of the downward trend in the lower tail after 2012 – a tendency that diminished somewhat during 2016.

In the U.S. the upper inflation quantile is mainly dominated by changes in expectations, although high unemployment and persistently tight credit conditions have also contributed to make 2 percent an effective ceiling – yet another dent left by the global financial crisis. The same can be said about the eurozone with the important difference that financial conditions exerted a stronger downward pressure on the upper tail, which thus resulted in a lower implicit inflation ceiling.

Figure J-4: Historical Contributions of Economic and Financial Factors.



NOTE: The figure shows historical contributions of average four-quarter-ahead euro area core HICP (left) and United States core CPI inflation (right) associated with the quantile regressions model (2). The contribution of a given inflation determinant is obtained by multiplying its time series with its estimated slope. Its relative share is then obtained by weighting the contribution with its relative magnitude vis-à-vis the sum of all contributions (the share of the constant term is distributed across the inflation determinants based on their relative share so as to not distort results). Shaded bars indicate NBER-dated recessions for the United States and OECD-based recession indicators for the euro area.

K Inflation-at-Risk

We refer to “Inflation-at-Risk” (*IaR*) as the probability that inflation falls above or below a certain threshold. These risks are two-sided, with upside risks coming from “excessive inflation” and downside risks from too low or even negative inflation (i.e., disinflation or deflation).³³ There are two key elements that characterize our measure of *IaR*: (i) a pre-specified threshold, i.e., an upper (lower) level of inflation above (below) which inflation is “at risk” and (ii) a time period (say, $t + k$) over which the risk to the inflation outlook is assessed. These elements are necessary to substantiate statements such as: “With $(100-\tau)$ percent confidence we shall not experience, on average, inflation below (above) the level $\bar{\pi}^*$ over the next $t + k$ periods.”

The conditional downside inflation-at-risk

$$\mathcal{P}_t^D(\bar{\pi}_{t+1,t+4}|x_t) \equiv Prob(\bar{\pi}_{t+1,t+4} < \bar{\pi}^*|x_t) \quad (\text{K-2})$$

is the probability mass below $\bar{\pi}^*$ in the conditional density $f(\bar{\pi}_{t+1,t+k}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t)$:

$$\mathcal{P}_t^D(\bar{\pi}_{t+1,t+4}|x_t) \equiv \int_{-\infty}^{\bar{\pi}^*} f(\bar{\pi}_{t+1,t+k}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t) d\bar{\pi}_{t+1,t+k}, \quad (\text{K-3})$$

where at $(100-\tau)$ percent confidence, inflation will not be, on average, below the level $\bar{\pi}^*$ over the next $t + k$ periods. In other words, this expression defines the (downside) inflation-at-risk through the integral of the PDF over the inflation support up to a specified threshold (or the CDF).

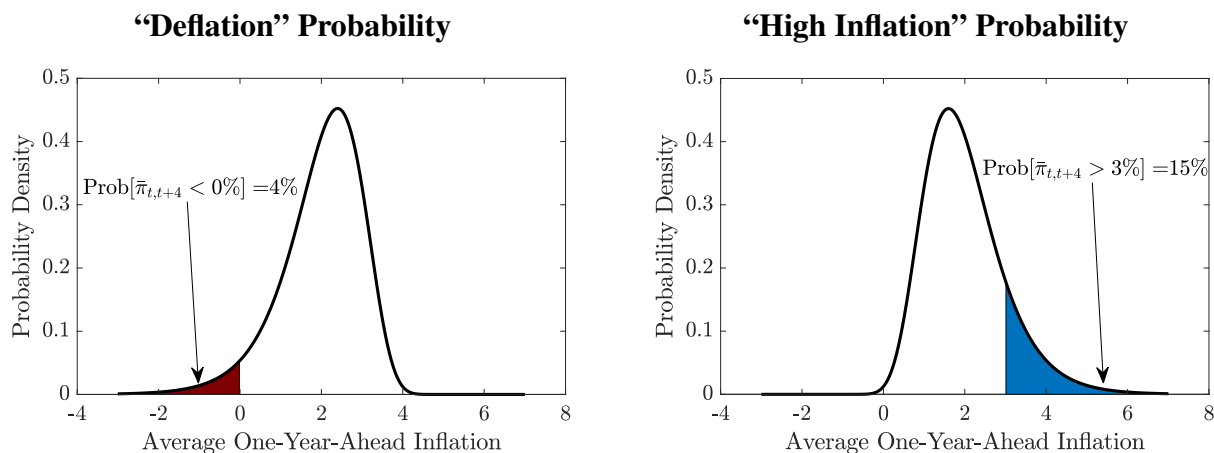
Figure K-1 illustrates the link between *IaR* and the quantiles of the inflation distribution. Downside risks to inflation can be characterized by the probability mass to the left tail of the distribution (left panel). The red area indicates that at 4 percent confidence level, inflation at risk is “zero percent”. Or, equivalently, that a zero (or below) inflation rate corresponds to the 4th quantile of the inflation distribution. Similarly, the right panel illustrates that, with a 15 percent probability,

³³Our approach differs from the *Value-at-Risk* literature in one key way. In that literature, $VaR(\tau)$ is not a probability but the threshold such that the probability of future returns (not) exceeding that threshold is equal to τ . In that sense, $VaR(\tau)$ is the τ^{th} quantile of future returns. Formally, according to that definition, inflation-at-risk $IaR(\tau)$ is thus the τ^{th} conditional inflation quantile, $\mathcal{Q}_\tau(\bar{\pi}_{t+1,t+4}|x_t)$, implicitly defined by the integral over the conditional inflation density $f(\bar{\pi}_{t+1,t+k}|x_t)$ that sums up to τ :

$$\int_{-\infty}^{\mathcal{Q}_\tau(\bar{\pi}_{t+1,t+4}|x_t)} f(\bar{\pi}_{t+1,t+k}|x_t) d\bar{\pi}_{t+1,t+k} = \tau. \quad (\text{K-1})$$

average future inflation can be above 3 percent – in other words, the upside (tail) risk associated with “excessive inflation” is 15 percent. More generally, measuring τ th-percent *IaR* is akin to estimating the τ th quantile of the probability distribution of inflation (or its outlook).

Figure K-1: Inflation-at-Risk.



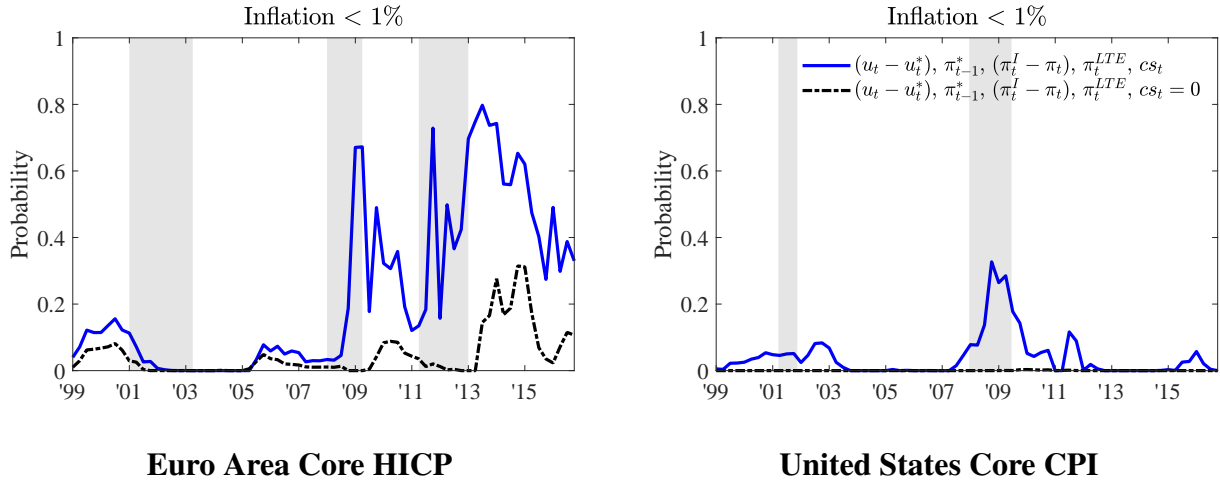
NOTE: The figure displays simulated distributions. In the left panel, the probability of average future inflation falling below 0% is 4 percent. In the right panel the probability of average future inflation exceeding 3% is 15 percent.

K.1 Inflation Probabilities

We now show how credit conditions affect the odds of low inflation. Figure K-2 displays, beginning in 1999, our estimates of the evolution of the probability of observing inflation rates below 1 percent over the next four quarters. The two columns are used to contrast the eurozone (left column) with the U.S. (right column). In each panel, we display the probabilities computed using our baseline model (blue solid line) and its version which omits the effects attributable to changes in credit conditions (black dash-dotted line).

Several conclusions emerge from these comparisons. Since 2000, the model omitting credit conditions would have assigned zero probability to inflation running below 1 percent in the U.S., whereas accounting for the financial meltdown had profound effects on the inflation outlook – with the probability of very low inflation (and deflation) temporarily reaching almost 40 percent in the last quarter of 2008 (upper right panel). Results for the eurozone are more striking on this account.

Figure K-2: Inflation Probabilities.



NOTE: The figure shows the time evolution of the probability of average one-year-ahead euro area core HICP inflation (left) and United States core CPI inflation (right) falling below 1%. These probabilities are computed from the skewed t -Student conditional densities of the average four-quarter-ahead inflation measures which were fitted on the estimated conditional quantiles for alternative specifications of the quantile regression model (2). Both panels are reported for the specification with and without the credit spread (in blue solid and black dash-dotted lines, respectively). Shaded bars indicate NBER-dated recessions for the U.S. and OECD-based recessions for the euro area.

Changes in the credit spreads in 2008-2009 and especially in late 2011 induced sharp increases in the odds of very low inflation and a remarkable divergence between the blue and the dash-dotted lines in the top-left panel of Figure K-2. By early-2014, this probability was slightly above 80 percent when the model includes financial variables, while it was around 30 percent in the model accounting for the effects of non-financial variables only.

K.2 The Role of Skewness

In this section, we show that tail risks to the inflation outlook are amplified following a change in an inflation driver if the latter not only shifts but also skews the distribution. We thus argue that since credit spreads are the predominant factor in the modern Phillips curve which introduces skewness in the inflation distribution, changes in credit conditions make the inflation outlook particularly vulnerable – more than any other inflation driver considered in conventional Phillips curves.

We start by characterizing the derivative of downside inflation-at-risk $\mathcal{P}_t^D(\bar{\pi}_{t+1,t+4}|x_t)$ defined

in (K-3), for a given probability cutoff $\bar{\pi}^*$, with respect to an inflation determinant x_t . Formally,

$$\frac{\partial \mathcal{P}_t^D(\bar{\pi}_{t+1,t+4}|x_t)}{\partial x_t} = \frac{\partial}{\partial x_t} \int_{-\infty}^{\bar{\pi}^*} f(\bar{\pi}_{t+1,t+k}, \mu_t, \sigma_t, \eta_t, \kappa_t | x_t) d\bar{\pi}_{t+1,t+k}, \quad (\text{K-4})$$

where we abstract from the dependence of the parameters $\mu_t, \sigma_t, \eta_t, \kappa_t$ on x_t .

Applying the Leibniz integral rule,

$$\frac{\partial \mathcal{P}_t^D(\bar{\pi}_{t+1,t+4}|x_t)}{\partial x_t} = \int_{-\infty}^{\bar{\pi}^*} \frac{\partial f(\bar{\pi}_{t+1,t+k}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t)}{\partial \bar{\pi}_{t+1,t+k}(x_t)} \frac{\partial \bar{\pi}_{t+1,t+k}(x_t)}{\partial x_t} d\bar{\pi}_{t+1,t+k}, \quad (\text{K-5})$$

and assuming a linear regression quantile model for the mean of $\bar{\pi}_{t+1,t+k}(x_t)$ simplifies to:

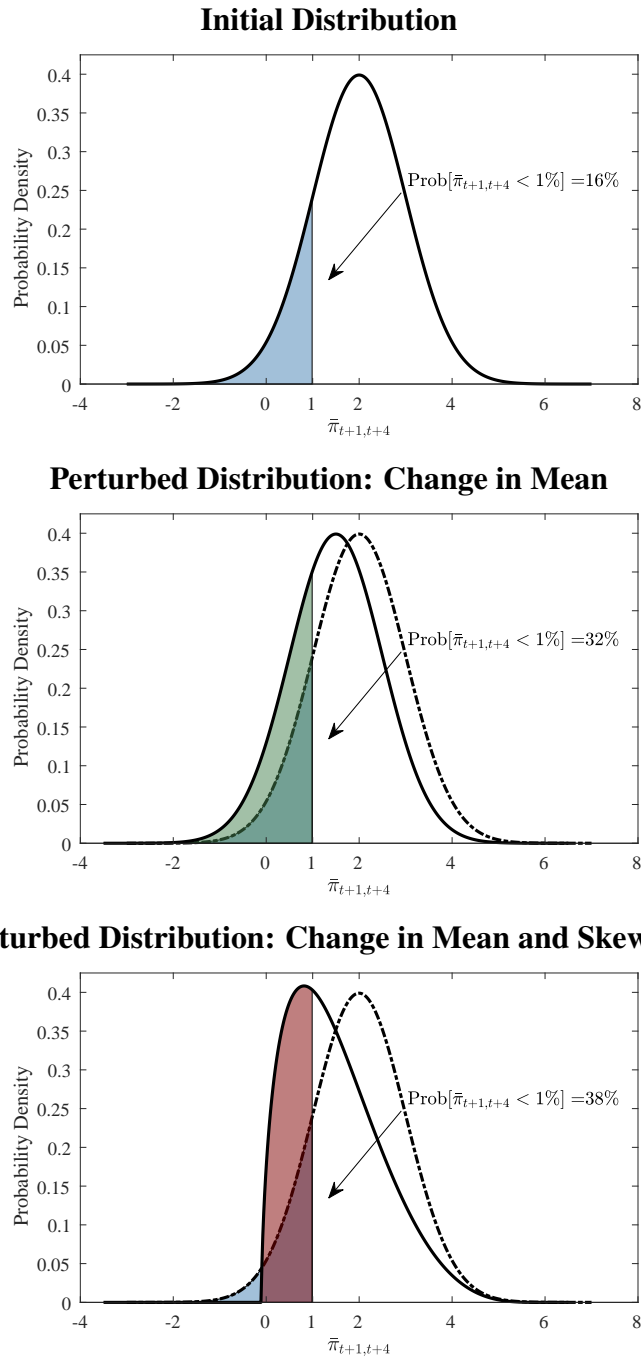
$$\frac{\partial \mathcal{P}_t^D(\bar{\pi}_{t+1,t+4}|x_t)}{\partial x_t} = \beta_{OLS} \int_{-\infty}^{\bar{\pi}^*} \frac{\partial f(\bar{\pi}_{t+1,t+k}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t)}{\partial \bar{\pi}_{t+1,t+k}(x_t)} d\bar{\pi}_{t+1,t+k}. \quad (\text{K-6})$$

From expression (K-6) it follows that changes in any variable x_t , besides affecting linearly the quantile of the distribution of inflation, introduces a “nonlinear” effect on downside inflation-at-risk. The first effect captures how a change in x_t scales (up or down) the support of the inflation distribution. The strength of this channel is measured by the coefficient β_{OLS} . This first effect gets amplified by the second term that cumulates the derivatives of the conditional density function with respect to its support up to cutoff level $\bar{\pi}^*$. It thus follows that the more right-skewed the distribution is (i.e., the more mass is on its left tail) at a given point in time, the stronger the density changes in the left part of the support and, in turn, the bigger the effect on downside risk caused by a change in x_t .

An Illustrative Example In Figure K-3 we illustrates the effect of a change in an inflation determinant on the probability of average one-year-ahead inflation falling below 1% (downside inflation-at-risk). The initial (normal) density is illustrated in the top panel. In this thought experiment, the change in economic/financial conditions induces a change in the mean (center panel) and then also in the skewness of the distribution (bottom panel). It is evident how the effect on

downside inflation-at-risk is amplified if the change in the inflation determinant increases the right-skewness of the distribution.

Figure K-3: Inflation Probabilities and The Role of Skewness.



NOTE: The figure displays the three states associated with a change in an inflation determinant that causes the initial normal density (top panel) to feature a lower mean (center panel) and then also a right-skew (bottom panel).