Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

Inflation at Risk

David López-Salido and Francesca Loria

2020-013

Please cite this paper as:
López-Salido, David, and Francesca Loria (2020). "Inflation at Risk," Finance and Economics Discussion Series 2020-013. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2020.013.

Inflation at Risk

David López-Salido Francesca Loria

Federal Reserve Board

February 6, 2020

Abstract

We investigate how macroeconomic drivers affect the *predictive* inflation distribution as well as the probability that inflation will run above or below certain thresholds over the near term. This is what we refer to as *Inflation-at-Risk* – a measure of the tail risks to the inflation outlook. We find that the recent muted response of the *conditional mean* of inflation to economic conditions does not convey an adequate representation of the overall pattern of inflation dynamics. Analyzing data from the 1970s reveals ample variability in the conditional *predictive* distribution of inflation that remains even when focusing on the post-2000 period of stable and low *mean* inflation. We also document that in the United States and in the Euro Area tight financial conditions carry substantial downside inflation risks, a feature overlooked by much of the literature. Our paper offers a new empirical perspective to existing macroeconomic models, showing that changes in credit conditions are also key to understand the dynamics of the inflation tails.

JEL Classification: C21, E31.
Keywords: Quantile Regression, Inflation Risks.

Special thanks to Ben Bernanke, Steve Cecchetti, David Cho, Jim Clouse, Deepa Datta, Giovanni Favara, Felix Galbis-Reig, Ed Herbst, Paul Lengermann, Antoine Lepetit, Ed Nelson, Claudia Pacella, Jeremy Rudd, Tatevik Sekhposyan and Srečko Zimic as well as seminar participants of the Research & Statistics Lunch Seminar at the Federal Reserve Board and of the DG-E Research Seminar at the European Central Bank 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.

*"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]

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. Yet, much of the economic analysis has 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 past and current labor market conditions to explain recent inflation outcomes. The Phillips curve linkages seem to be breaking down. In this paper, we find that some of the macroeconomic factors covered under the "Phillips curve umbrella" are still at work in the *tails* of the inflation distribution. Moreover, we also show that looking at the entire inflation distribution uncovers – after controlling for the state of the labor market and inflation expectations – that tight financial conditions carry substantial low-inflation risks, an aspect of inflation behavior overlooked by much of the literature.

President Draghi's quote is an excellent reminder that, in the presence of tail risks, the conditional inflation mean does not necessarily adequately represent the inflation outlook. Indeed, it has been documented that 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) – with its left tail being the most sensitive to macroeconomic shocks (see Loria, Matthes, and Zhang, 2019a) – and implied medium-term upside risks to unemployment (see Kiley, 2018).[2] Yet, against this backdrop, the modal outlook for inflation seemed to remain somewhat insensitive to these developments. We show that the contrasting response of the tails and the median of the inflation distribution reveals a more complete picture of the effects that real and financial shocks impinge on inflation. The objective of this paper is to see what conclusions can be drawn from a closer look at the entire conditional inflation distribution, using data both for the United States and the euro area.

We investigate how macroeconomic drivers affect the *predictive* inflation distribution as well as the probability that inflation will run above or below certain thresholds over the near term. This is what we refer to as *Inflation-at-Risk* – a measure of the tail risks to the inflation outlook.[3] Our econometric strategy first constructs inflation quantiles conditioning on observed economic and financial variables ("risk factors") using a quantile regression model, where the commonly used approach is to fit a linear curve (e.g., Koenker, 2005).[4]

[1]Mario 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).

[2]In a cross-section of countries, Cecchetti (2008) finds that asset price booms increase growth and inflation risks.

[3]Readers may glean some lack of novelty in this label, as we adapted this name from the Value-at-Risk literature.

[4]As shown in the literature, this is done mainly because linear models enjoy good approximation properties as well as desirable computational properties (e.g., Chernozhukov, Fernandez-Val, and Galichon, 2010).

We frame the effects of different risk factors on inflation within an "augmented" quantile Phillips curve model using data since the 1970s.[5] That is, we extend the standard regression analysis – designed to ascertain the drivers of the conditional mean of inflation – to different inflation quantiles.[6] This setup allows to relate the risks to the inflation outlook to variations in labor market slack, and changes in the persistence of inflation and inflation expectations, as well as movements in relative prices (imported goods and/or oil). More importantly, we extend the analysis to consider the effect of financial conditions on the inflation distribution and on the odds of low inflation.

The conditional distribution of future inflation is constructed by fitting a flexible distribution on the estimated conditional inflation quantile distribution. Thus, in our econometric approach, variations in the conditional inflation distribution depend on the evolution of economic and financial factors and how they (a)symmetrically affect the inflation quantiles. These variations are not limited to a change in the mean and in the variance. Indeed, we find that periods featuring a relatively tight and centered inflation distribution evolve into periods in which the tails of the distribution increase substantially, leading to a change in the kurtosis of the distribution; or to periods in which the symmetry of the distribution is tilted towards the left (or the right), leading to a change in the skewness. Finally, our framework links the persistence in the evolution of economic and financial conditions to the persistence, not just of the conditional mean, but also of the inflation tails. This is illustrated with historical contributions of economic and financial conditions to the evolution of the median, the lower and the upper quantiles of inflation.

As noted above, our "augmented" quantile Phillips-curve model considers changes in credit spreads as an additional factor affecting the entire inflation distribution. In this regard, the recent global financial crisis is an ideal case study for illustrating how both economic and financial headwinds influenced the inflation outlook both in the United States and in the euro area.

In the United States, average inflation experienced only a modest reduction despite the fall in output triggered by the financial crisis. Likewise, it featured only a shy rebound despite the recovery and subsequent growth of the U.S. economy. The inability of conventional economic wisdom – derived from the historical Phillips-curve relationship between inflation and economic conditions – to explain this phenomenon is referred to by economists as the "missing disinflation and inflation puzzle". There are two lead explanations supporting this lack of response in average inflation. First, inflation expectations were well anchored (Yellen, 2013). Second, as financial shocks increased the cost of external finance, liquidity-constrained firms have restrained from cutting price below marginal cost to support their cash-flow and thus hedge against these risks (Gilchrist, Schoenle, Sim, and Zakrajšek, 2017).

[5]Our approach hence 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. An important departure from our approach is that, in general, they lack an explicit link of these risks to specific macroeconomic outcomes.

[6]Korobilis (2017) finds that the predictive densities coming 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. Manzan and Zerom (2013) find that incorporating macroeconomic variables into quantile regressions improves the accuracy of inflation density forecasts.

We find that this view is incomplete, as it ignores how inflation risks moved over this period. Indeed, to anticipate some of our results, we show that although the conditional mean of inflation held up during the crisis, the inflation distribution shifted to the left and was characterized by a substantial left tail. As a result, the odds of very low inflation or even deflation increased. These odds stayed at considerable levels for some time: the downside risks to the inflation outlook were quite persistent, as they only vanished well into the recovery. Most interestingly, their dramatic increase was mainly the reflection of soaring credit spreads during the financial meltdown. Subsequently, stable long-term inflation expectations sustained the recovery of the left tail of the inflation distribution, accompanied by easing of credit conditions (as well as, to a lesser extent, improvements in the labor market).[7]

These patterns have been less benign in the euro area, where the sovereign debt crisis triggered an increase in the odds of low inflation which was more prolonged to the more limited role of inflation expectations in counteracting downside risks to inflation posed by the economic slowdown and financial distress.

New research by Christiano, Motto, and Rostagno (2014), Christiano, Eichenbaum, and Trabandt (2015), Del Negro, Giannoni, and Schorfheide (2015), and Gilchrist, Schoenle, Sim, and Zakrajšek (2017) argues that, in models with financial frictions, firms' financial conditions help to explain inflation dynamics. However, these papers have almost exclusively focused on explaining the response of the conditional mean of inflation and hardly paid any attention to how financial conditions affect the tails of the inflation distribution. Our paper offers a new empirical perspective on these issues, one that shows that changes in credit conditions are key to understand the tail-risk dynamics of inflation.

Finally, we check whether the distribution of United States inflation embodied in financial options is consistent with some of the conclusions about inflation risks derived from our analysis. Not only do inflation probabilities coming from financial markets and from our quantile Phillips curve model point in the same direction but they also share a defining feature, namely that tight financial conditions carry substantial downside inflation risks and most strongly so for the left tail.

Outline The paper is structured as follows. In Section 2 we organize ideas by presenting our theoretical framework and empirical strategy. We then study the role of economic and financial conditions for the risks to the United States inflation outlook in Section 3 using our full sample. As time-variation emerges in the characterization of the determinants of the inflation distribution, we illustrate subsample results in Section 4 and use them to shed new light on the "missing disinflation and inflation" debate. In Section 5 we then compare the United States and euro area inflation experiences in the last 20 years and explore the role of financial conditions in affecting the odds of low inflation during and after the global financial crisis. We then perform external validation of our approach in Section 6. Concluding remarks and policy implications are offered in Section 7.

[7]As Yellen (2013) noted: "After the onset of the financial crisis, these stable [long-run inflation] expectations also helped the United States avoid excessive disinflation or even deflation."

2 Quantile Regressions and 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" to the inflation outlook 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 (usually, 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.

2.1 (Phillips-Curve) Quantile Regressions

Quantile regression models are a flexible tool for studying the determinants of *IaR*.[8] Our inflation measure of interest is the (annualized) average inflation rate between quarter t and quarter $t+4$, $\bar{\pi}_{t,t+4}$.[9] We consider a linear model for the conditional inflation quantiles whose predicted value

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

is a consistent linear estimator[10] of the quantile function of $\bar{\pi}_{t,t+4}$ 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. Accordingly, a determinant x_t may exert non-linear effects on inflation dynamics if it affects differently the median and the tails.

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

[9] A similar approach is taken in Adrian, Boyarchenko, and Giannone (2019) for the average growth rate of GDP. An alternative approach is taken by Ghysels, Iania, and Striaukas (2018) who use a Quantile Autoregressive Distributed Lag Mixed-Frequency Data Sampling (QADL-MIDAS) regression model to construct measures of inflation risk.

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

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

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

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) – a recent paper that nicely 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:

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

where risk factors affecting the distribution of future inflation can be divided in different blocks.[11]

First, the variables π^*_{t-1} and π^{LTE}_t 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.[12] 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.[13]

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 $(\pi^R_t - \pi_t)$, which reflects variations in relative prices. We use the quarterly change in relative import prices $(\pi^I_t - \pi_t)$. 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).

[11] A full description of the data is provided in Appendix A.

[12] 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.

[13] To preserve the notion that actual 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 quantile model becomes an extension of the accelerationist Phillips curve, where changes in inflation are a function of the unemployment gap. We impose $(1-\lambda_\tau)+\lambda_\tau = 1$, $0 \le (1-\lambda_\tau) \le 1$ and $0 \le \lambda_\tau \le 1$ following Koenker and Ng (2005).

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.

In this sense, these papers do not completely analyze the potential nonlinearities in the response of inflation to financial distress as a useful diagnostic of the models. For instance, Gilchrist and Zakrajšek (2015) and Gilchrist, Schoenle, Sim, and Zakrajšek (2017) noted that there is still room for changes in firms' financial conditions to influence not just average inflation, but also the distribution of inflation. In particular, these authors observed that a rich heterogeneity in price settings will arise from changes in firms' credit conditions: When external financing is expensive, a liquidity constrained firm may have to sacrifice current and future demand to get liquid funds today by not lowering prices in response to tighter credit conditions; but firms with abundant liquidity may cut prices substantially.[14] Some of these distributional effects may be reflected not just in the median but also in the tails of the aggregate inflation distribution when the model is solved at third (or higher) order so as to capture changes in skewness (and kurtosis). 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 especially recommended by Gilchrist and Zakrajšek (2012), we approximate F_t by the credit spread, cs_t.

2.2 Quantile Function of Inflation

The estimated conditional quantiles are approximations to the so-called "quantile function", that is, $\mathcal{Q}_\tau(\bar{\pi}_{t,t+4}|x_t) = F^{-1}(\bar{\pi}_{t,t+4}|x_t)$, where $F^{-1}(\cdot)$ is the conditional inverse cumulative distribution function (CDF) of average future inflation. As noted by Adrian, Boyarchenko, and Giannone (2019), in practice it is challenging to map these estimates into a probability distribution function (PDF) because of approximation error and estimation noise. We therefore follow their approach by 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:

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

[14] See also Chevalier and Scharfstein (1996) for an earlier analysis of this mechanism.

where $z_{t,t+4} = \frac{\bar{\pi}_{t,t+4}(x_t)-\mu_t}{\sigma_t}$ 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. As in Adrian, Boyarchenko, and Giannone (2019), we compute these parameters at each point in time t to minimize the squared distance between our estimated quantile function $\widehat{\mathcal{Q}}_\tau(\bar{\pi}_{t,t+4}|x_t)$, obtained from the quantile Phillips-curve model (2), and the quantile function of the skewed $t-$distribution $F^{-1}(\bar{\pi}_{t,t+4}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t)$ to match the 5$^{\text{th}}$, 25$^{\text{th}}$, 75$^{\text{th}}$ and 95$^{\text{th}}$ quantiles.[15]

2.3 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., deflation).[16] 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-τ) 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, $\mathcal{P}_t^D(\bar{\pi}_{t,t+4}|x_t) \equiv Prob\,(\bar{\pi}_{t,t+4} < \bar{\pi}^*|x_t)$[17] , is the probability mass below $\bar{\pi}^*$ in the conditional density $f(\bar{\pi}_{t,t+k}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t)$:

$$\mathcal{P}_t^D(\bar{\pi}_{t,t+4}|x_t) \equiv \int_{-\infty}^{\bar{\pi}^*} f(\bar{\pi}_{t,t+k}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t) d\bar{\pi}_{t,t+k}, \tag{5}$$

where at (100-τ) percent confidence, inflation will not be, on average, below the level $\bar{\pi}^*$ over the next $t + k$ periods.[18] 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, equivalently, through the CDF.

[15]The parameters are thus functions of the conditioning variables x_t. This dependence is not made explicit for notational convenience.

[16]Our approach differs from the *Value-at-Risk* literature in two ways. First, 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, $\mathcal{Q}_\tau(\bar{\pi}_{t,t+4}|x_t)$, implicitly defined by the integral over the conditional inflation density $f(\bar{\pi}_{t,t+k}|x_t)$ that sums up to τ:

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

Second, (*VaR*) analysis and the recent "Growth-at-Risk" literature focus on measuring one-sided "downside risks" to the economic outlook (e.g., Adrian, Boyarchenko, and Giannone, 2019).

[17]For simplicity, our notation suppresses the dependence of $\mathcal{P}_t^D(\bar{\pi}_{t,t+4}|x_t)$ and $\mathcal{P}_t^U(\bar{\pi}_{t,t+4}|x_t)$ on the parameters $\mu_t,\ \sigma_t,\ \eta_t,\ \kappa_t$ which, in turn, depend on x_t.

[18]Similarly, we can define the conditional upside inflation-at-risk $\mathcal{P}_t^U \equiv Prob\,(\bar{\pi}_{t,t+4} > \bar{\pi}^*|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t)$ as $\int_{\bar{\pi}^*}^{\infty} f(\bar{\pi}_{t,t+k}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t) d\bar{\pi}_{t,t+k} = \mathcal{P}_t^U(\bar{\pi}_{t,t+4}|x_t)$, which is the probability that future inflation will be, on average, above the level $\bar{\pi}^*$ over the next $t + k$ periods.

Figure 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 4^{th} quantile of the inflation distribution. Similarly, the right panel illustrates that, with a 15 percent probability, 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 τth-percent *IaR* is akin to estimating the τ^{th} quantile of the probability distribution of inflation (or its outlook).

Figure 1: Inflation-at-Risk.

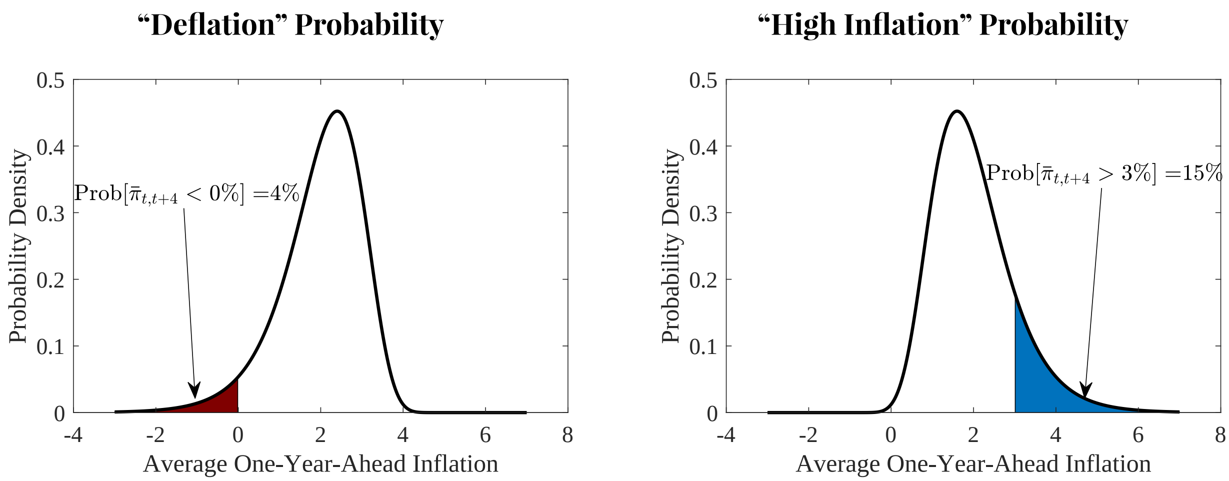

NOTE: The figure displays simulated distributions. In the left panel, the probability of the average future inflation rate falling below 0% is 4 percent. In the right panel the probability of average future inflation exceeding 3% is 15 percent.

3 Inflation-at-Risk Across Time

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 complement the analysis by reporting the historical decomposition of the lower and upper inflation tails. 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. We also present in Appendix C results for core PCE inflation and for Stock and Watson (2019) "Cyclically Sensitive Inflation", which are qualitatively the same. 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 2 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).[19] 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 10$^{\text{th}}$, 50$^{\text{th}}$ and 90$^{\text{th}}$ 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 is indicative of the presence of non-linearities in inflation dynamics.[20]

The top-left panel of Figure 2 presents the quantile-specific Phillips curve coefficients associated with variations in the unemployment gap. The results are in line with the recent evidence 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 2 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.[21]

The second row of panels in Figure 2 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.

[19] The quantile slopes and OLS estimates as well as their confidence intervals can be found in Figure B-1.

[20] In section 5, we complement the information in these figures by showing the confidence bands of the estimated slopes constructed by "blocks-of-blocks" bootstrapping. See also Appendix B for details.

[21] Results are similar using relative oil price inflation $(\pi_t^O - \pi_t)$ (see Appendix D.1), share-weighted core import prices and the real exchange rate (results available upon request).

Figure 2: Quantile Regression Slopes.

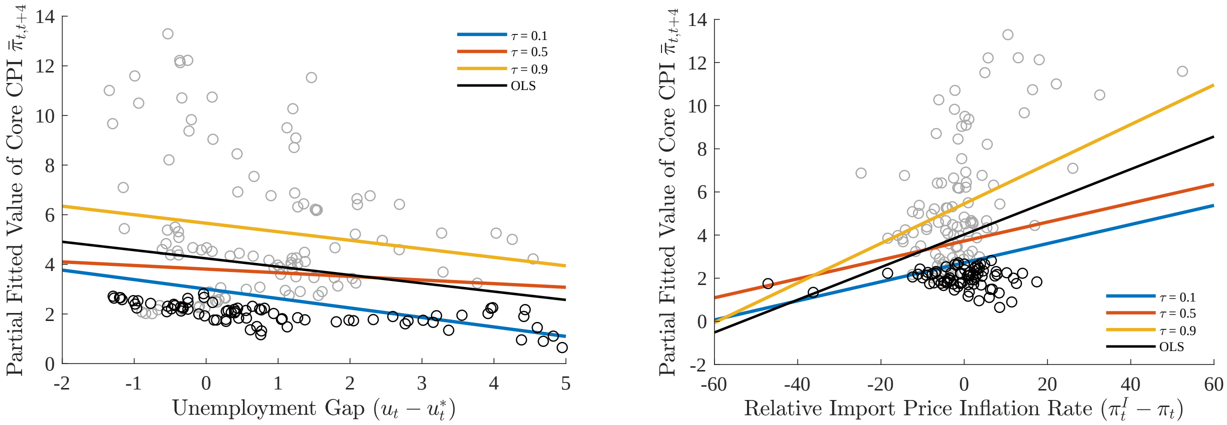

$$\hat{\theta}_\tau = \{\hat{\theta}_{0.1} = -0.38, \hat{\theta}_{0.5} = -0.15, \hat{\theta}_{0.9} = -0.34\} \qquad \hat{\gamma}_\tau = \{\hat{\gamma}_{0.1} = 0.04, \hat{\gamma}_{0.5} = 0.04, \hat{\gamma}_{0.9} = 0.09\}$$

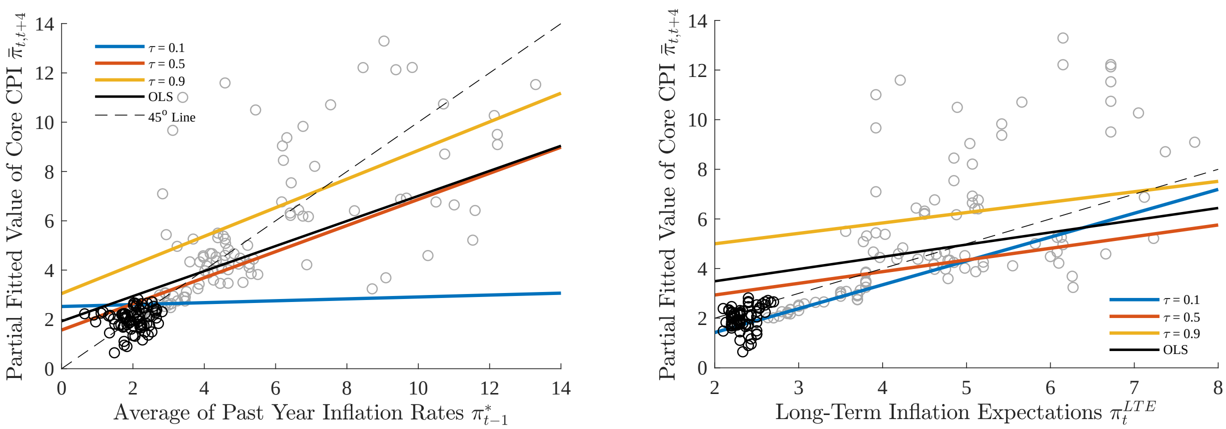

$\hat{\lambda}_\tau = \{\hat{\lambda}_{0.1} = 0.96, \hat{\lambda}_{0.5} = 0.47, \hat{\lambda}_{0.9} = 0.42\}$, where λ_τ is coefficient on π_t^{LTE} and $(1-\lambda_\tau)$ on π_{t-1}^*

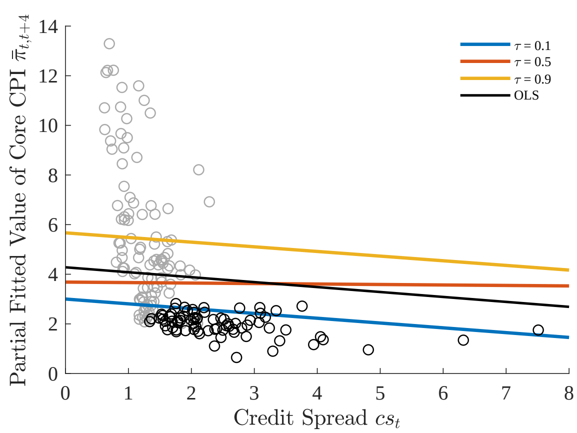

$$\hat{\delta}_\tau = \{\hat{\delta}_{0.1} = -0.19, \hat{\delta}_{0.5} = -0.02, \hat{\delta}_{0.9} = -0.19\}$$

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 10$^{\text{th}}$ (blue) and the 90$^{\text{th}}$ (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 4 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 2 shows the effects of changes in credit spreads on the inflation quantiles.[22] 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 will show in Section 5, 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.

3.1 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 10$^{\text{th}}$ 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 3 displays the evolution over time of the 10$^{\text{th}}$ inflation quantile in the baseline model – which includes the effects of credit spreads (straight blue line) – and the 10$^{\text{th}}$ 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

[22] In Appendix D.2 we show that these results are robust to the choice of other financial variables.

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.

Figure 3: Time Evolution of 10th Inflation Quantile Across Models.

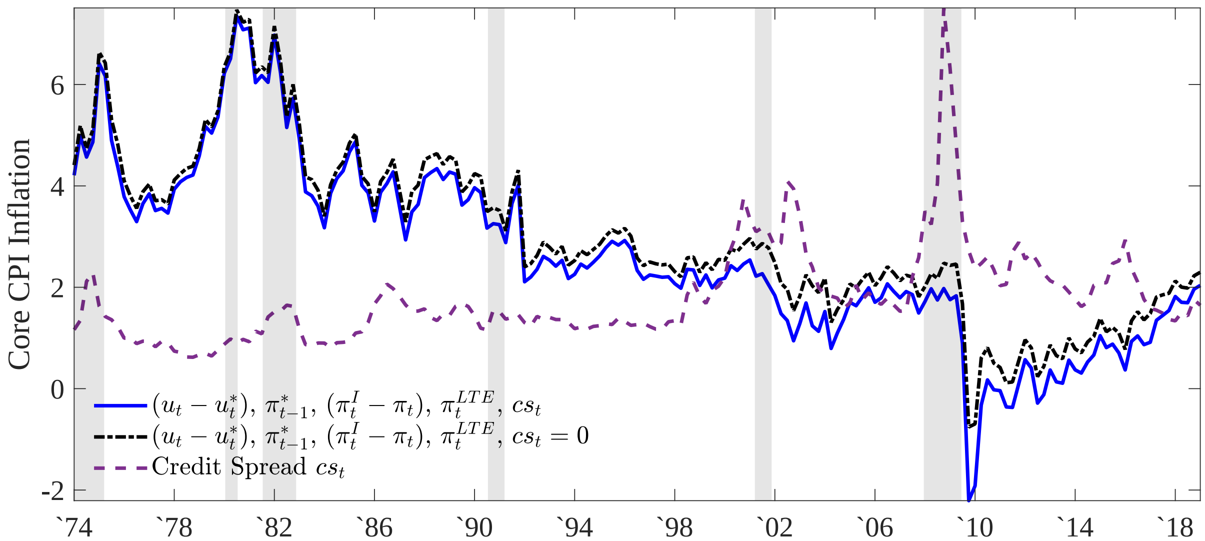

NOTE: The figure displays the time evolution of the 10th inflation quantile estimated from the quantile regressions model (2), in its baseline version (blue straight) and in its version where the effect of credit spreads is set to zero (black dash-dotted). The credit spread (purple dashed) is also reported.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

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 distribution 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.

Predictive Ability 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.[23] In this sense, both the parametric model and the estimation technique employed are being evaluated.

[23] See Rossi (2014) for an excellent summary of density forecast evaluations.

To run the test, we first define the probability integral transform (PIT), i.e., the conditional quantile z_t that corresponds to the realized observation $\bar{\pi}^*_{t,t+4}$:

$$z_t \equiv F^{-1}\left(\bar{\pi}^*_{t,t+4}|x_t\right) = Prob\left(\bar{\pi}_{t,t+4} < \bar{\pi}^*_{t,t+4}|x_t\right), \tag{6}$$

where $F^{-1}\left(\bar{\pi}^*_{t,t+4}|x_t\right)$ 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.

Figure 4: Rossi and Sekhposyan (2019) Test for Correct Calibration of Predictive Density.

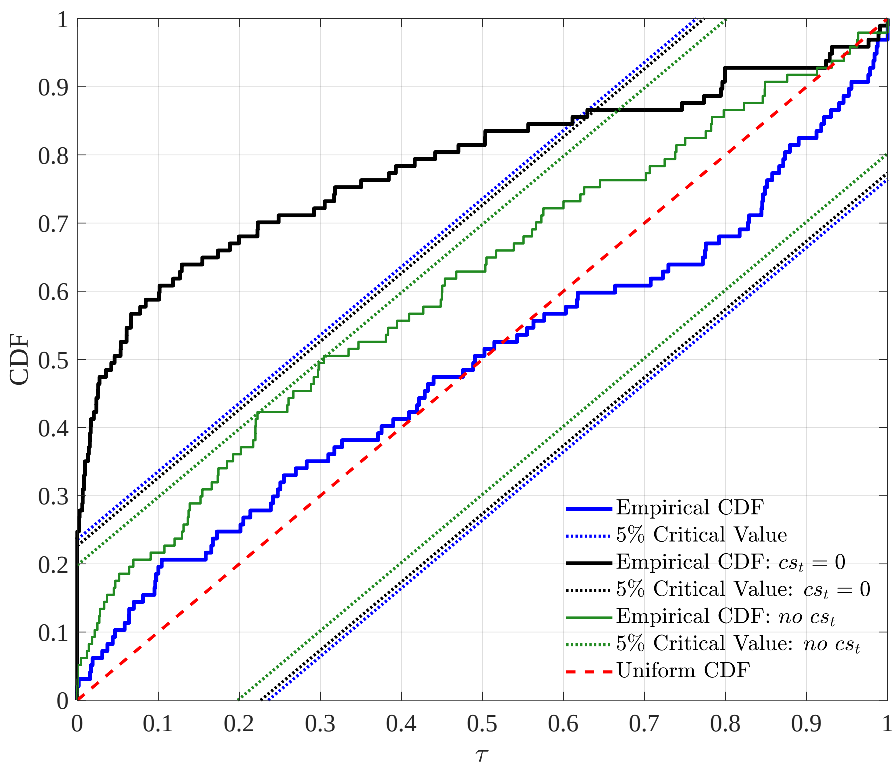

NOTE: The figure illustrates the CDF of a uniform distribution along with the empirical CDFs of out-of-sample PITs obtained from the quantile regressions model (2), in its baseline version (blue) and in two versions which either do not condition on financial variables in estimation (green) or in the 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.

In Figure 4 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). These 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.[24]

Unlike the baseline model, the model that disregards the role of financial variables in the construction of the quantiles does not pass the test for correct calibration – as it poorly specifies the predictive inflation distribution by placing too little mass on its lower tail. The model neglecting financial variables in estimation performs worse than the baseline along the entire inflation distribution except on the upper tail.

3.2 The Predictive Distribution of Inflation

Figure 5 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.[25] 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 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 1975:Q1, right after the recession triggered by the first wave of oil shocks and the easing cycle of the Federal Reserve. The second quarter of 1981 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 1985: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 the last quarter of 1999.

Overall, the estimated quantile models are able to 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

[24]Nevertheless, if we use a rolling window scheme which uses 20 years of data we can confirm that the model which does not consider the credit spread 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).

[25]As formally described in Section 2.3, we construct the skewed t-Student probability density function of inflation using the quantiles estimated using the regression model (2).

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 which we analyze in Section 4.

The bottom panel of Figure 5 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 2008:Q4, with the one shown for the last quarter of 1999 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 5 to develop this issue in depth, the three dates chosen in the bottom panel of Figure 5 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 2009: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 2015 (yellow distribution shown in the bottom panel of Figure 5). These and other considerations are more deeply explored in the next section that delves into the role played by labor market dynamics, and especially credit spreads and inflation expectations across the two subsamples just considered.

Figure 5: Conditional Predictive Inflation Densities at Selected Time Episodes.

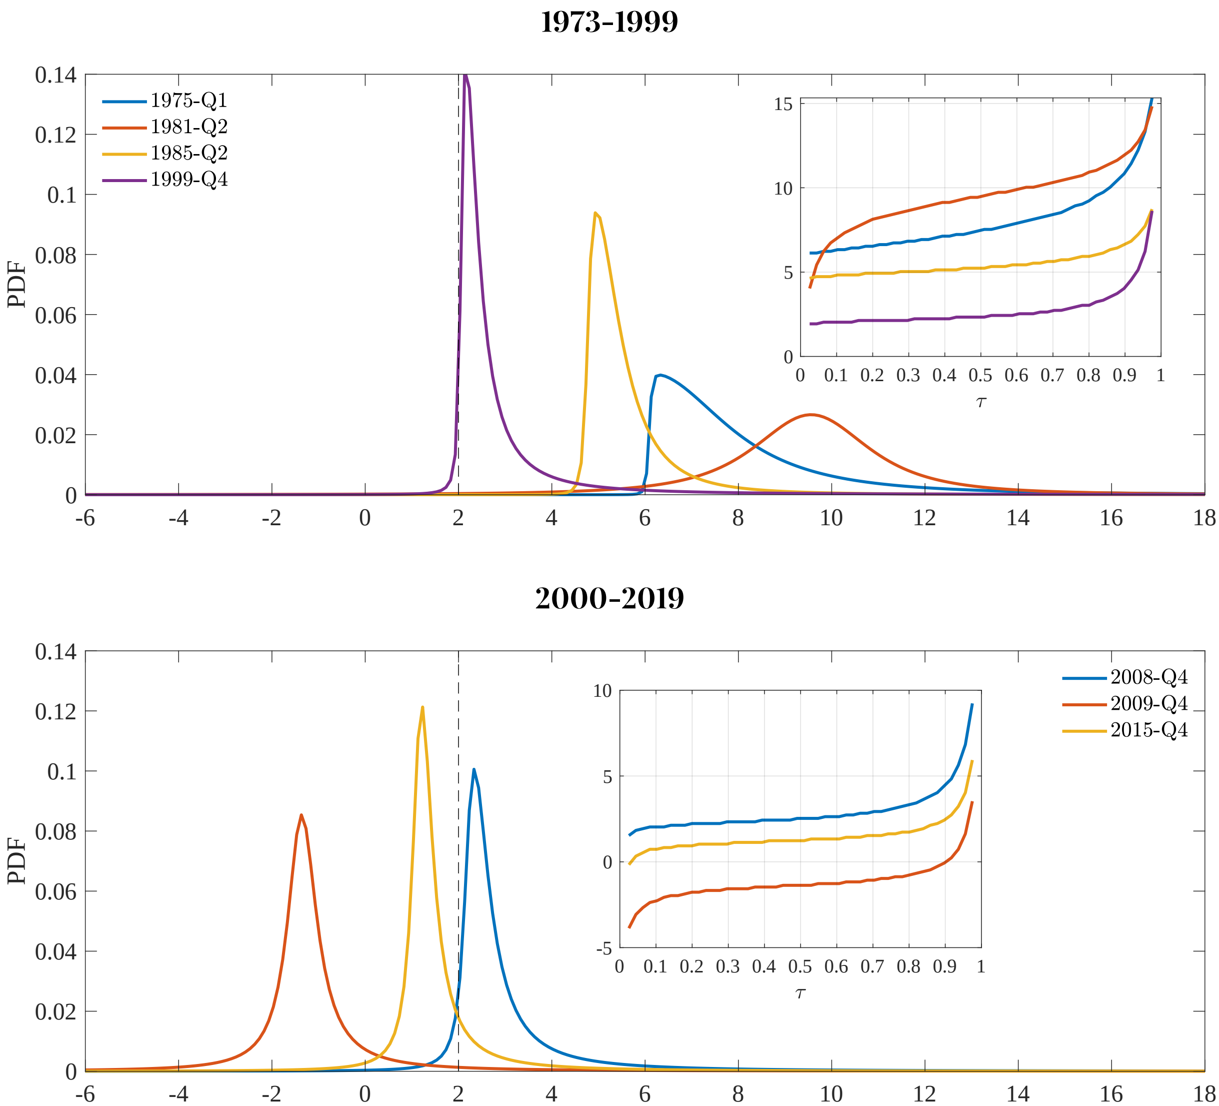

NOTE: The figures shows for selected time episodes the estimated skewed-t conditional densities of average four-quarter-ahead core CPI inflation associated with the quantile regressions model (2). The inset box reports the values of average future inflation across quantiles at the same selected periods.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

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

The full sample results anticipated that two distinct subsamples emerge when characterizing the determinants of the inflation distribution. The first period, running from 1973 to 1999, 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 were ubiquitous themes in the description of inflation dynamics.

The first period contrasts with the second subsample, from 2000 to 2019, 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, 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. 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, which point to a greater role for the labor market recovery and well-anchored inflation expectations in supporting average inflation. This point will be further investigated in the next section using contrasting evidence from the United States and the euro area.

4.1 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 6.

Three results stand out. First 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.[26] Second, long-run inflation expectations exert a symmetric effect on the inflation distribution. In fact, in this context well-anchored long-run inflation expectations lower the response of average inflation to labor market slack, financial conditions and relative price changes. However, it would be misleading to dismiss the role of these factors focusing on the conditional mean only. Instead, financial conditions and, to a lesser extent, labor market outcomes are key drivers of downside inflation risk, which in turn are important to characterize the entire

[26]We obtain long-term inflation expectations from Consensus Economics. Results are similar if we use long-term inflation expectations from the SPF or Michigan survey, which are respectively available from 1987Q1 and 1981Q1.

inflation distribution and its dynamics. Finally, relative import prices still pose threats to the upper inflation quantile, though to a lesser degree than prior to the Great Moderation.[27]

Figure 6: Quantile Regression Slopes Across Subsamples.

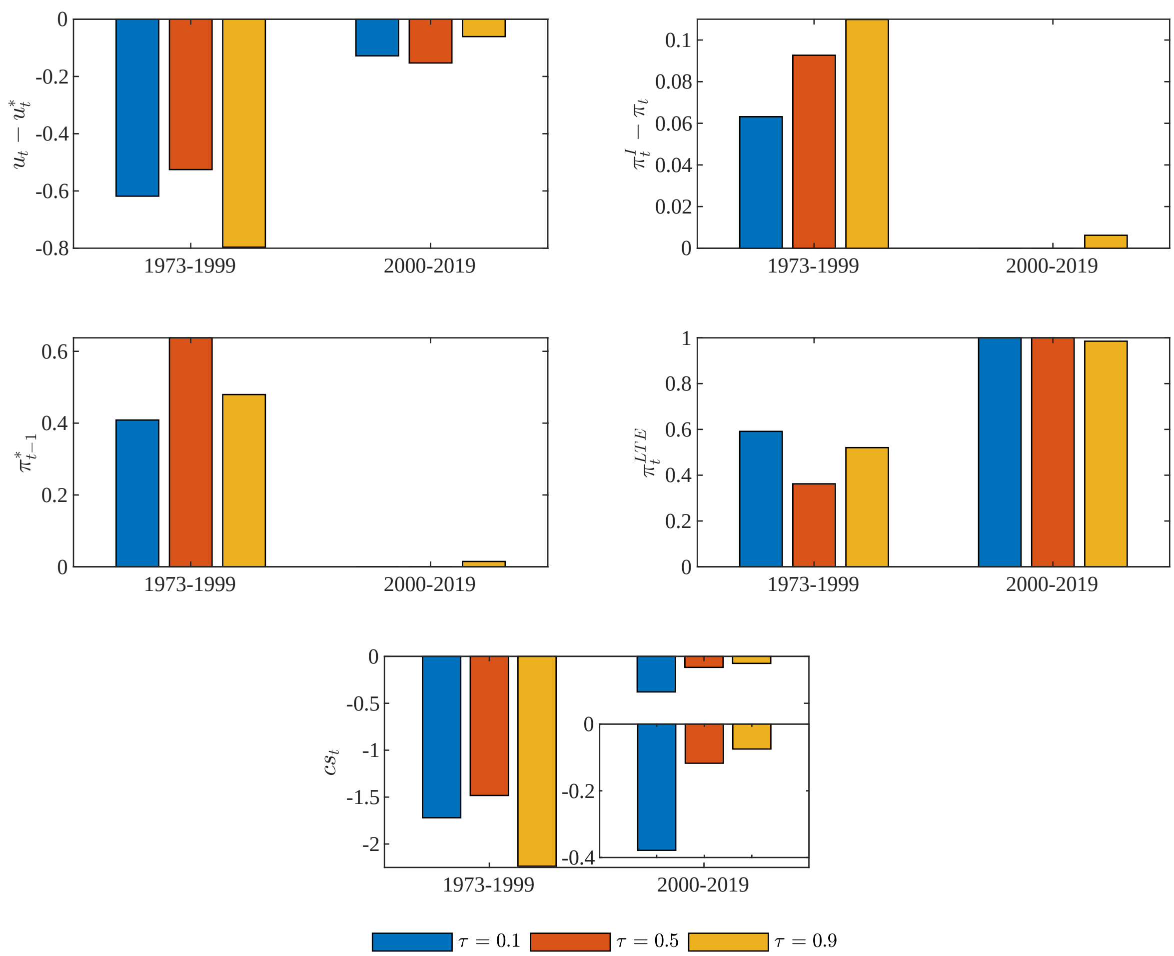

NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead core CPI inflation defined in expression (2). Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate the coefficients associated with the 10^{th} quantile (blue),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

[27]In Appendix C we show that these results are robust to alternative measures of inflation, either core PCE or CSI. Using current method core CPI or core PCE also delivers similar results (results are available upon request).

Inflation Probabilities We now 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 7 we present the conditional predictive densities (left panel) and inflation probabilities (right panel) 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 left-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 straight 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 left-skewed distributions. Conversely, the post-2000 economy wouldn't have suggested any change in the inflation odds. This tension is related to debate the on "missing inflation", which is 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.

Figure 7: "Counterfactual" Predictive Densities (left) and Inflation Probabilities (right).

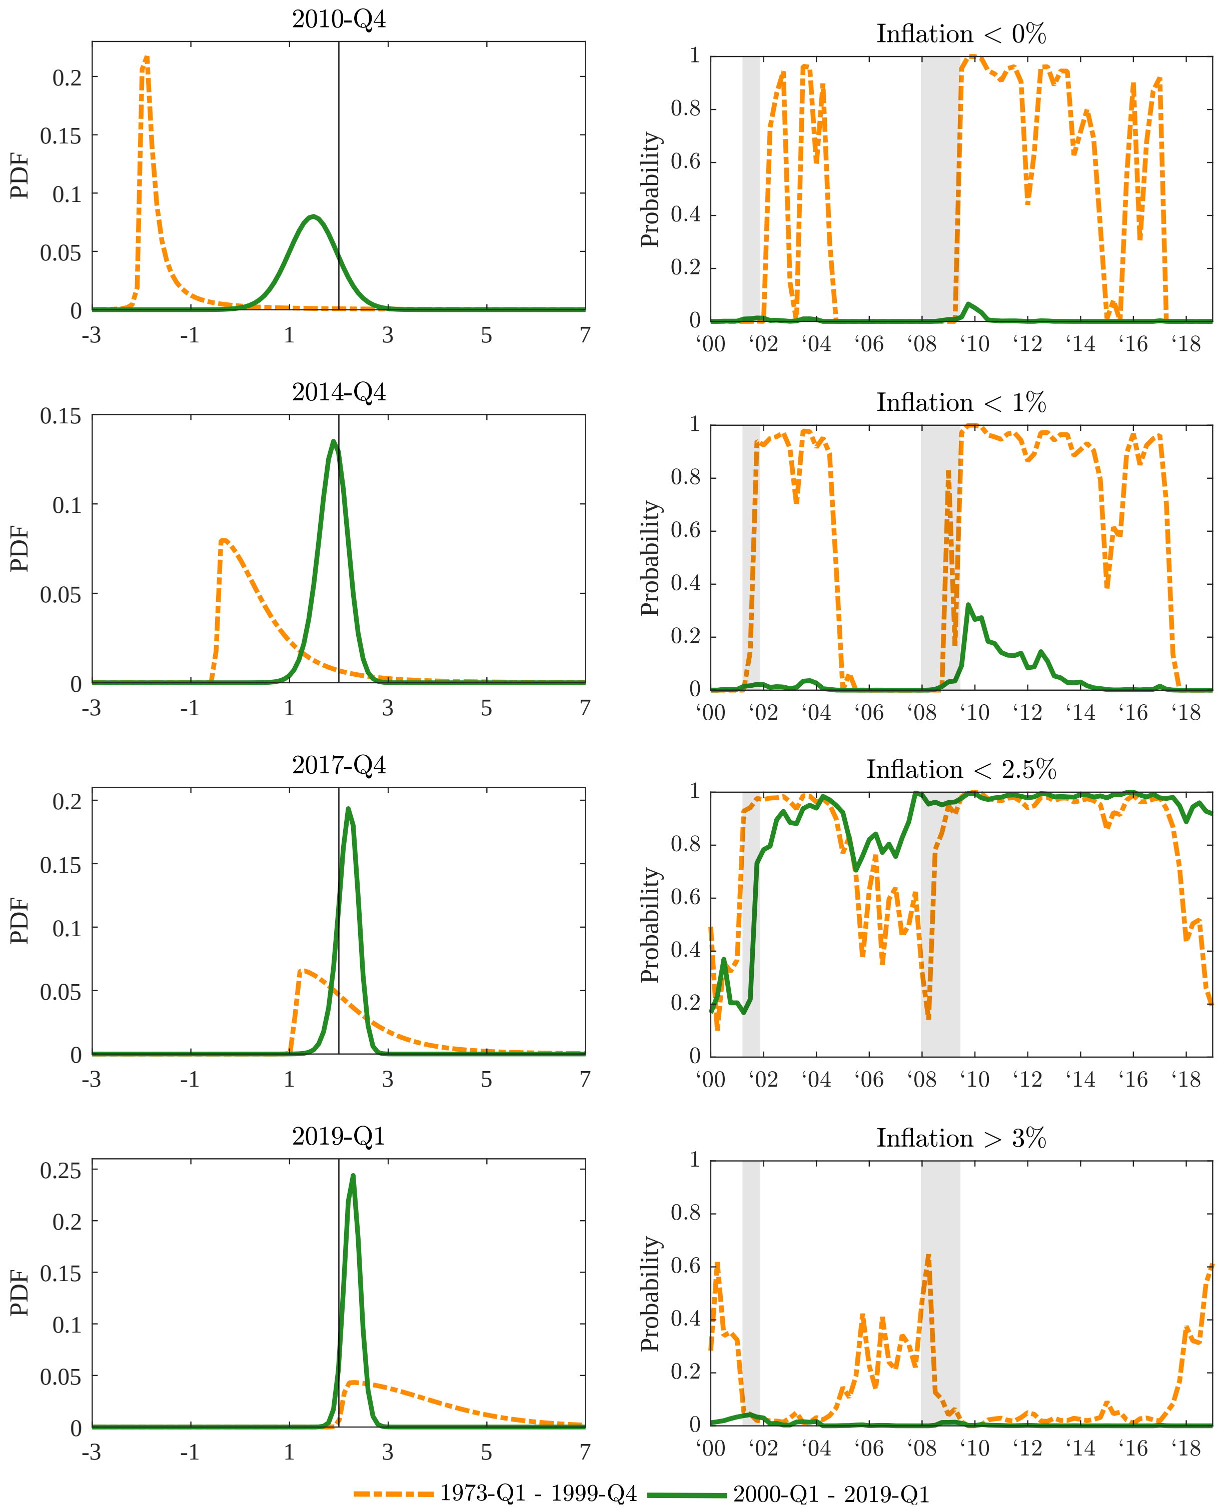

NOTE: The left panels show "counterfactual" skewed t-Student conditional densities of average four-quarter-ahead core CPI inflation computed using "counterfactual" conditional quantiles over 2000-2019 which were obtained using different subsample estimates of the quantile regressions model (2). The right 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.

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 deflation risks by supporting the job market and easing credit conditions. We discuss the relative role of these risk factors for the inflation outlook next.

5 Euro Area vs. United States During the Great Recession

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.[28] 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.

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 quantile model. The black squares report the point estimates of the 10$^{\text{th}}$, 50$^{\text{th}}$, and 90$^{\text{th}}$ 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 B). 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.[29]

[28] Busetti, Caivano, and Rodano (2015) use dynamic quantile regressions to estimate and forecast the conditional distribution of euro-area inflation. Tagliabracci (2018) estimates this distribution conditioning on Eurocoin.

[29] Busetti, Caivano, Monache, and Pacella (2019) investigate the role of domestic and global determinants of euro area core inflation by estimating a Phillips curve using an expectile regression approach.

Figure 8: Quantile Regression Slopes and Confidence Intervals.

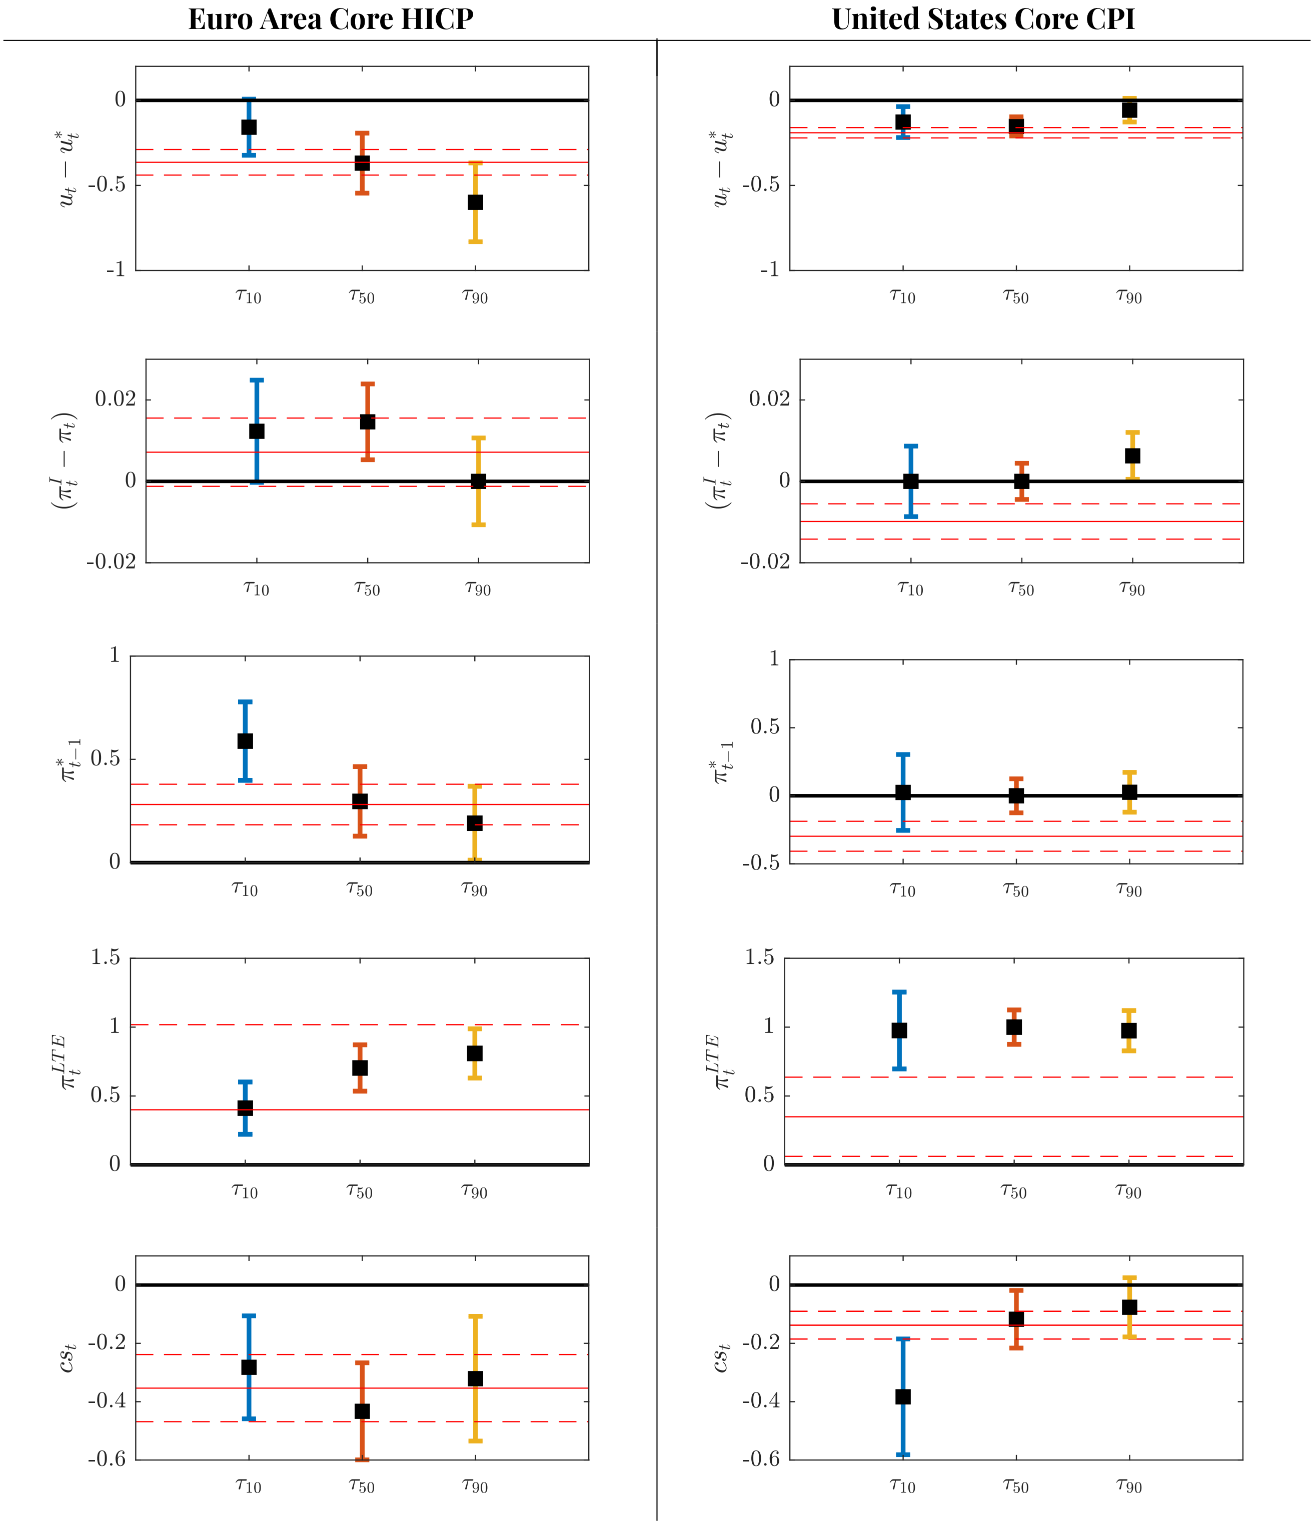

NOTE: The figure displays the slope coefficients of the quantile regression of average four-quarter-ahead euro area core HICP (left) and United States Core CPI inflation (right) defined in (2). The black squares correspond to the point estimates whereas the vertical lines to the 68% confidence intervals computed via "blocks-of-blocks" bootstrap using 10,000 replications (see Appendix B) 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 represented by the red lines.

Longer-term inflation expectations influence strongly the overall inflation distribution in the U.S., and to a lesser extent in the eurozone. It is striking that, in the U.S., the estimated slope coefficient is not statistically different from unity across all quantiles. Accordingly, inflation inertia – the backward-looking inflation component – plays a negligible role in characterizing inflation in the United States. In contrast, in the euro area the effects of long-term inflation expectations on the inflation distribution are more asymmetric: Inflation expectations play a major role in explaining the upper tail of inflation, while odds of low inflation are less sensitive to changes in long-run inflation expectations as they are also driven by past inflation. In other words, in the eurozone, unmoored reductions in inflation expectations result in more persistent increases in downside inflation risks as their negative effect is propagated over time through a higher inflation inertia than in the United States.

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.[30] Most importantly, in the U.S. financial conditions are the only significant source of asymmetry in the inflation distribution.

Inflation Quantiles Figure 9 highlights key aspects of the evolution of the inflation outlook by displaying the time series of the median, the 10$^{\text{th}}$ and the 90$^{\text{th}}$ inflation quantiles. The top panel shows the evolution for the euro area while the bottom panel focuses on the U.S.

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.

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 10$^{\text{th}}$ 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.

[30] This result is also robust to limiting the sample to 1995:Q1-2007:Q4 so as to limit the Great Recession period and to using the short-term instead of the long-term CBO NAIRU measure for u_t^* which features a higher level in the aftermath of the Great Recession. Further, in Appendix C 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.

Figure 9: Time Evolution of Selected Conditional Inflation Quantiles.

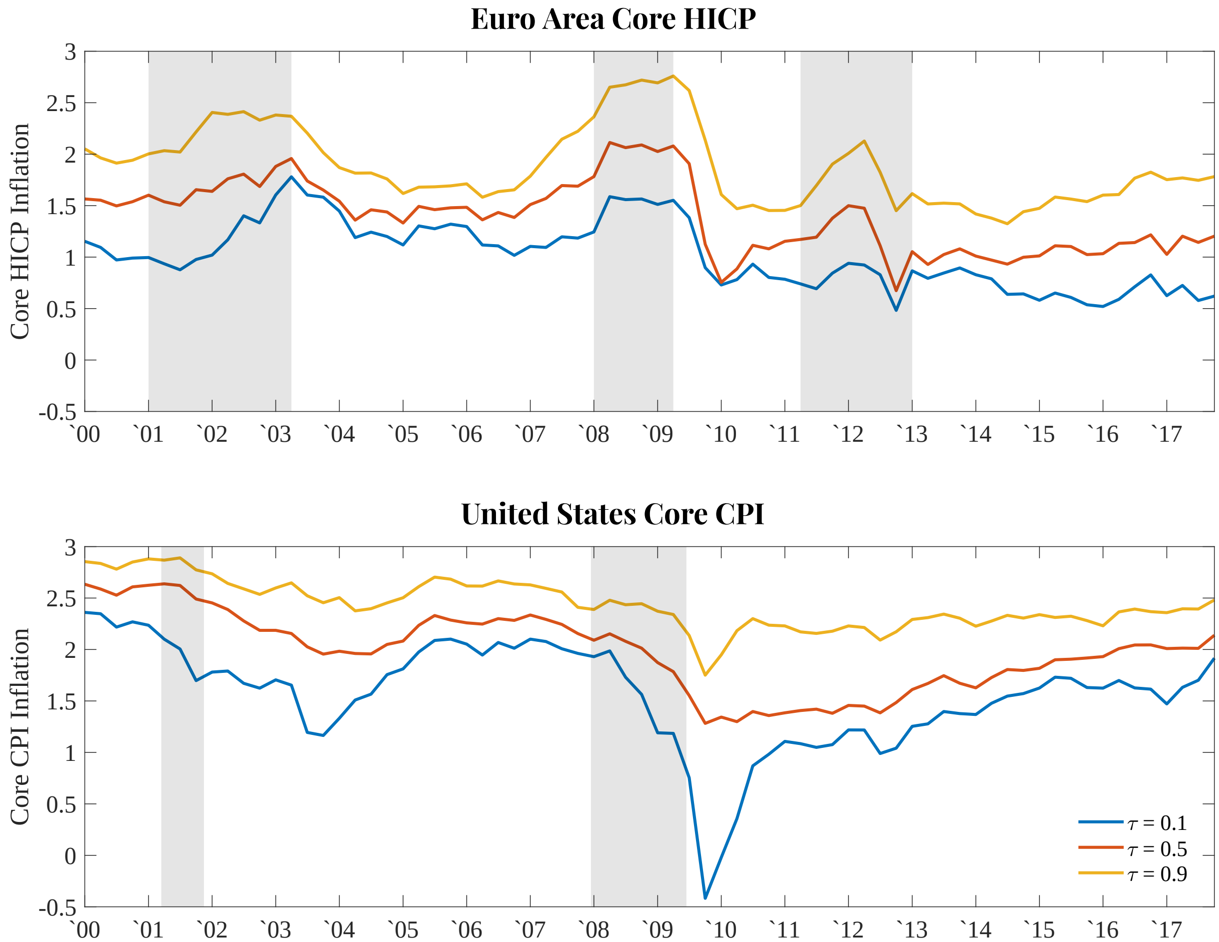

NOTE: The figure displays the time evolution of the conditional quantiles of euro area core HICP inflation (top) and of Unites States core CPI inflation (bottom) estimated from the quantile regression defined in (2). Shaded bars indicate NBER-dated recessions for the United States and OECD-based recession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

This opens up the question of 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 10 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 10) 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).

Figure 10: Historical Contributions of Economic and Financial Factors.

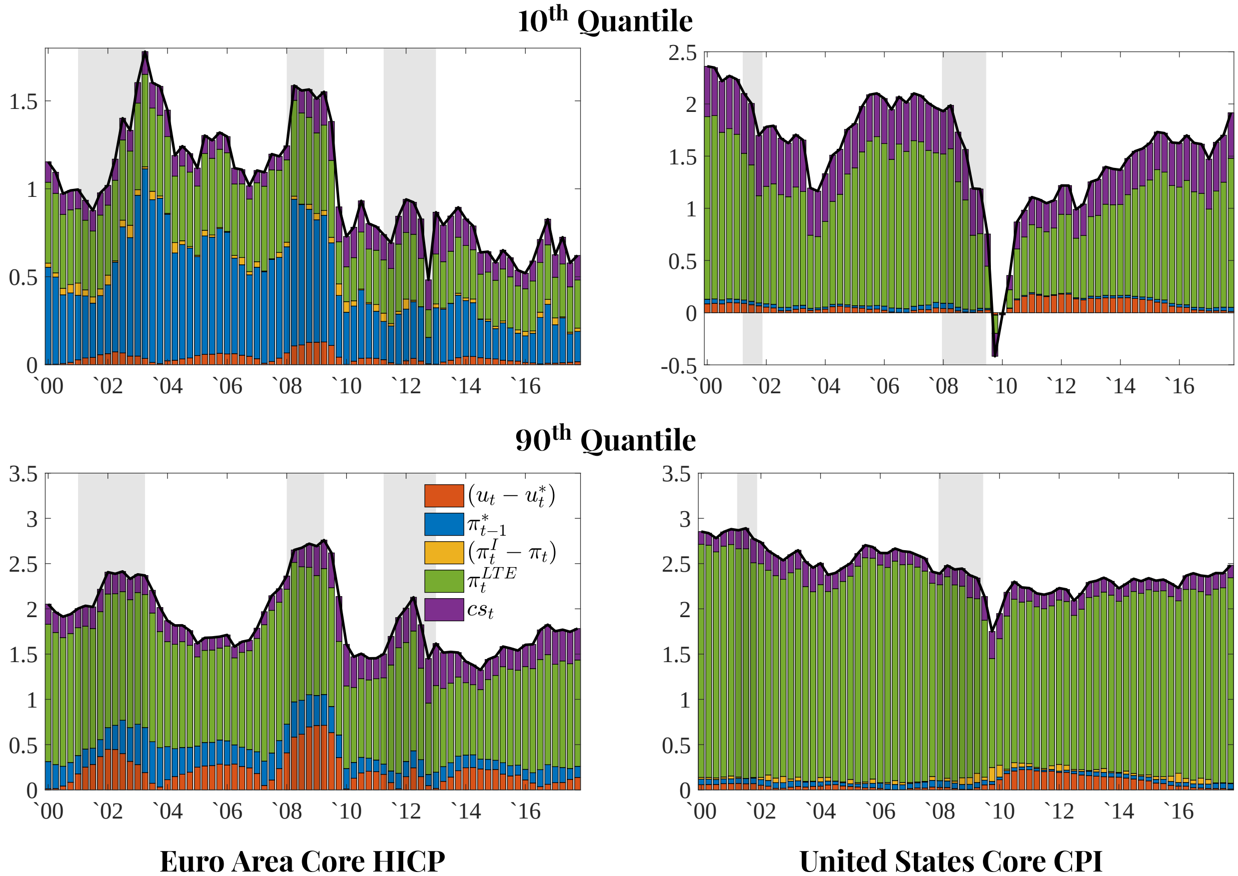

NOTE: The figure shows historical contributions of average four-quarter-ahead euro area core HICP (left) and United States core CPI inflation (right) associated with the quantile regressions model (2). The contribution of a given inflation determinant is obtained by multiplying its time series with its 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.

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 2017.

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.

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 11, we illustrate the time evolution of the 10th conditional inflation quantile of euro area core HICP inflation (left) and of Unites States core CPI inflation (right), both for the baseline version of the model (blue straight) 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 11) pushed down the lower inflation quantile, which slowly moved back to about 2 percent by the end of 2017.

Figure 11: Partial Effect of Credit Spread on 10th Inflation Quantile.

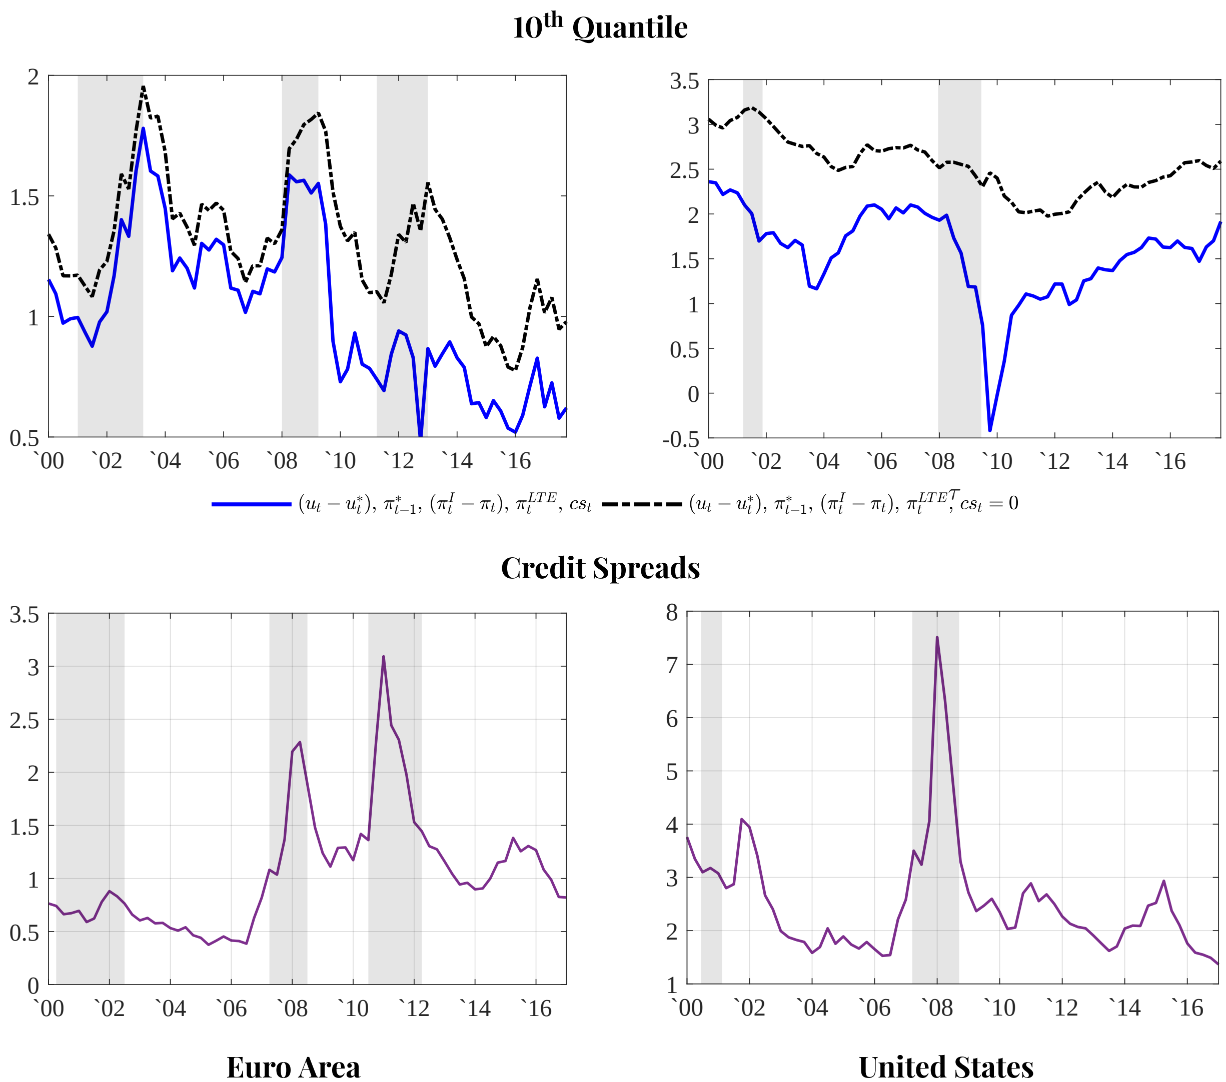

NOTE: The figure displays the time evolution of the 10th conditional inflation quantile of euro area core HICP inflation (left) and of Unites States core CPI inflation (right) estimated from the quantile regressions model (2), in its baseline version (blue straight) and in its version in which the effect of credit spreads is set to zero (black dash-dotted). Shaded bars indicate NBER-dated recessions for the United States 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 11. 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.[31]

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 12, which plots the estimated euro area core HICP predictive inflation distributions (left column) and their associated quantile functions (right column) across four selected dates. We start at the dawn of the global financial crisis (2008:Q4), then explore those periods in which downside risks were most acute (2009:Q4 and 2012:Q4) and finally zoom in the end of our sample (2017:Q4). The blue straight 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 late 2009 and remarkably so in late 2012 – at the time of Draghi's speech. It is noteworthy how by the end of 2008, 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 2017 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 12, which shows that the inflation quantiles which condition on credit spreads were significantly below those that solely rely on economic factors.

[31] In Appendix C we show that, for the U.S. economy the results shown in the top-right panel of Figure 11 are robust to alternative measures of inflation, either core PCE or CSI.

Figure 12: Selected Time Episodes of Predictive Densities (Left) and Skewness (Right). Euro Area Core HICP Inflation.

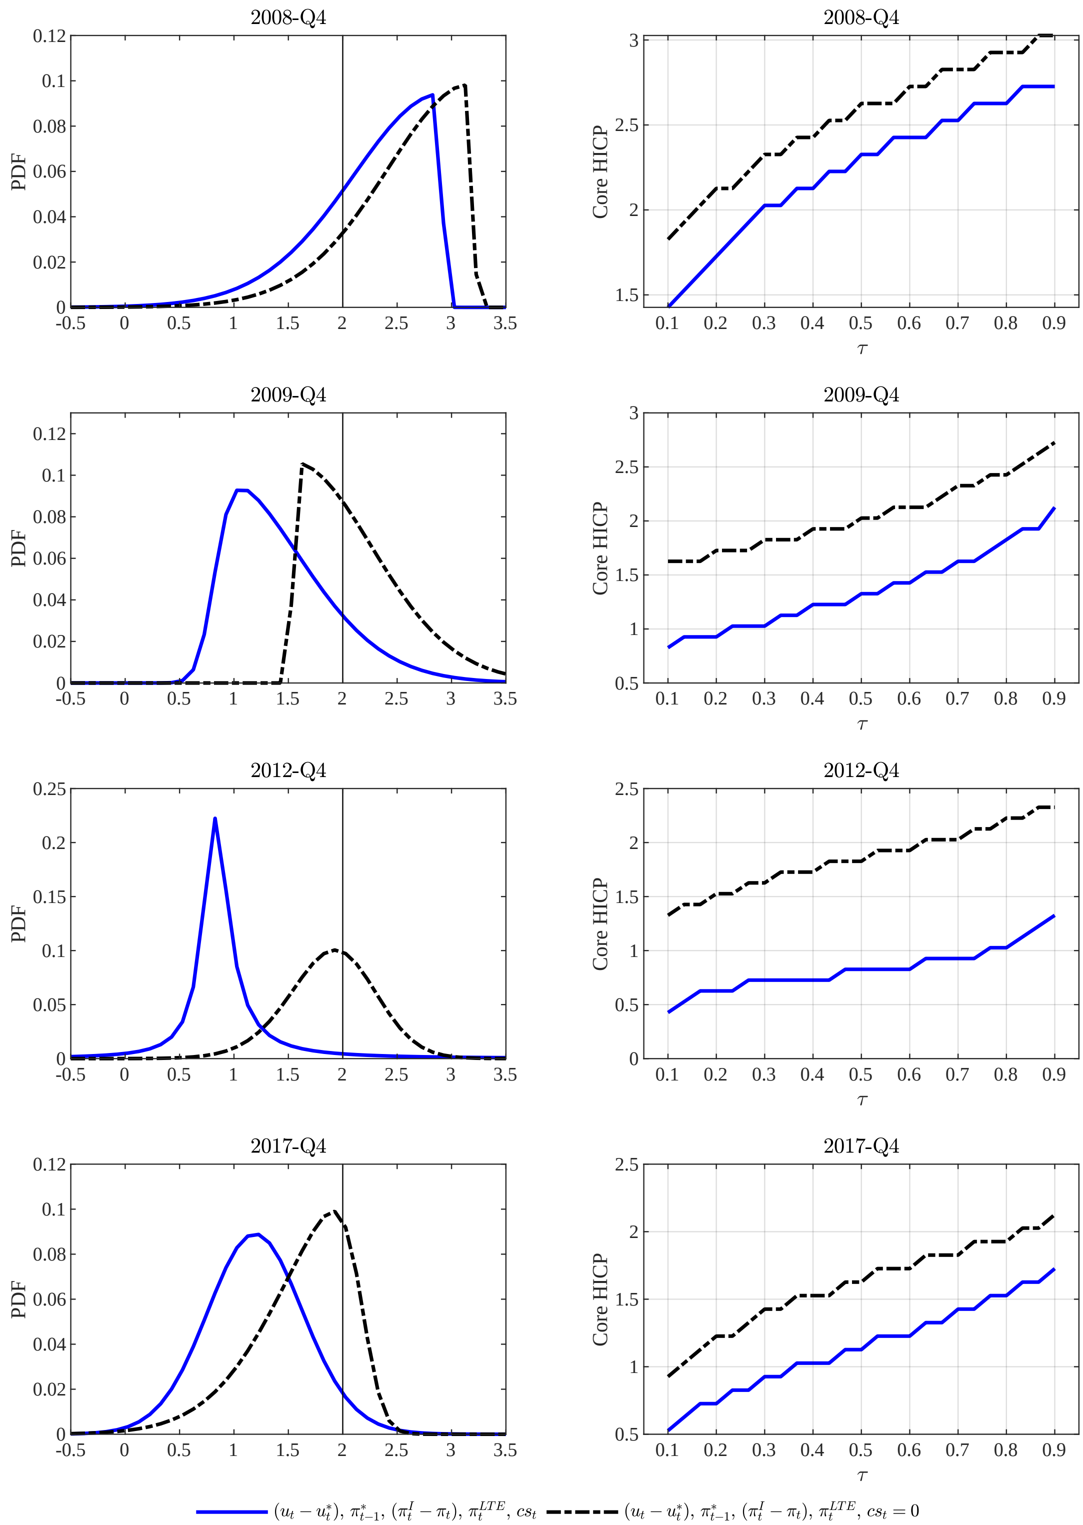

NOTE: The left panels show the estimated skewed t-Student densities of average four-quarter-ahead euro area core HICP inflation for alternative specifications of the quantile regressions model (2), in its baseline version (blue straight) and in its version where the effect of credit spreads is set to zero (black dash-dotted). The right panels show the estimated skewed t-inverse cumulative associated with the conditional densities in the left panels.

Figure 13: Selected Time Episodes of Predictive Densities (Left) and Skewness (Right). United States Core CPI Inflation.

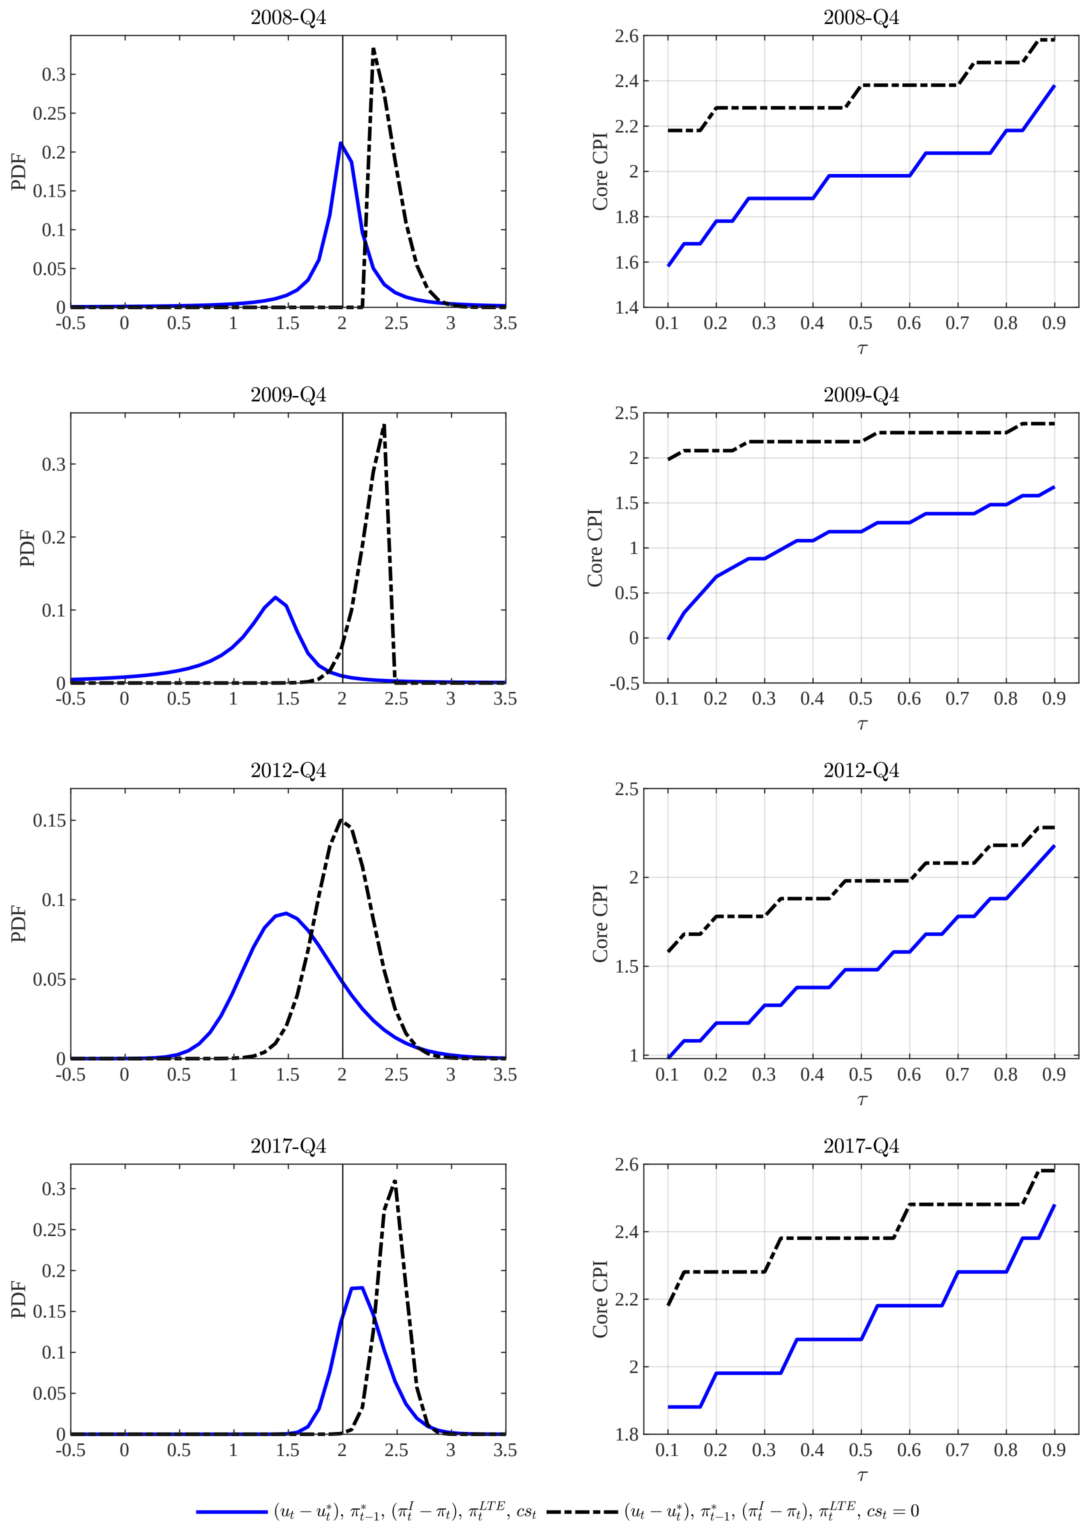

NOTE: The left panels show the estimated skewed t-Student densities of average four-quarter-ahead United States core CPI inflation for alternative specifications of the quantile regressions model (2), in its baseline version (blue straight) and in its version where the effect of credit spreads is set to zero (black dash-dotted). The right panels show the estimated skewed t-inverse cumulative associated with the conditional densities in the left panels.

In Figure 13, we compare these results with the experience of the United States, for which we consider similar dates. We focus on 2009: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 2009: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).[32]

Looking at the right columns of Figures 12 and 13, 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).

Inflation-at-Risk: The Role of Skewness We first show that the effects of changes in economic and financial conditions on downside inflation-at-risk depend on the skewness of the inflation distribution.

To see this, we need to characterize the derivative of $\mathcal{P}_t^D(\bar{\pi}_{t,t+4}|x_t) \equiv Prob(\bar{\pi}_{t,t+k}|x_t) < \pi^*$ defined in (5), for a given probability cutoff $\bar{\pi}^*$, with respect to an inflation determinant x_t. Formally,

$$\frac{\partial \mathcal{P}_t^D(\bar{\pi}_{t,t+4}|x_t)}{\partial x_t} = \frac{\partial}{\partial x_t} \int_{-\infty}^{\bar{\pi}^*} f(\bar{\pi}_{t,t+k}, \mu_t, \sigma_t, \eta_t, \kappa_t|x_t)\, d\bar{\pi}_{t,t+k}, \tag{7}$$

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,

$$\frac{\partial \mathcal{P}_t^D(\bar{\pi}_{t,t+4}|x_t)}{\partial x_t} = \int_{-\infty}^{\bar{\pi}^*} \frac{\partial f(\bar{\pi}_{t,t+k}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t)}{\partial \bar{\pi}_{t,t+k}(x_t)} \frac{\partial \bar{\pi}_{t,t+k}(x_t)}{\partial x_t}\, d\bar{\pi}_{t,t+k}, \tag{8}$$

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

$$\frac{\partial \mathcal{P}_t^D(\bar{\pi}_{t,t+4}|x_t)}{\partial x_t} = \beta_{OLS} \int_{-\infty}^{\bar{\pi}^*} \frac{\partial f(\bar{\pi}_{t,t+k}|x_t, \mu_t, \sigma_t, \eta_t, \kappa_t)}{\partial \bar{\pi}_{t,t+k}(x_t)}\, d\bar{\pi}_{t,t+k}. \tag{9}$$

From expression (9) 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

[32]In Appendix C we show that, for the U.S. economy the results shown Figure 13 are robust to alternative measures of inflation, either core PCE or CSI.

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. We offer an illustrative example of this effect in Appendix E.

We now show how credit conditions affect the odds of low and very low inflation. Figure 14 displays, beginning in 2000, our estimates of the evolution of the probability of observing inflation rates below 1 or 2 percent over the next four quarters, respectively. 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 straight line) and its version which omits the effects attributable to changes in credit conditions (black dash-dotted line).

Figure 14: Inflation Probabilities for Alternative Cutoff Values.

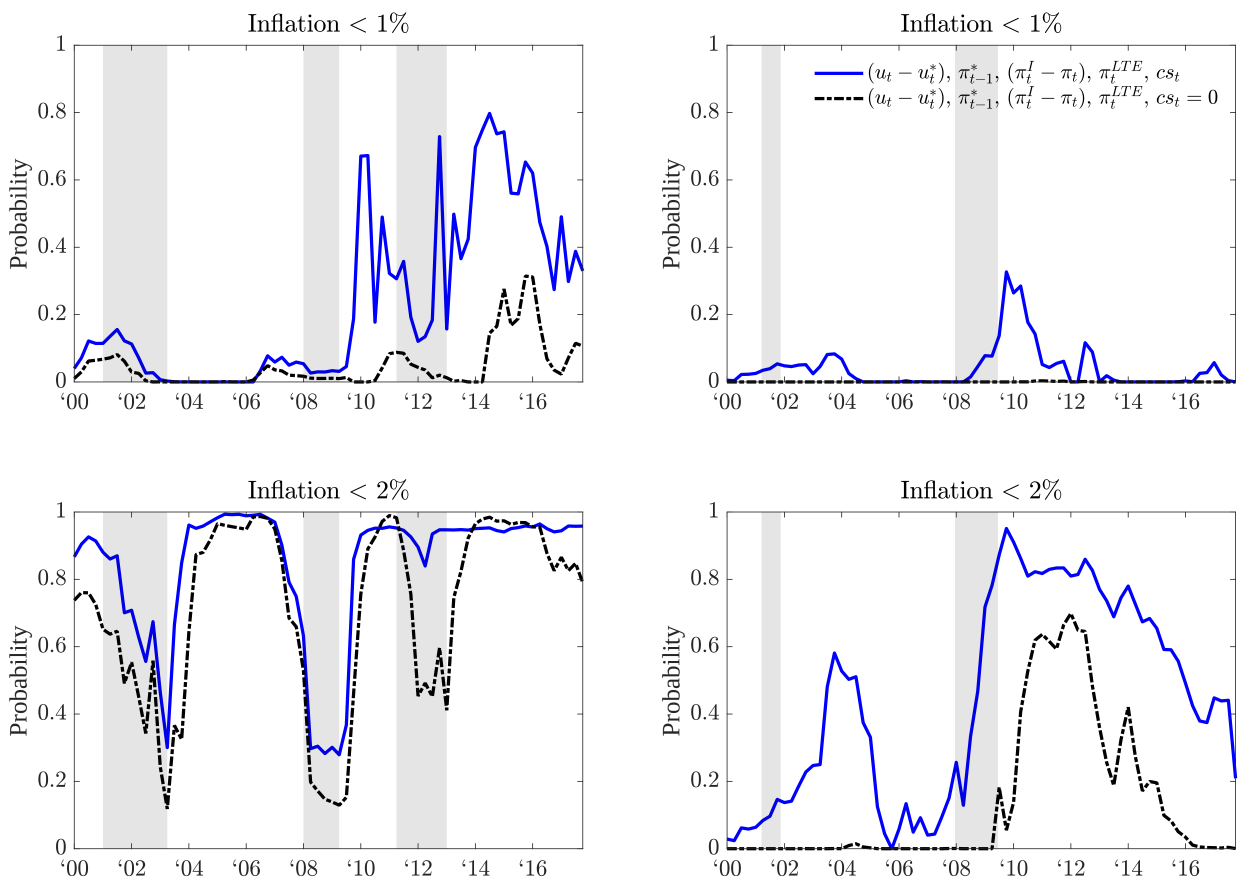

Euro Area Core HICP **United States Core CPI**

NOTE: The figure shows the time evolution of inflation probabilities of euro area core HICP inflation (left) and United States core CPI 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 without and with the credit spread (in blue straight and black dash-dotted lines, respectively). Shaded bars indicate NBER-dated recessions for the United States and OECD-based recession indicators for the euro area.

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 2010 (upper right panel). Results for the eurozone are more striking on this account. Changes in the credit spreads in 2009 and especially in 2012 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 14. 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. The bottom panels of Figure 14 considers the partial effect that credit conditions impinged on the probability that inflation could run below 2 percent. Even the narrower credit spread that prevailed in the U.S. during 2015, translated into a non-negligible probability of inflation running below 2 percent during 2016 and 2017, before subsiding at the end of the sample. As for the euro area, the credit spread accounts for the probability of inflation being lower than 2 percent remaining at sustained levels around the sovereign debt crisis. Indeed, in the absence of this channel, during that period the same probability would have been around 50 percent instead of about 90 percent, on average.

6 External Validation: The Case of the United States

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.[33]

The black-continuous line shown in the top panel of Figure 15 displays the options-implied probability of *headline* CPI inflation next year being below 1% – from mid-2009 until the end of 2017. A straight reading from this market portraits 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 (see Loria, Matthes, and Zhang, 2019b) – especially since the onset of the global financial crisis.

The top panel of Figure 15 corroborates this claim by displaying the inflation probabilities along with 3-months- and 6-months-ahead oil price surprises – computed using the oil market price ex-

[33]The 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.

pectations that Baumeister and Kilian (2016) recovered from oil futures prices and after controlling for changes in the risk premium.[34] 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 coincides with a period in which financial markets have been steadily surprised to the downside in their oil price expectations.

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-1 in Appendix F).[35]

The dashed red line displayed at the bottom panel of Figure 15 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 (previously displayed in the top right panel of Figure 14). The figure is very suggestive as it shows how both measures point 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 2018, 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-1 in Appendix F).

The small remaining differences between the two measures 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).

Financial Market Inflation Probabilities and the Credit Spread Two defining features of our quantile Phillips curve model are that tight financial conditions carry substantial and persistent 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 a regression of the options-implied inflation probabilities on the credit spread. We confirm

[34] The 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.

[35] Since 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.

Figure 15: Inflation Probabilities, Quantile Model vs. Financial Markets.

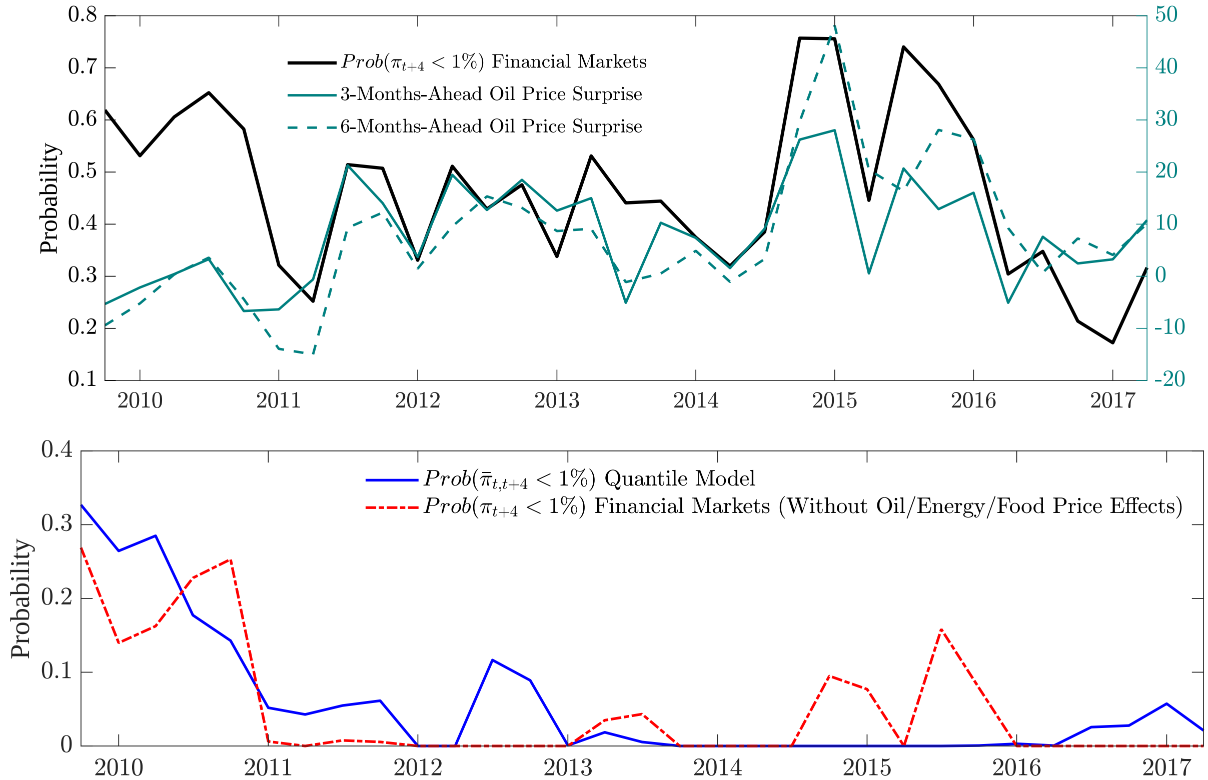

NOTE: The top panel shows the options-implied inflation probabilities of United States headline CPI inflation next year being above 1% along with the 3-months- and 6-months-ahead oil price surprises computed using the oil market price expectations that Baumeister and Kilian (2016) recovered from oil futures prices (top panel). The 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 frequency.

that higher credit spreads are associated with a downward pressure on the inflation distribution and that as the cutoffs for the inflation probabilities increase the relationship weakens. This result is established via standard OLS regressions and is robust to the inclusion of the regressors used to purify the inflation probabilities from the oil price effects in the previous figure.

In the left panel of Figure 16 we present the coefficients of the credit spread for various inflation probabilities derived from financial market, along with their 95 percent confidence intervals. 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. Despite the vast disparities in the construction of the tails of the inflation distribution, the estimated slopes are very similar to each other. Finally, in Appendix F we plot the time series of these financial-market-based probabilities together with the credit spread (see Figure F–2). 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.

Figure 16: Credit Spread Coefficients, Quantile Regression vs. Financial Markets.

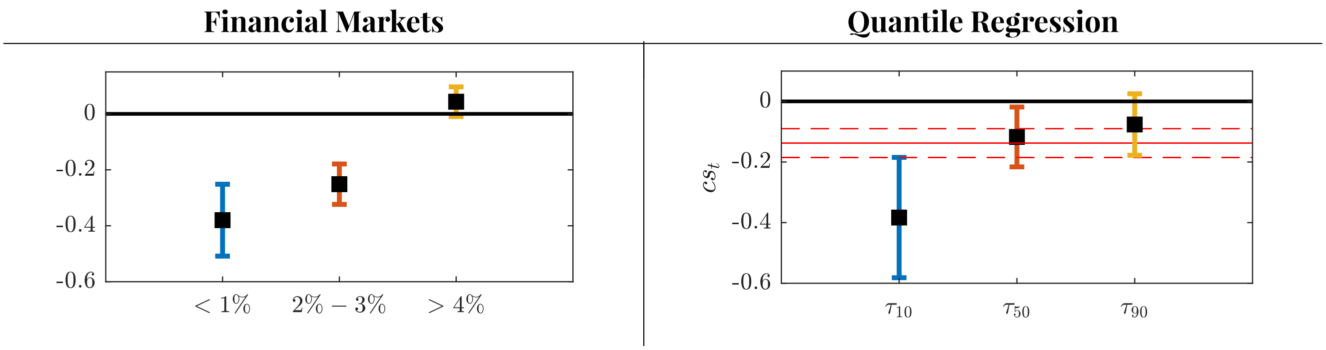

NOTE: The left panel reports the slopes of separate regressions of inflation probabilities on the credit spread (at monthly frequency), along with their 95% confidence interval. The coefficient for the probability of future inflation being below 1% is rescaled to match the slope estimated on the lowest 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.

7 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. Including data from the 1970s shows ample variability in the entire conditional distribution of inflation, a result which we confirm when restricting our analysis to the post-2000 stable and low *mean* inflation period. In particular, using time-series data for the United States and the euro area, we document that – after controlling for the state of the labor market and inflation expectations – tight financial conditions generated times of substantial and persistent downside risks to inflation, an analysis forsaken by much of the literature. Our paper thus offers a new empirical perspective to existing macroeconomic models and to policymakers, showing that changes in credit conditions are also key to understand the tail-risk dynamics of inflation.

In future research, we plan to construct a "risk-adjusted" inflation measure which can be derived from a central bank's preferences not only about deviations of inflation from its target but also about how tolerable (downside and upside) inflation risks are. Finally, we will further study inflation-at-risk by exploiting time variation in sectoral inflation rates and ask whether particular sectors make inflation particularly vulnerable.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109(4), 1263–89.

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.

Baumeister, C., and L. Kilian (2016): "A General Approach to Recovering Market Expectations from Futures Prices with an Application to Crude Oil," Discussion paper.

Berkowitz, J., I. Biegean, and L. Kilian (1999): "On the Finite-Sample Accuracy of Nonparametric Resampling Algorithms for Economic Time Series," Discussion paper.

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.

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," Temi di discussione (Economic working papers) 1027, Bank of Italy, Economic Research and International Relations Area.

Cecchetti, S. G. (2008): "Measuring the Macroeconomic Risks Posed by Asset Price Booms," *Asset Prices and Monetary Policy*, pp. 9–43.

Chernozhukov, V., I. Fernandez-Val, and A. Galichon (2010): "Quantile and Probability Curves without Crossing," *Econometrica*, 78(3), 1093–1125.

Chevalier, J. A., and D. S. Scharfstein (1996): "Capital-Market Imperfections and Countercyclical Markups: Theory and Evidence," *The American Economic Review*, 86(4), 703–725.

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2015): "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, 7(1), 110–67.

Christiano, L.J., R. Motto, and M. Rostagno (2014): "Risk Shocks," *American Economic Review*, 104(1), 27–65.

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.

Ghysels, E., L. Iania, and J. Striaukas (2018): "Quantile-Based Infflation Risk Models," Working Paper Research 349, National Bank of Belgium.

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), 1692–1720.

——— (2015): "Inflation Dynamics Trough Firms' Pricing Behavior," in *Inflation Dynamics and Monetary Policy*, ed. by F. R. B. of Kansas City, Jackson Hole Economic Policy Symposium. Federal Reserve Bank of Kansas City.

KILEY, M. T. (2018): "Unemployment Risk," Finance and Economics Discussion Series 2018-067, Board of Governors of the Federal Reserve System (US).

KILIAN, L., AND H. LÜTKEPOHL (2018): *Structural Vector Autoregressive Analysis*, no. 9781107196575 in Cambridge Books. Cambridge University Press.

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.

LORIA, F., C. MATTHES, AND D. ZHANG (2019a): "Assessing Macroeconomic Tail Risk," Finance and Economics Discussion Series 2019-026, Board of Governors of the Federal Reserve System (US).

——— (2019b): "The Effect of Oil Price Shocks on Inflation Risk," *Mimeo*.

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," Discussion Papers 1815, Centre for Macroeconomics (CFM).

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. Chapter 1: The Vulnerability of Euro Area Inflation Forecasts," Doctoral thesis, Universitat Autònoma de Barcelona.

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?". A speech at the "Rethinking Macro Policy II", a conference sponsored by the International Monetary Fund, Washington, D.C., April 16, 2013," Speech 622, Board of Governors of the Federal Reserve System (U.S.).

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: `CPILFESL_PCA`, 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: `PCEPILFE_PCA`, 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. Before that date we use the series from Blanchard, Cerutti, and Summers (2015).
- 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), Continuously Compounded Annual Rate of Change, Quarterly, Seasonally Adjusted.

- Oil Price
  - Source: FRED.
  - Details: `WTISPLC_CCA`, Spot Crude Oil Price: West Texas Intermediate (WTI), Continuously 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: `NFCI`, 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).
- 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.

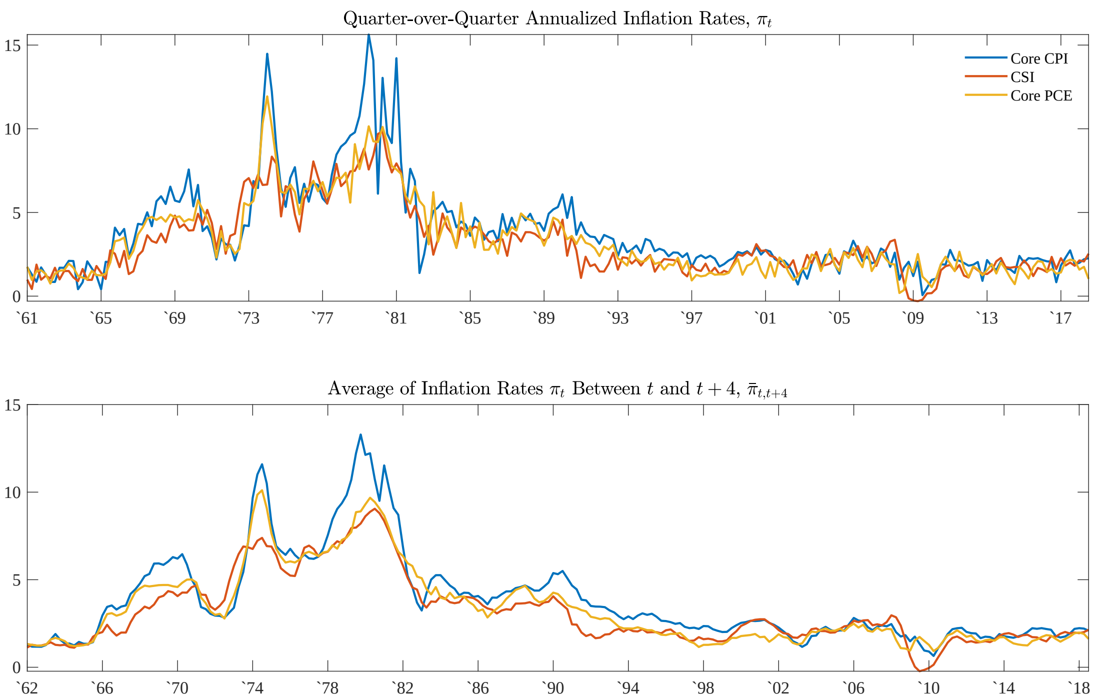

Figure A-2: Regressors, United States.

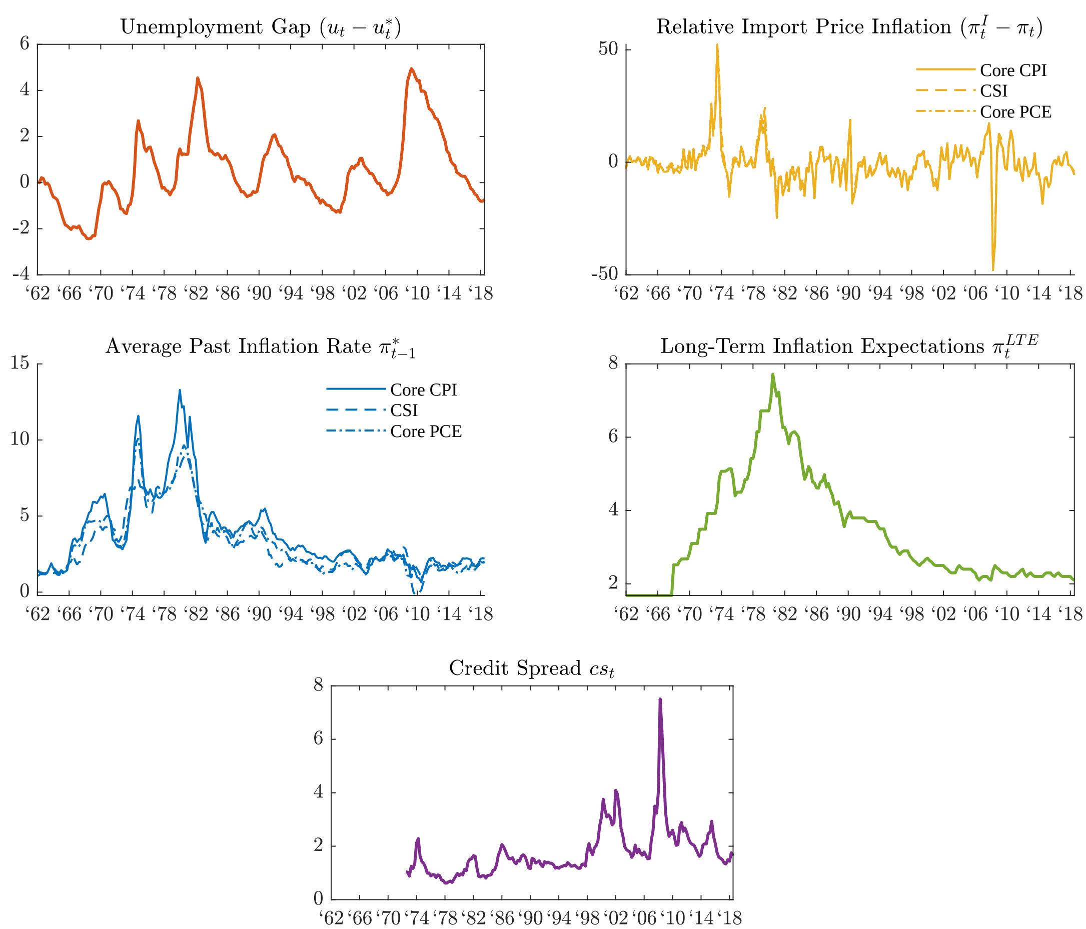

A.2 Euro Area

- 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.

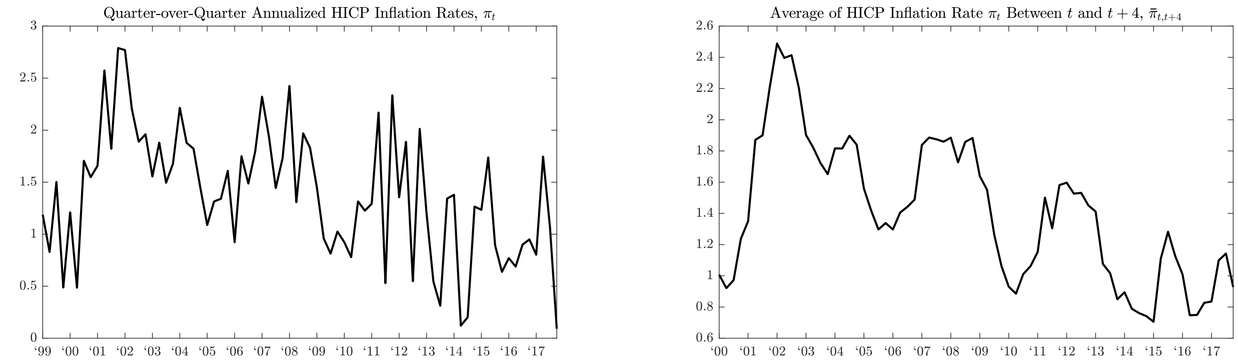

Figure A-4: Regressors, Euro Area.

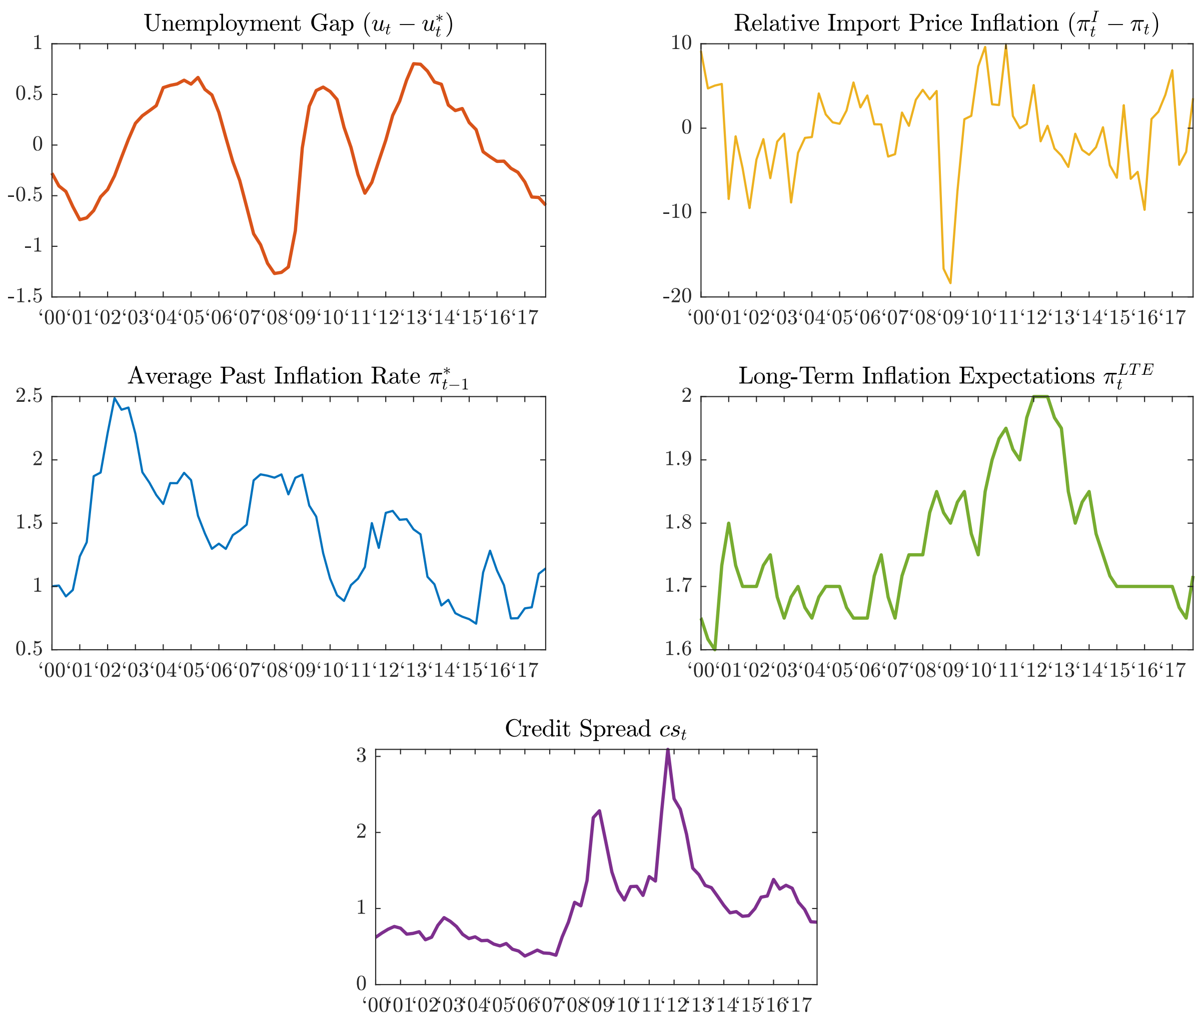

B Bootstrap Method

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 l 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 $T \to \infty$. Following Berkowitz, Biegean, and Kilian (1999) we set $l = m = \sqrt[3]{T}$, 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.

Figure B-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 10$^{\text{th}}$ quantile (blue), median (red) and 90$^{\text{th}}$ quantile (yellow). The estimation period is 1973:Q1 to 2019:Q1. The OLS estimates and their 95% confidence intervals are represented by the red lines.

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

United States Core CPI

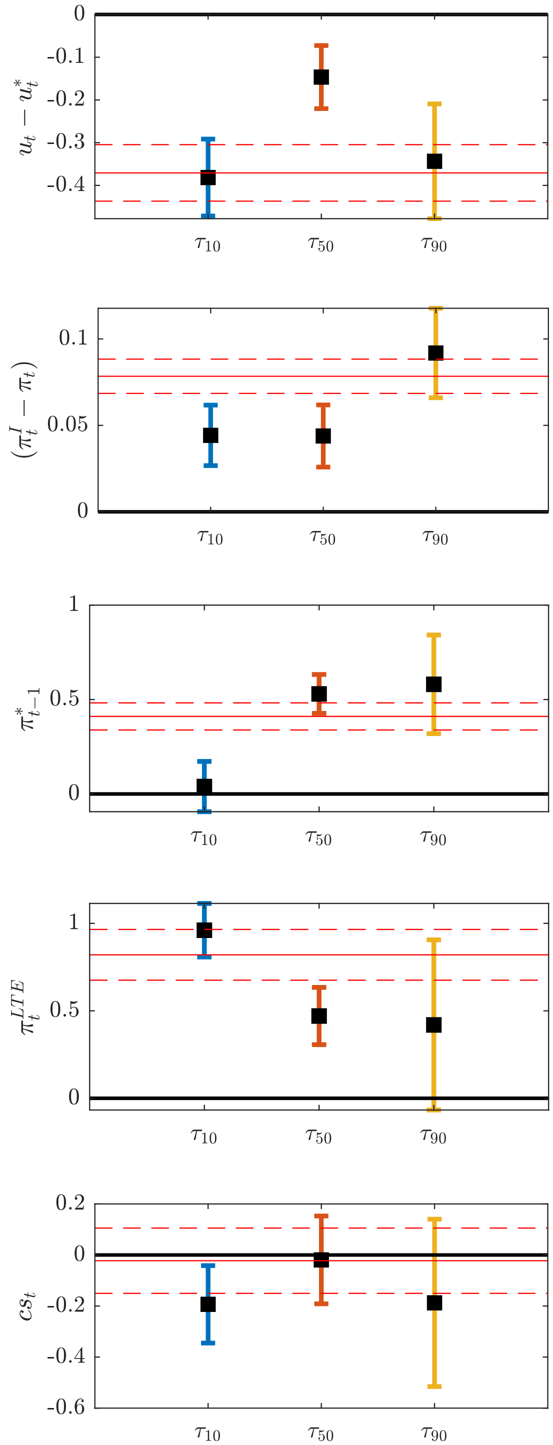

NOTE: The figure displays the slope coefficients of the quantile regression of average four-quarter-ahead United States Core CPI inflation defined in (2). The black squares correspond to the point estimates whereas the vertical lines to the 68% confidence intervals computed via "blocks-of-blocks" bootstrap 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 represented by the red lines.

C 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). The estimated slopes of the quantile regressions are plotted in Figures C-1 and Figure C-2, respectively. The slopes are very robust for the economic variables. The effects of the credit spread on CSI are consistent with those obtained for core CPI, but less so for core PCE. 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 C-3 and Figure C-4. These figures confirm 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 6). Importantly, once we control for subsample stability, the results are extremely similar across these different inflation measures.

Figure C-5 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 B. 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 C-6 confirms the important influence of credit spreads on the 10^{th} quantile of the distribution both for core PCE and for CSI inflation. This figure mimics the top-right panel of Figure 6 in the main text.

Figure C-7, Figure C-8, and Figure C-9 display similar exercises to those presented in the main text for these two alternative measures of inflation, core PCE and CSI, respectively.

Figure C-1: Quantile Regression Slopes, Core PCE.

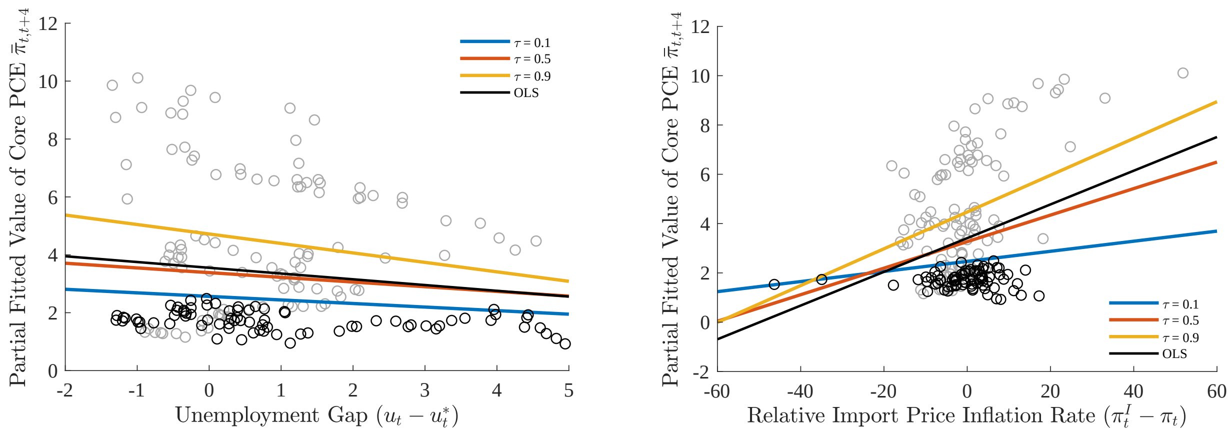

$\hat{\theta}_\tau = \{\hat{\theta}_{0.1} = -0.12, \hat{\theta}_{0.5} = -0.16, \hat{\theta}_{0.9} = -0.33\}$ $\hat{\gamma}_\tau = \{\hat{\gamma}_{0.1} = 0.02, \hat{\gamma}_{0.5} = 0.05, \hat{\gamma}_{0.9} = 0.07\}$

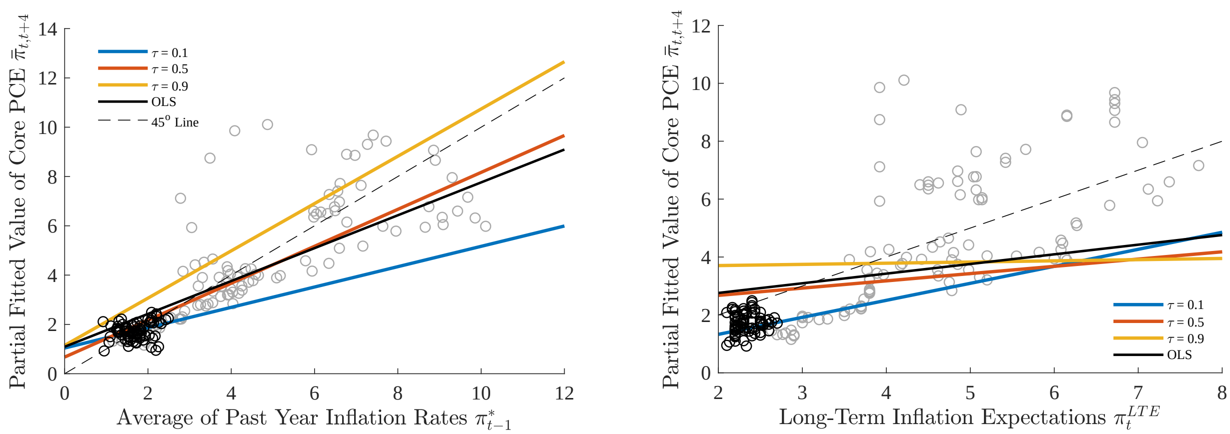

$\hat{\lambda}_\tau = \{\hat{\lambda}_{0.1} = 0.41, \hat{\lambda}_{0.5} = 0.75, \hat{\lambda}_{0.9} = 0.96\}$, where λ_τ is coefficient on π_t^{LTE} and $(1 - \lambda_\tau)$ on π_{t-1}^*

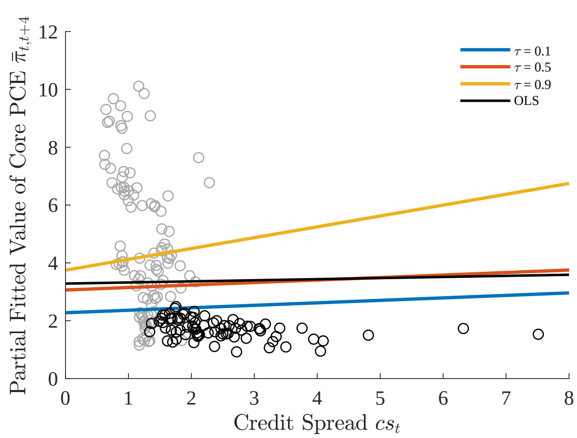

$\hat{\delta}_\tau = \{\hat{\delta}_{0.1} = 0.09, \hat{\delta}_{0.5} = 0.09, \hat{\delta}_{0.9} = 0.38\}$

NOTE: The figure displays the slope coefficients of the quantile regression of average four-quarter-ahead core PCE 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.

Figure C-2: Quantile Regression Slopes, Stock and Watson (2019) CSI.

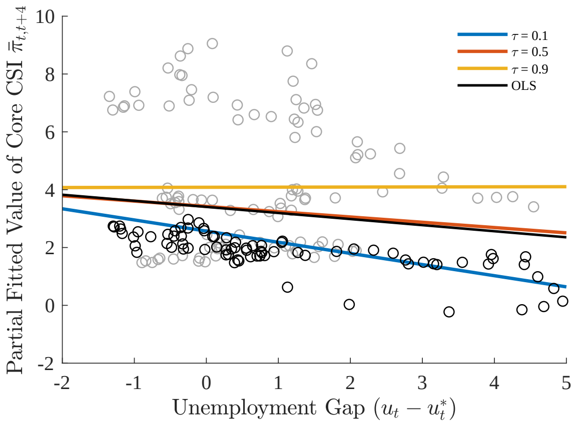

$\hat{\theta}_\tau = \{\hat{\theta}_{0.1} = -0.38, \hat{\theta}_{0.5} = -0.18, \hat{\theta}_{0.9} = 0\}$

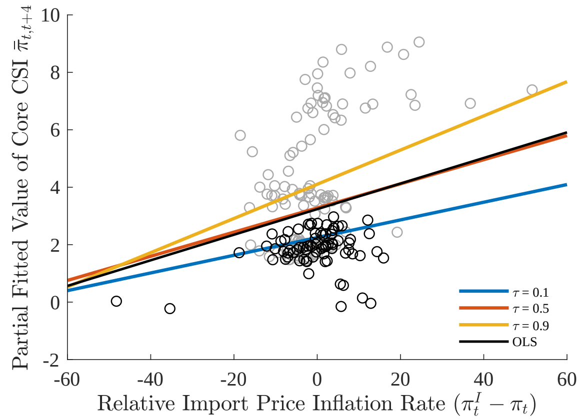

$\hat{\gamma}_\tau = \{\hat{\gamma}_{0.1} = 0.03, \hat{\gamma}_{0.5} = 0.04, \hat{\gamma}_{0.9} = 0.06\}$

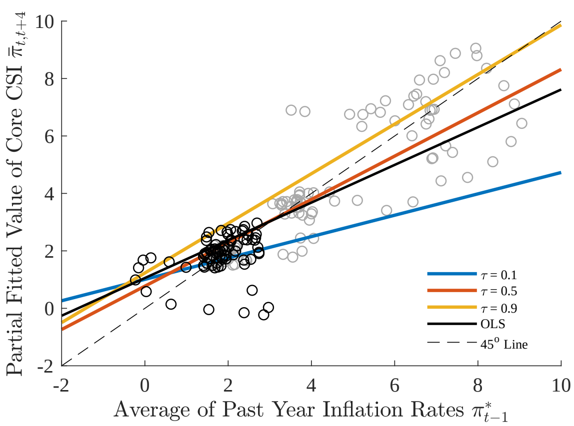

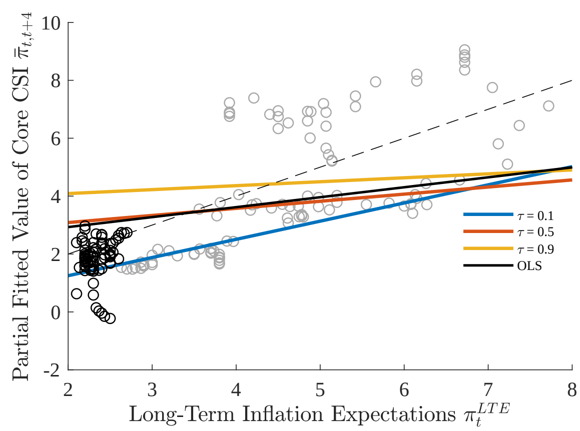

$\hat{\lambda}_\tau = \{\hat{\lambda}_{0.1} = 0.63, \hat{\lambda}_{0.5} = 0.25, \hat{\lambda}_{0.9} = 0.14\}$, where λ_τ is coefficient on π_t^{LTE} and $(1 - \lambda_\tau)$ on π_{t-1}^*

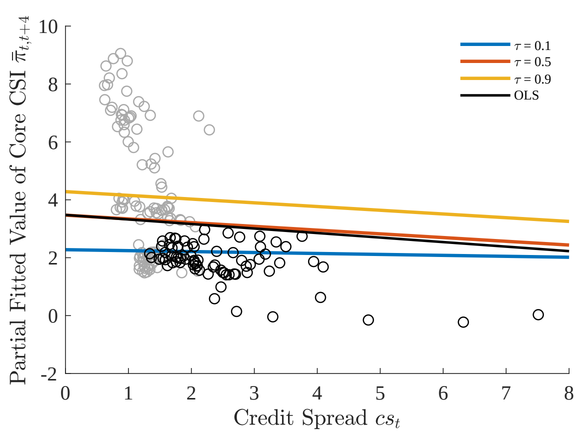

$\hat{\delta}_\tau = \{\hat{\delta}_{0.1} = -0.03, \hat{\delta}_{0.5} = -0.13, \hat{\delta}_{0.9} = -0.13\}$

NOTE: The figure displays the slope coefficients of the quantile regression of average four-quarter-ahead Stock and Watson (2019) Cyclically Sensitive 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.

Figure C-3: Quantile Regression Slopes Across Subsamples, Core PCE.

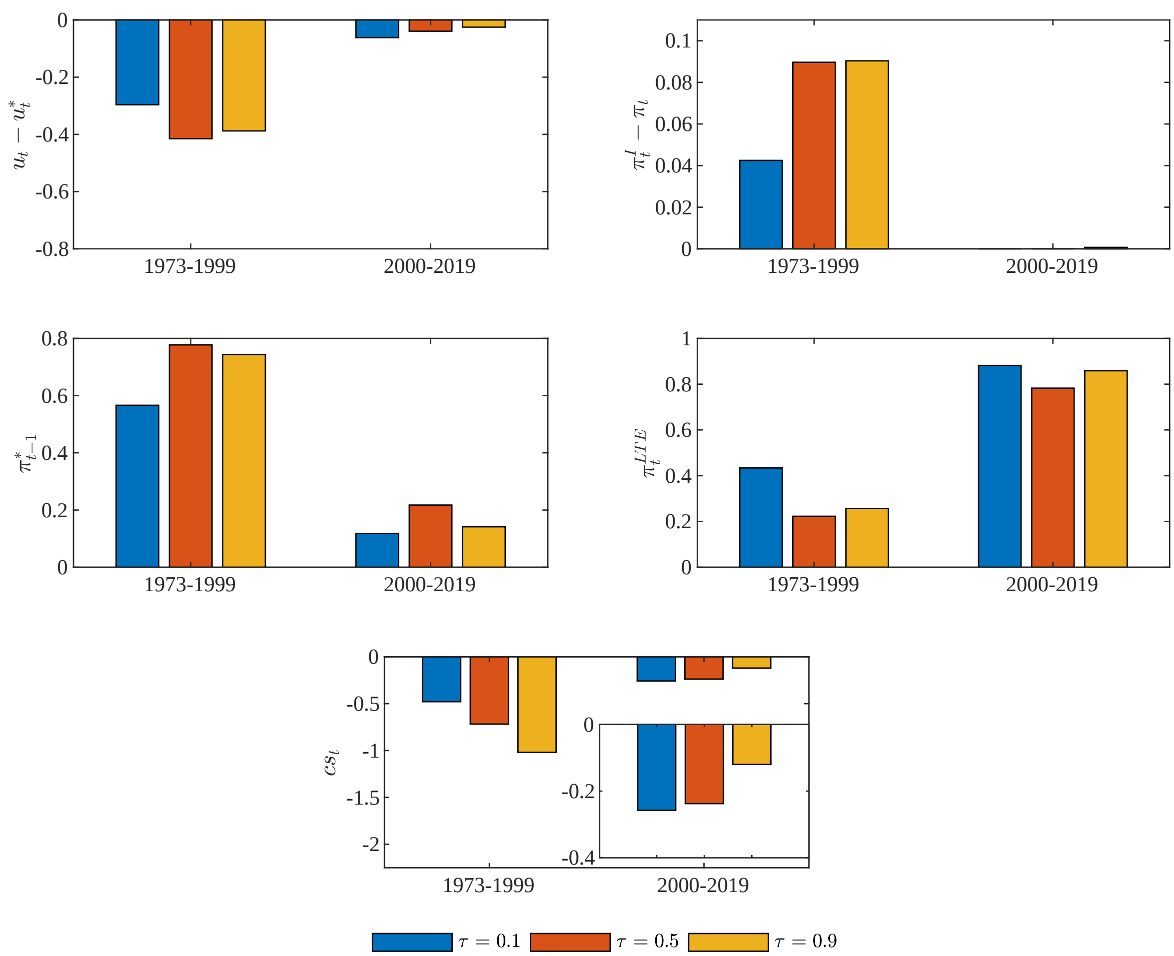

NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead core PCE inflation defined in expression (2). Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate the coefficients associated with the 10$^{\text{th}}$ quantile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

Figure C-4: Quantile Regression Slopes Across Subsamples, Stock and Watson (2019) CSI.

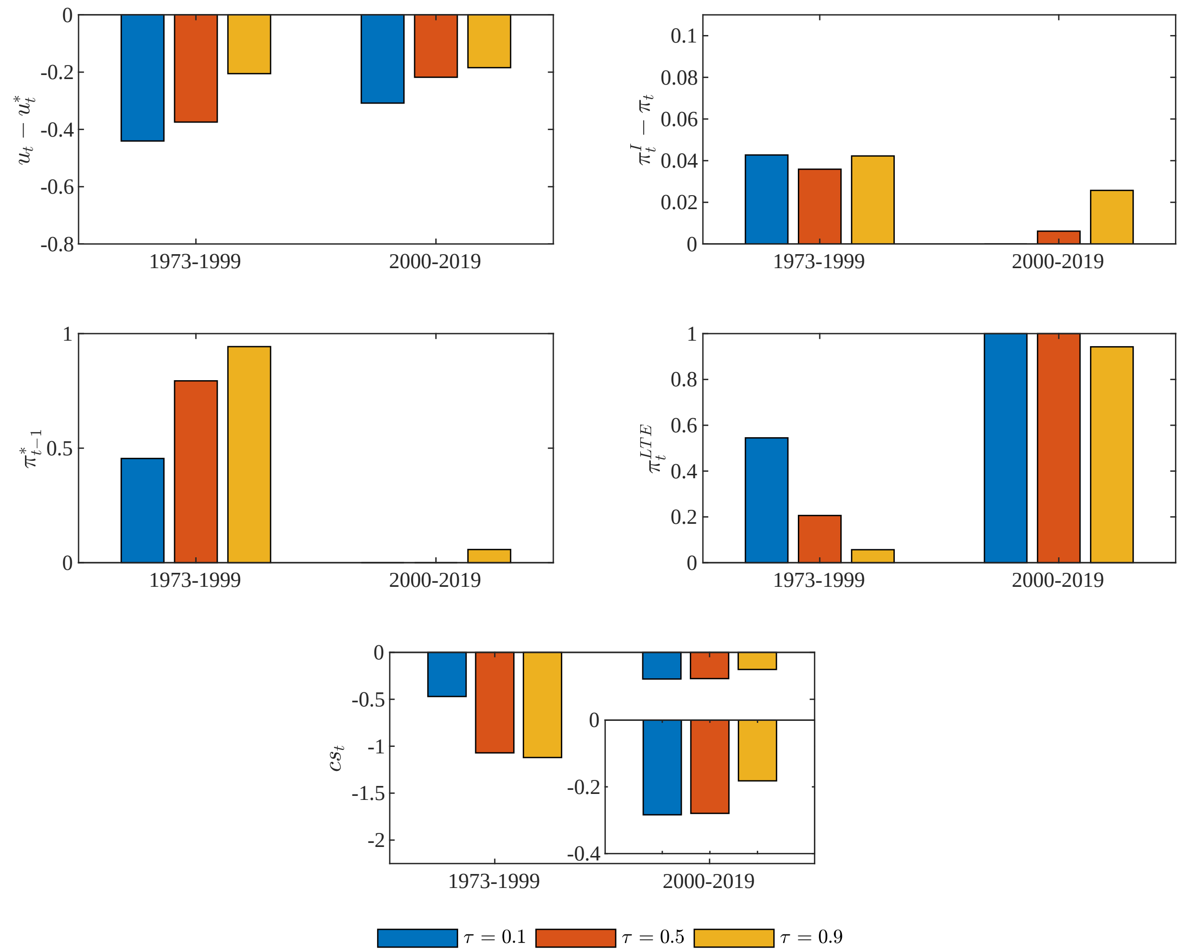

NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead Stock and Watson (2019) Cyclically Sensitive Inflation defined in expression (2). Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate the coefficients associated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

Figure C-5: Quantile Regression Slopes and Confidence Intervals.

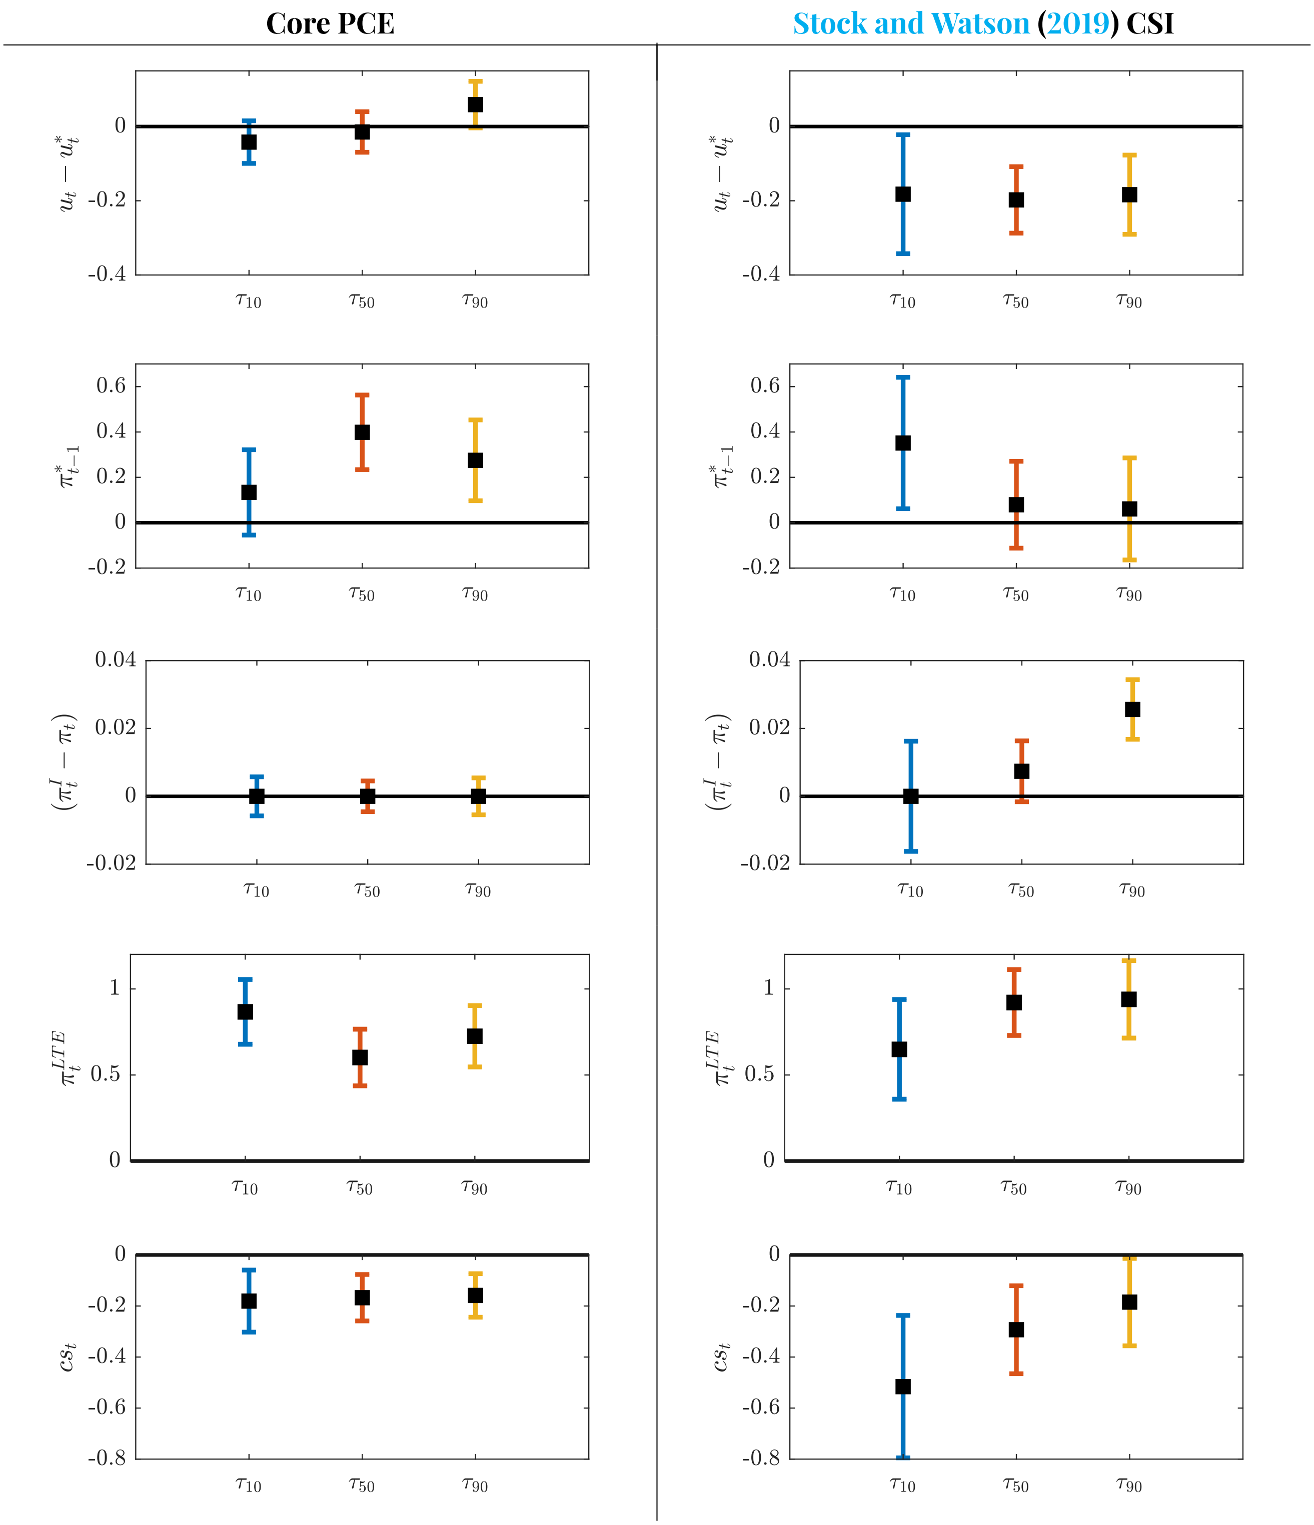

NOTE: The figure displays the slope coefficients of the quantile regression of average four-quarter-ahead of core PCE inflation (left) and Stock and Watson (2019) Cyclically Sensitive Inflation (right) defined in (2). The black squares correspond to the point estimates whereas the vertical lines to the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

Figure C-6: Partial Effect of Credit Spread on 10th Inflation Quantiles.
Core PCE and Stock and Watson (2019) CSI Inflation.

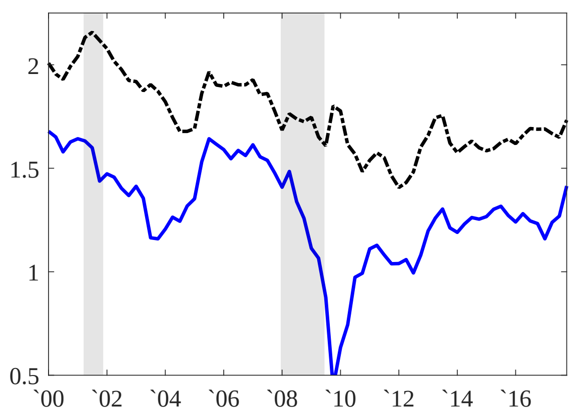

Core PCE

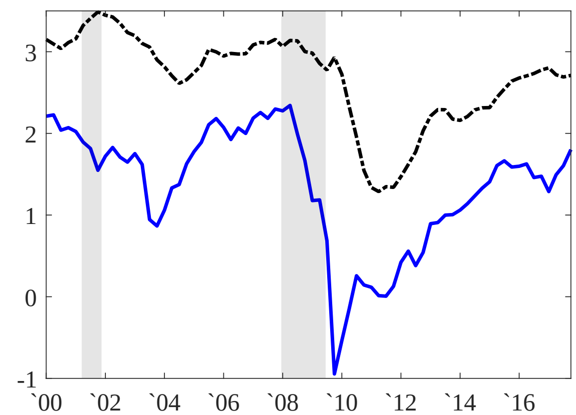

Stock and Watson (2019) CSI

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 straight) and in its version where the effect of credit spreads is set to zero (black dash-dotted). Shaded bars indicate NBER-dated recessions.

Figure C-7: Selected Time Episodes of Predictive Densities (Left) and Skewness (Right). Core PCE Inflation.

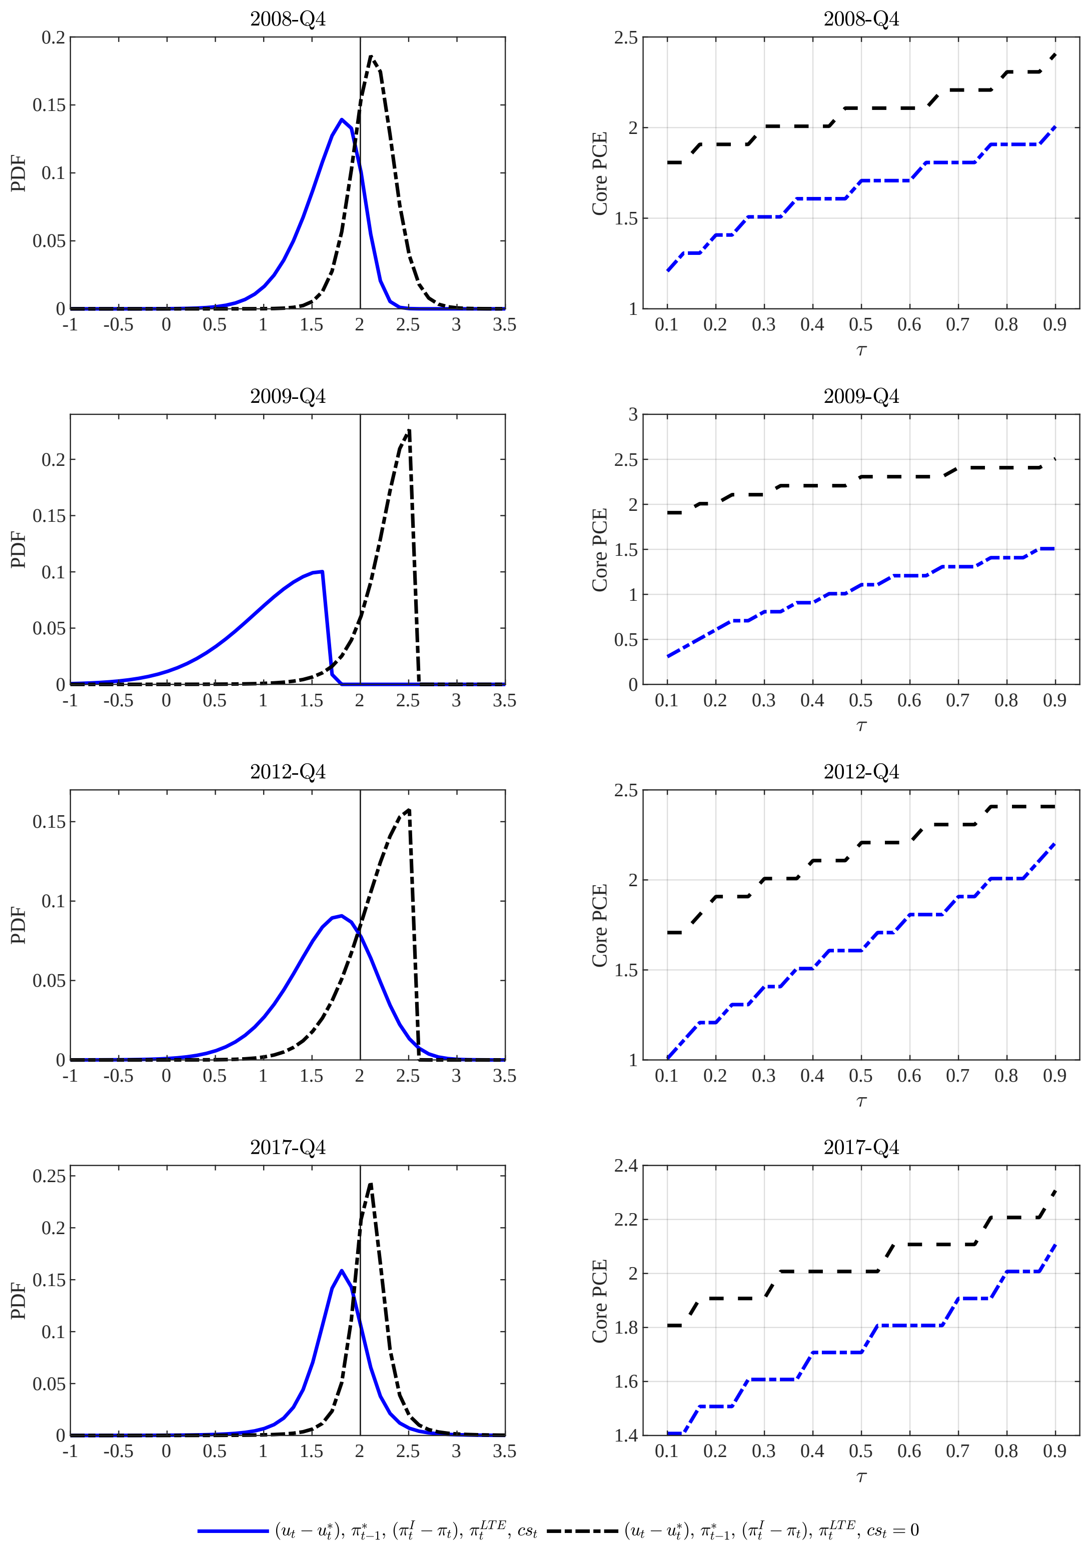

NOTE: The left panels show the estimated skewed t-Student densities of average four-quarter-ahead core PCE inflation for alternative specifications of the quantile regressions model (2), in its baseline version (blue straight) and in its version where the effect of credit spreads is set to zero (black dash-dotted). The right panels show the estimated skewed t-inverse cumulative associated with the conditional densities in the left panels.

Figure C-8: Selected Time Episodes of Predictive Densities (Left) and Skewness (Right). Stock and Watson (2019) CSI Inflation.

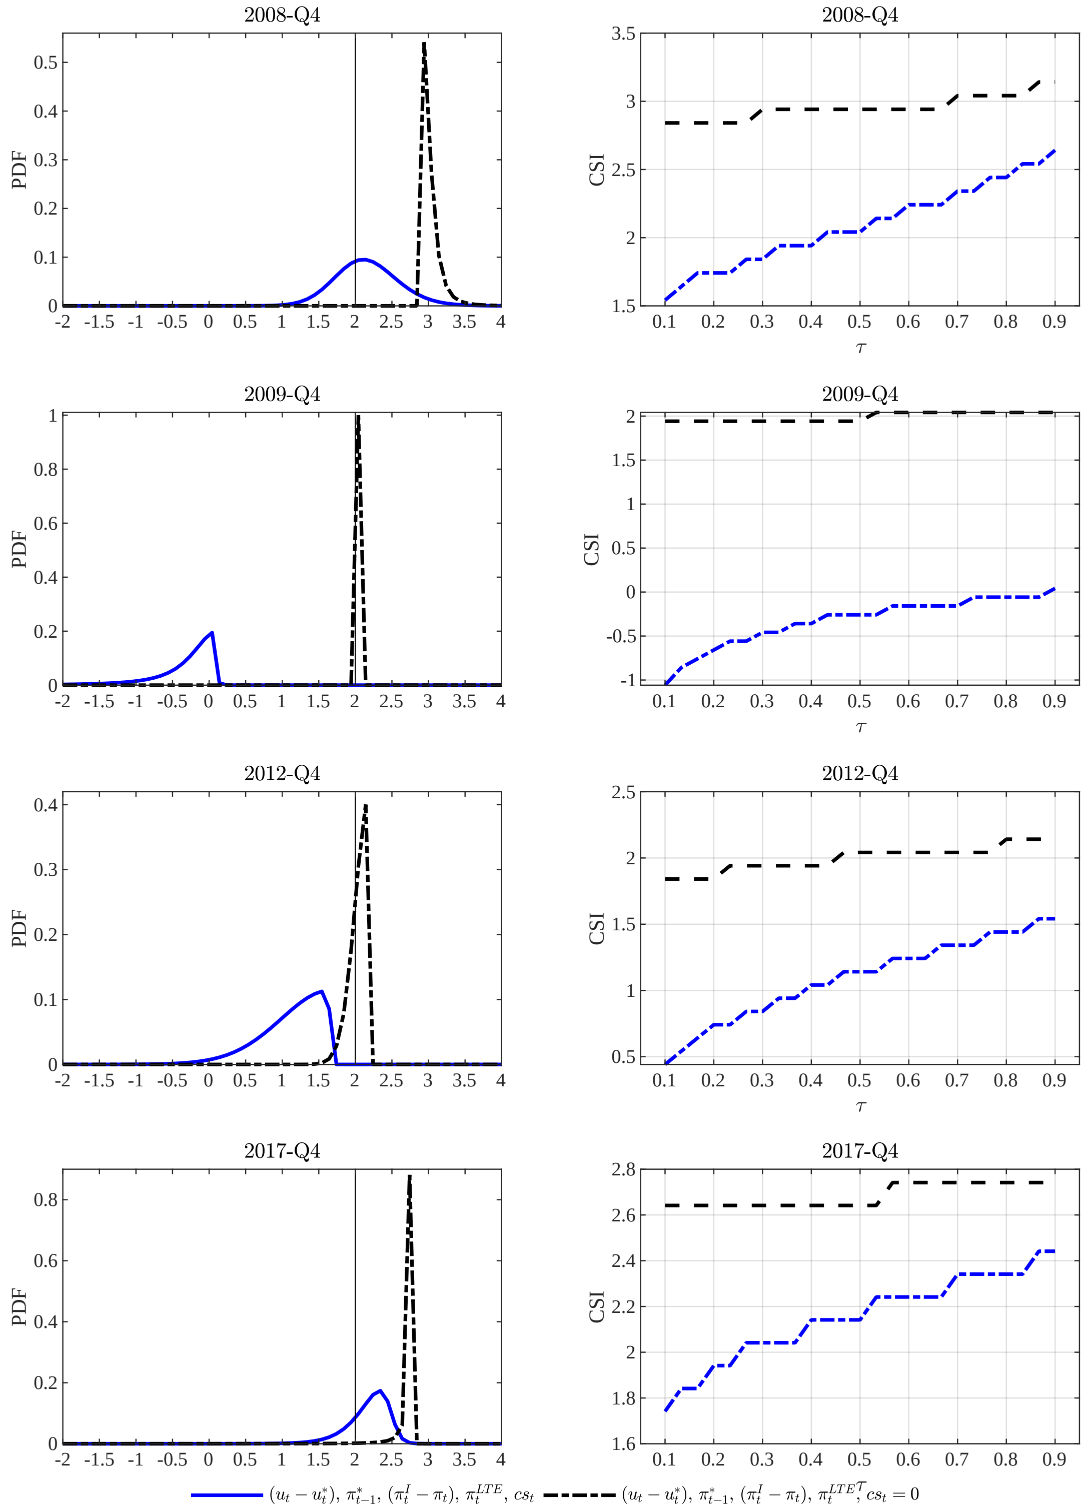

NOTE: The left panels show the estimated skewed t-Student densities of average four-quarter-ahead Stock and Watson (2019) Cyclically Sensitive Inflation for alternative specifications of the quantile regressions model (2), in its baseline version (blue straight) and in its version where the effect of credit spreads is set to zero (black dash-dotted). The right panels show the estimated skewed t-inverse cumulative associated with the conditional densities in the left panels.

Figure C-9: Inflation Probabilities for Alternative Cutoff Values.
Core PCE and Stock and Watson (2019) CSI Inflation.

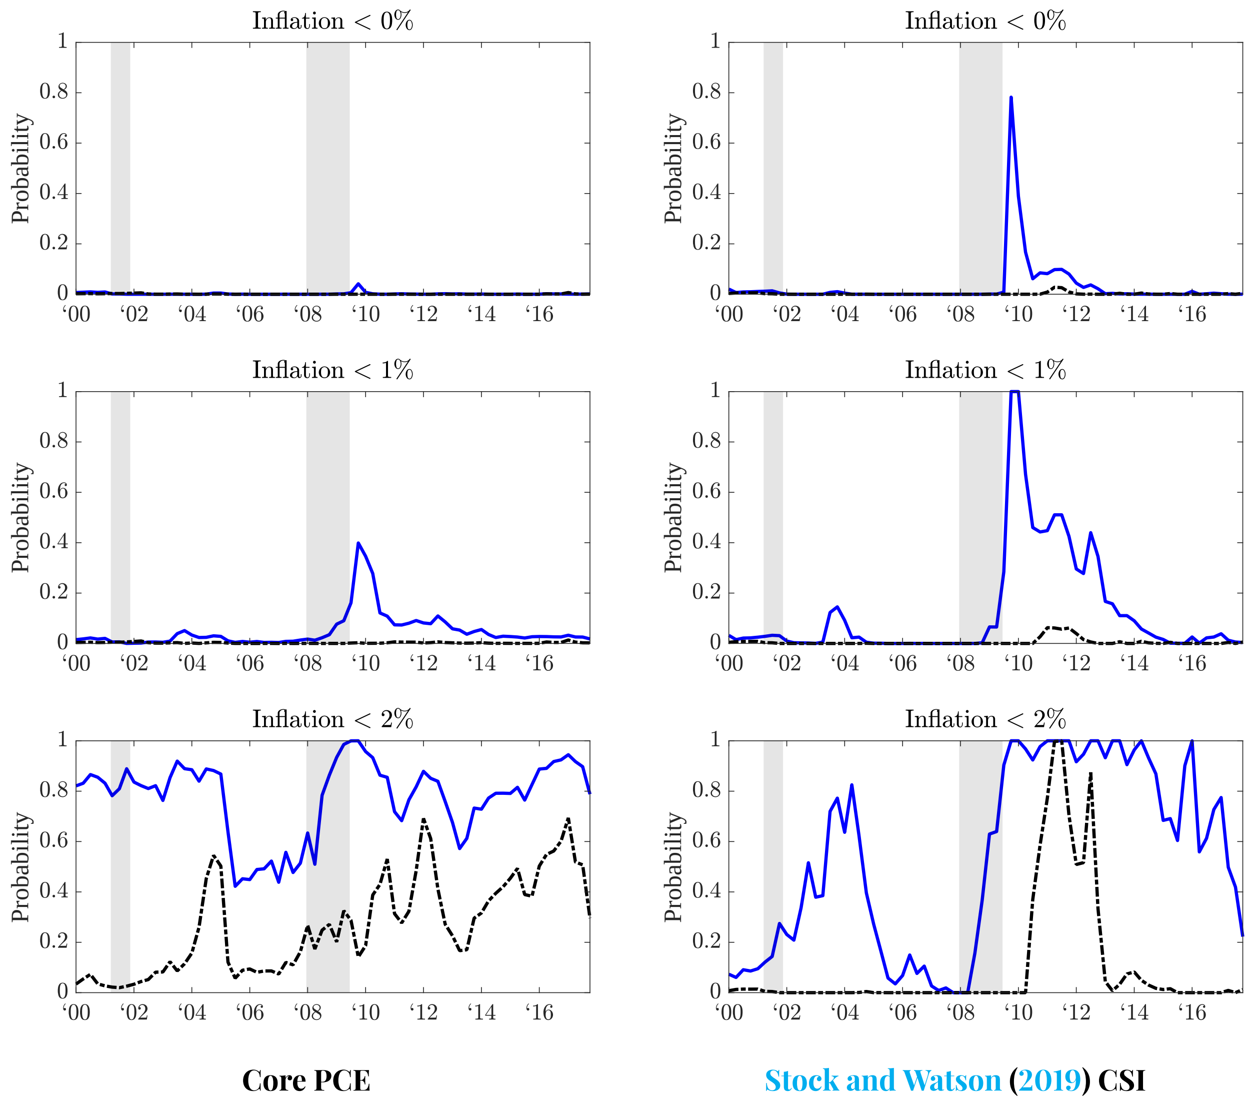

NOTE: The figure shows the time evolution of inflation probabilities of core PCE inflation (left) and Stock and Watson (2019) Cyclically Sensitive Inflation (right) for different cutoffs. These probabilities are computed from the skewed t-Student conditional densities of the average four-quarter-ahead 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 without and with the credit spread (in blue straight and black dash-dotted lines, respectively). Shaded bars indicate NBER-dated recessions.

D Robustness

In this Appendix we present additional robustness results for the U.S. economy. 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 D-1). Figure D-2 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. As noted, the substantial subsample instability is responsible for fuzzing the effects on the upper tail. This is confirmed by the subsample stability results shown in Figure D-3, Figure D-4, and Figure D-5.

D.1 Oil vs Import

Figure D-1: Quantile Regressions Slopes Across Relative Price Measures.

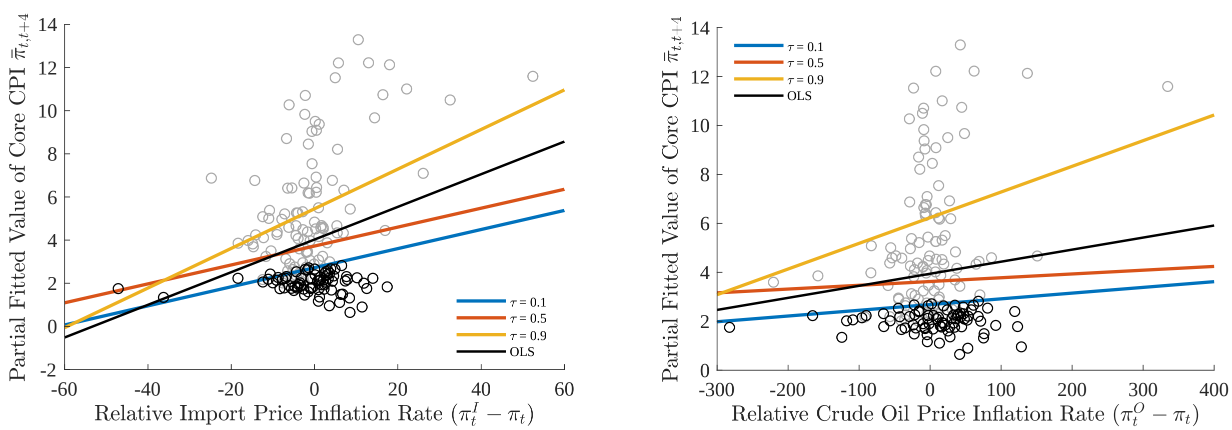

Slope γ_τ^I on $(\pi_t^I - \pi_t)$ Slope γ_τ^O on $(\pi_t^O - \pi_t)$

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 10^{th} (blue) and the 90^{th} (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.

D.2 Different Financial Variables

Figure D-2: Quantile Regressions Slopes Across Financial Variables.

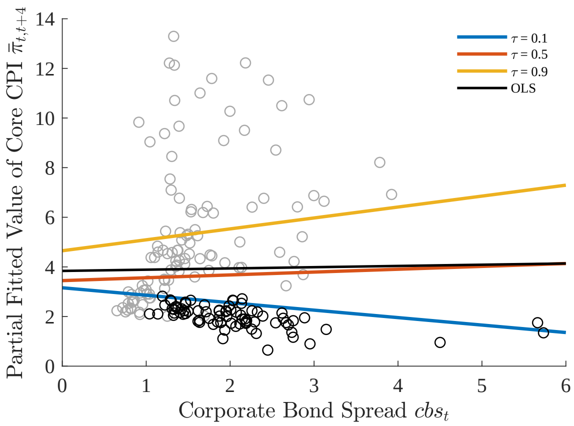

Slopes δ_τ on cbs_t

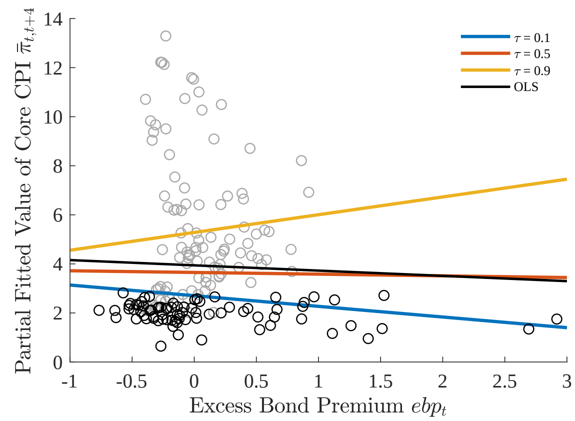

Slopes δ_τ on ebp_t

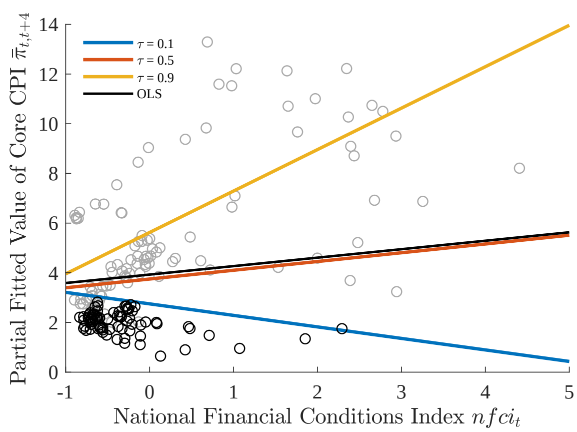

Slopes δ_τ on $nfci_t$

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 10$^{\text{th}}$ (blue) and the 90$^{\text{th}}$ (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.

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

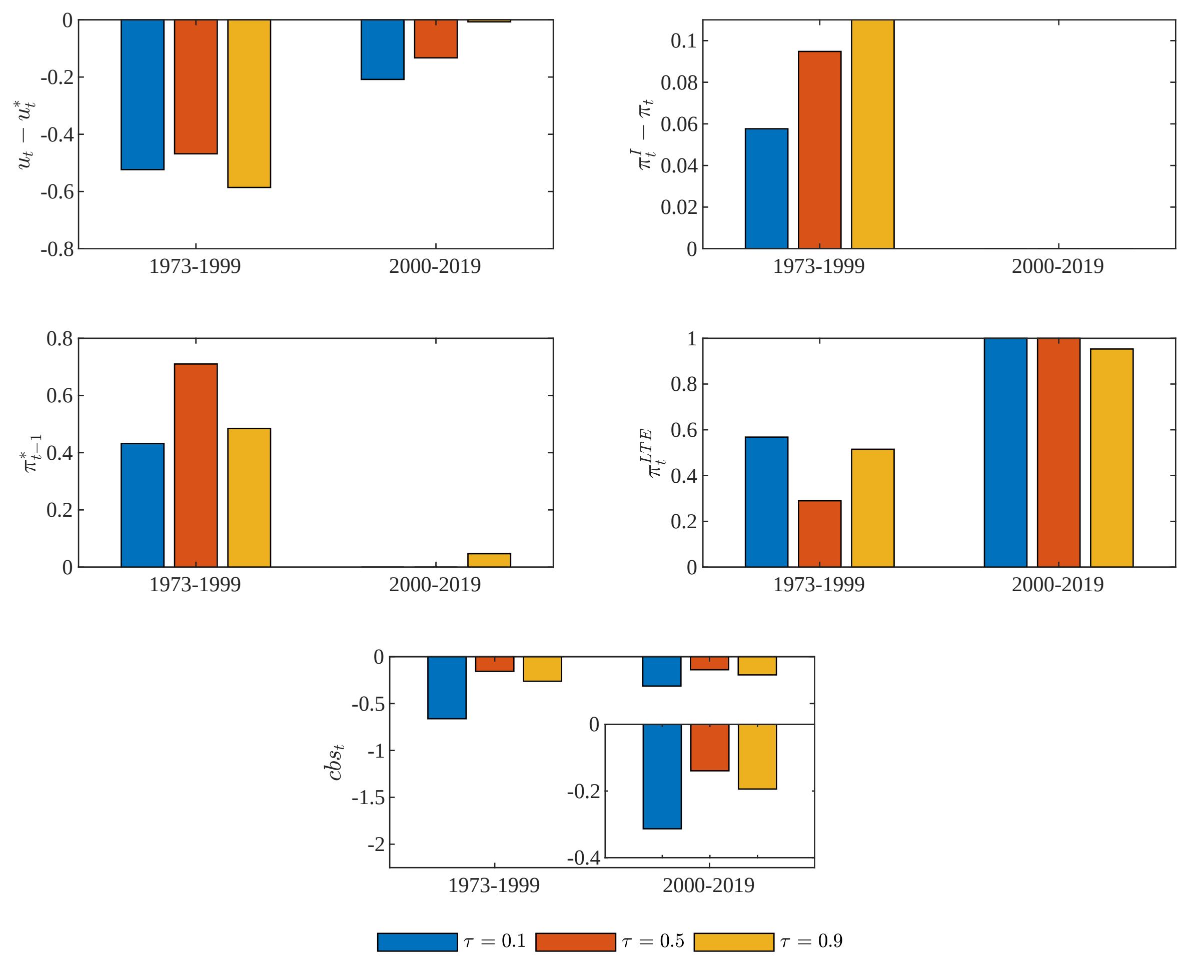

NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead core CPI inflation defined in expression (2), using corporate bond spreads and the same sample period as the baseline. Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The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

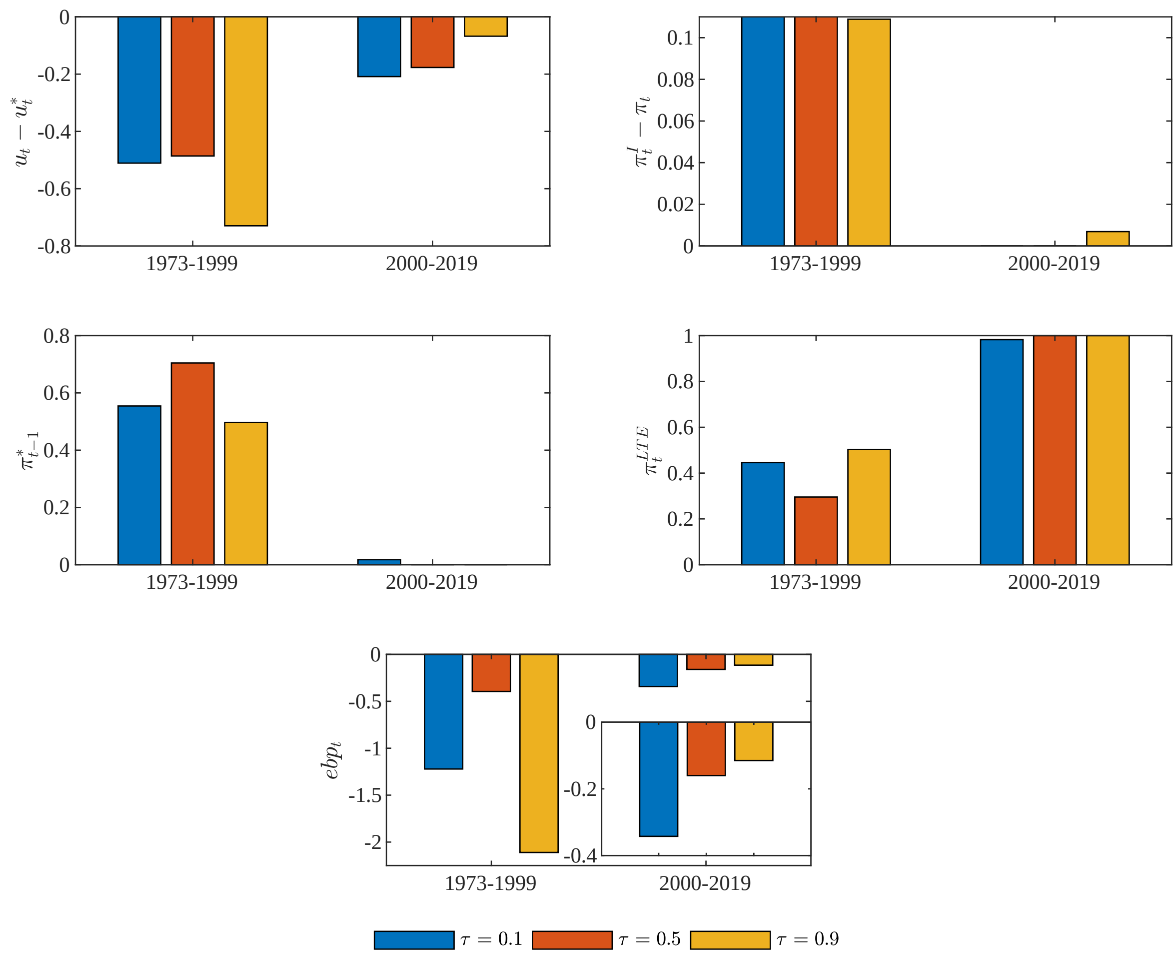

NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead core CPI inflation defined in expression (2), using the Gilchrist and Zakrajšek (2012) excess bond premium and the same sample period as the baseline. Two different subsamples are considered: (i) 1973-1999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

Figure D-5: Quantile Regression Slopes Across Subsamples, NFCI.

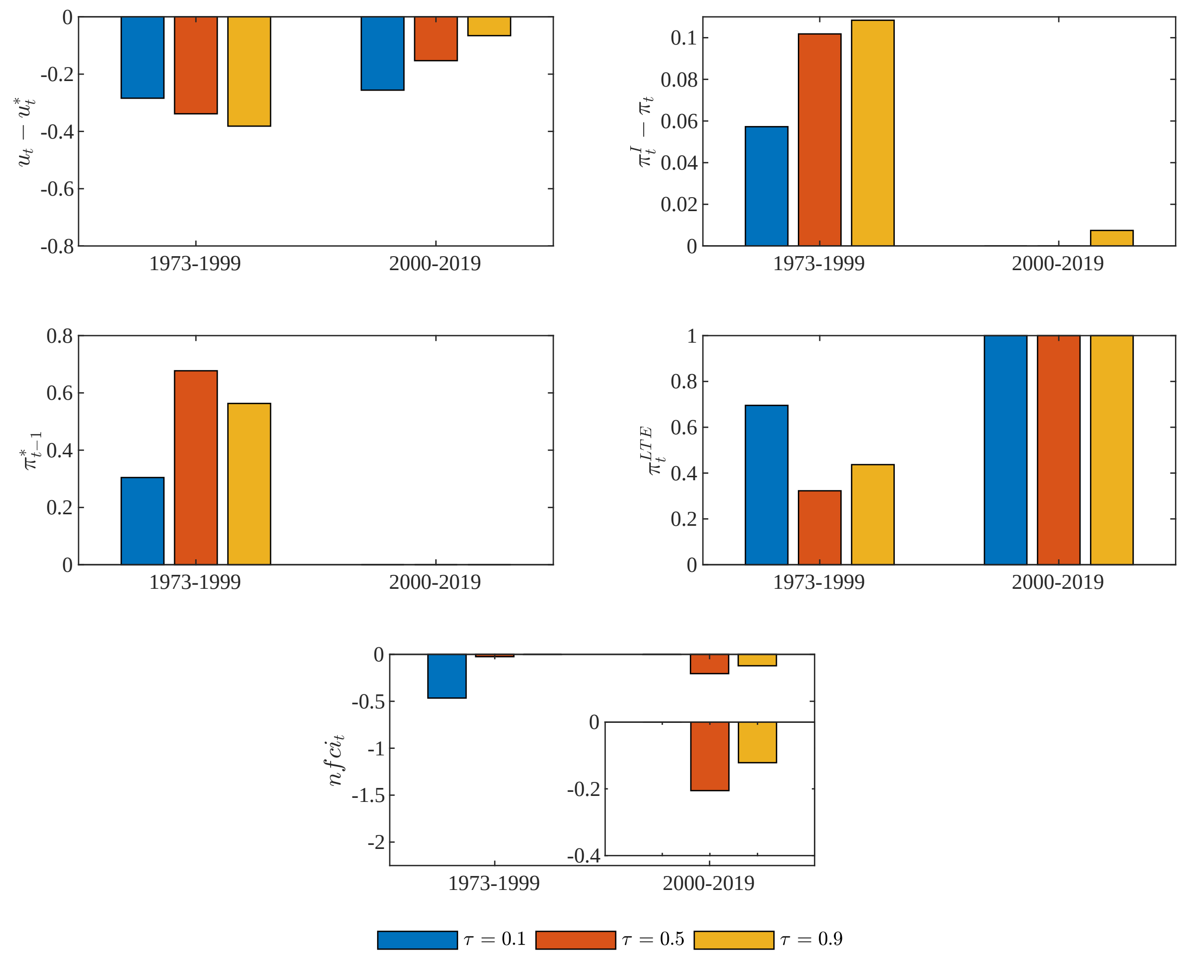

NOTE: The figure displays the estimated slopes of the quantile regression of average four-quarter-ahead core CPI inflation defined in expression (2), using the NFCI and the same sample period as the baseline. Two different subsamples are considered: (i) 1973-1999 and (iii) 2000-2019. The bars illustrate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

E The Role of Skewness

Figure E-1 below 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 E-1: Inflation Probabilities and The Role of Skewness.

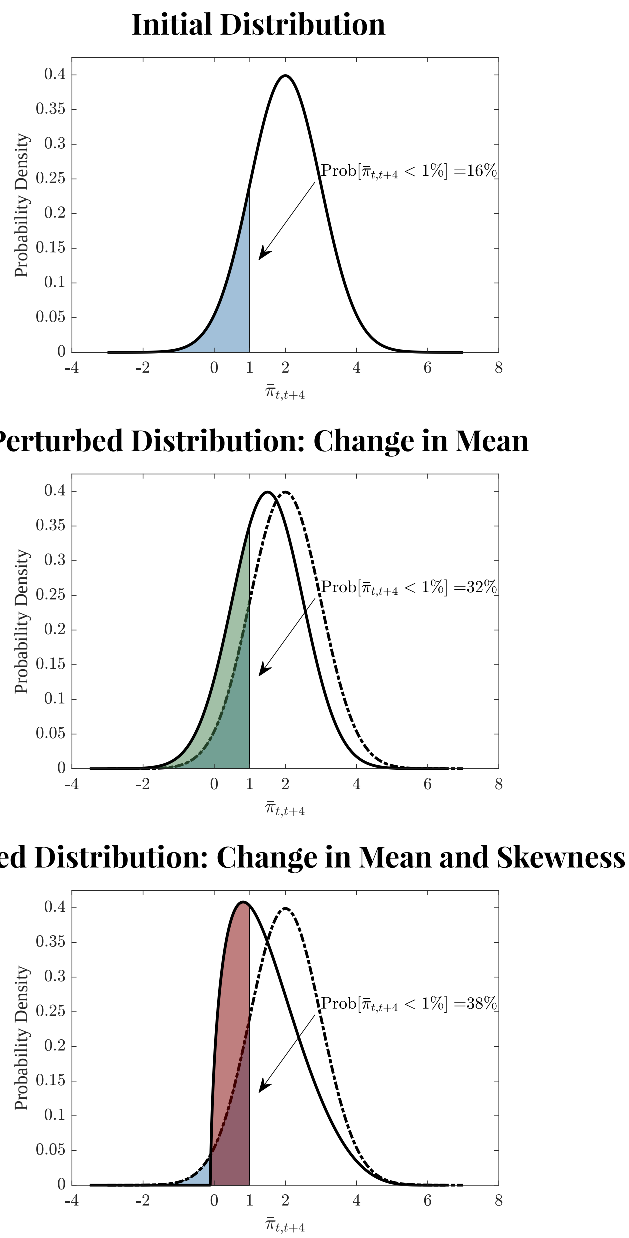

NOTE: The figure displays the three states associated with a change in an inflation determinant that causes the initial normal density (top panel) to feature a lower mean (center panel) and then also a right-skew (bottom panel).

F Financial Market Probabilities

Figure F-1: 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-2: Inflation Probabilities from Financial Markets vs. Credit Spread.

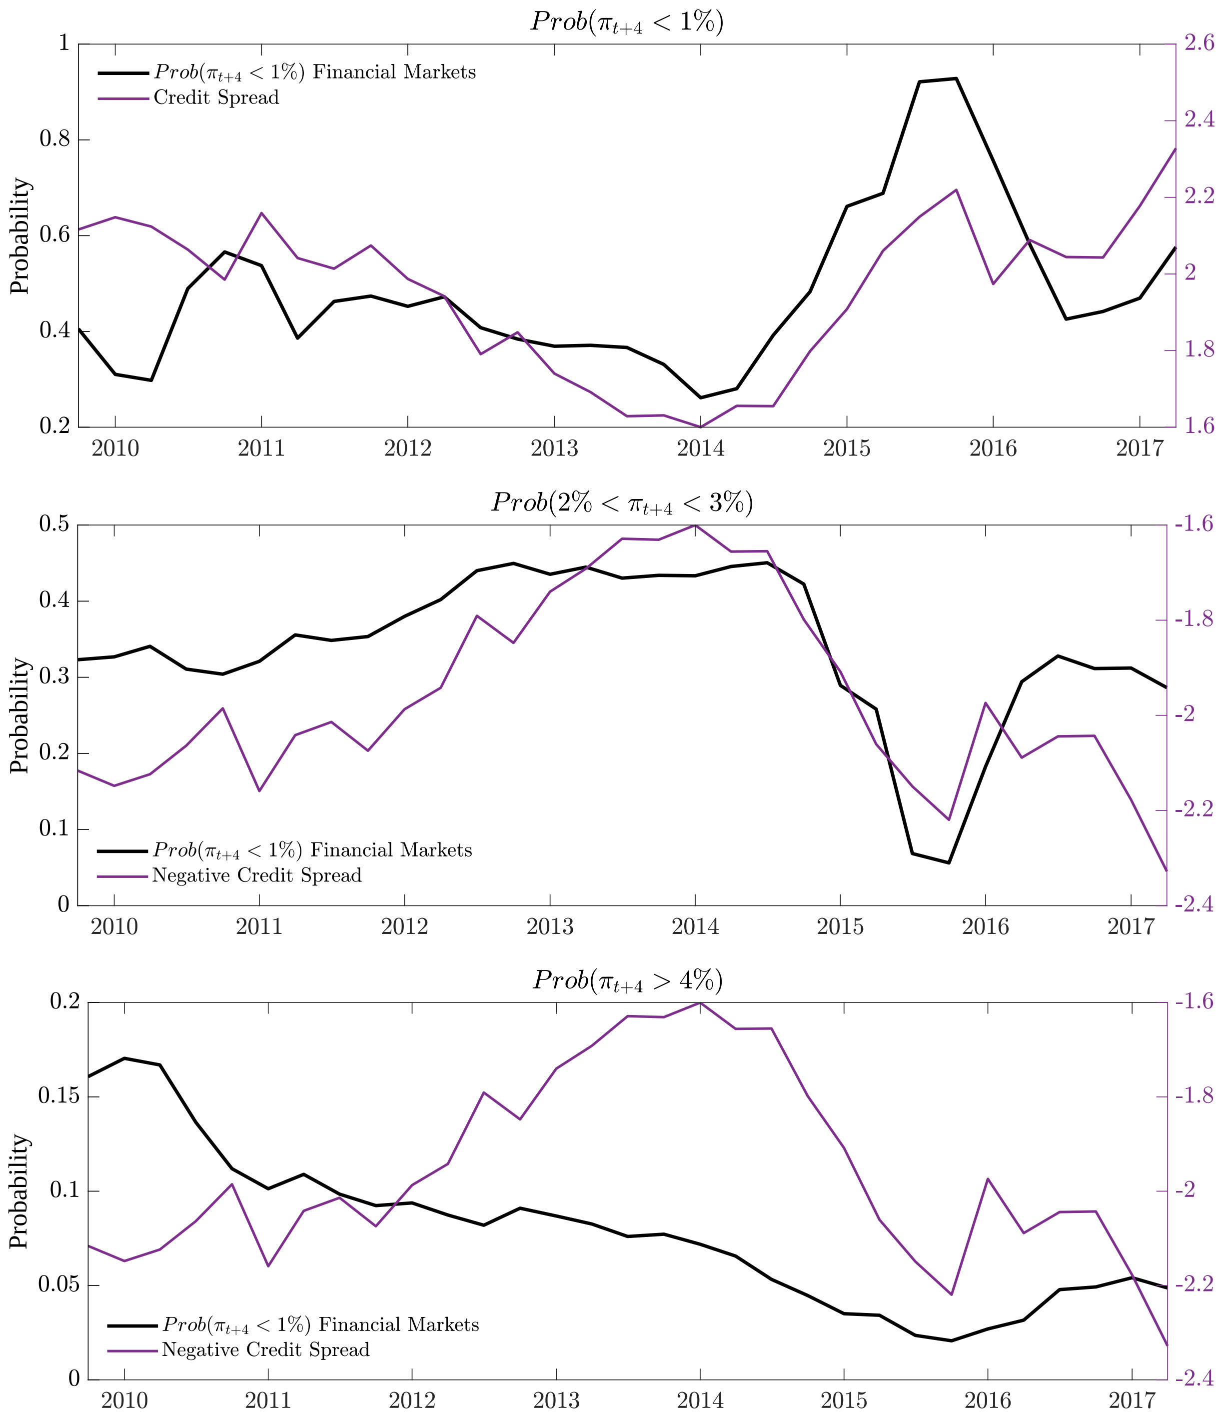

NOTE: The figures shows the credit spread against quarterly options-implied inflation probabilities of headline CPI inflation next year being above 1% (top panel), between 2% and 3% (mid panel) and above 4% (top panel).