Conditional Inflation-at-Risk

Markeljan Blacëri*

University of Padova

November 2026

Abstract

This paper introduces Conditional Inflation-at-Risk (CoIaR), a measure of how the tails of the future inflation distribution change when a given economic variable is in a stressed state. The measure attributes inflation tail risk to a specific source and decomposes it into the sensitivity of a given inflation quantile to that source and the risk in the source itself. I apply the framework to the labor market, using labor market tightness as the conditioning variable. Using monthly U.S. data, I find that the labor market is primarily a source of downside inflation risk. A slack labor market lowers the 10th percentile of core inflation one year ahead by about half a percentage point, while a tight labor market leaves the upper tail essentially unchanged. Upside inflation risk is instead driven by inflation persistence and expectations. These results suggest that labor market slack is informative about the risk of undershooting the inflation target, but provides limited warning of upside inflation risk.

Keywords: inflation risk, Phillips curve, labor market tightness, CoIaR
JEL codes: E31, E24, E37, C21

*Department of Economics and Management, University of Padova, Via del Santo 33, 35123 Padova, Italy. Email: markeljan.blaceri@phd.unipd.it.

I would like to thank Giovanni Caggiano and participants at the Padova Macro Brown Bag at the University of Padova for valuable discussions and suggestions that greatly improved this manuscript. All remaining errors are my own.

1 Introduction

The inflation surge of 2021–2023 shifted the attention of central banks from the most likely path of inflation to the risks around it. Policymakers now routinely describe the balance of risks to the inflation outlook, and a growing literature measures these risks directly, by estimating how the tails of the inflation distribution respond to economic conditions (López-Salido and Loria, 2024; Banerjee et al., 2024). This *Inflation-at-Risk* approach, which builds on the *Growth-at-Risk* framework of Adrian et al. (2019), tells us how much inflation risk there is at a given point in time. It is less informative about where this risk comes from. Knowing that the upper tail of inflation is elevated, for instance, does not tell us whether this reflects pressure in the labor market, disruptions to supply, or a drift in expectations, and the appropriate policy response may differ across these sources.

This paper proposes a way to attribute inflation tail risk to a specific source. I introduce *Conditional Inflation-at-Risk* (CoIaR), defined as a quantile of future inflation conditional on a given economic variable being in a particular state, and the associated *contribution to inflation risk*, ΔCoIaR, the change in that quantile when the variable moves from its median to one of its tails. The measure adapts the conditional value-at-risk methodology of Adrian and Brunnermeier (2016), developed to quantify the contribution of individual financial institutions to systemic risk, to the measurement of inflation risk. It holds the inflation quantile fixed and varies only the state of the conditioning variable, so that it isolates the component of inflation tail risk that co-moves with that variable. The concept is general, as one can condition on financial conditions, energy prices or supply disruptions in the same way.

I apply the framework to the labor market, which is the natural starting point from the perspective of the Phillips curve. I measure labor market conditions by tightness, the ratio of job vacancies to unemployment, which has proven to be a better predictor of inflation than the unemployment rate or the output gap (Barnichon and Shapiro, 2022) and featured prominently in accounts of the post-pandemic inflation surge (Ball et al., 2022; Benigno and Eggertsson, 2023). The question I ask is simple: *how would the tails of future inflation change if the labor market were unusually slack or unusually tight today?*

Estimation proceeds in two steps. I first estimate the conditional quantiles of labor market tightness, which I call *Tightness-at-Risk*, and then quantile regressions of future inflation on tightness, controlling for past inflation, inflation expectations, and supply

shocks. The resulting *labor market contribution to inflation risk* decomposes into the product of a *tail sensitivity*, the slope of a quantile Phillips curve that measures how strongly a given quantile of inflation responds to tightness, and a measure of *labor market risk*, the distance between a tail and the median of the conditional distribution of tightness. Both regressions are estimated with the smoothed quantile regression estimator of Fernandes et al. (2021), and confidence bands for ΔCoIaR are obtained with a block bootstrap that accounts for estimation uncertainty in both steps.

Using monthly U.S. data for headline and core CPI inflation from 1978 to 2019, I obtain three main findings. First, labor market tightness matters for the distribution of future core inflation, but not uniformly across it. The tail sensitivity is positive and significant in the lower tail and at the median at horizons of 3, 6, and 12 months, and it increases with the horizon, while in the upper tail it is small and never significant. Second, the labor market contribution to inflation risk is economically meaningful and asymmetric. When the labor market is in its slack tail rather than at its median, the 10th percentile of core inflation is lower by about 0.25 percentage points at a horizon of 3 months and by about 0.55 percentage points at a horizon of 12 months, and this effect is statistically different from zero over almost the entire sample. When the labor market is in its tight tail, by contrast, the 90th percentile of core inflation is essentially unchanged. Third, the tails of headline inflation are shaped mainly by other sources. Labor market slack lowers the lower tail of headline inflation at short horizons, but the distribution of headline inflation is dominated by inflation expectations and supply shocks, and upside risk is governed by expectations and the persistence of past inflation rather than by the labor market. These results are robust to extending the sample to include the pandemic period, which adds some evidence that very tight labor markets raise upside inflation risk at intermediate horizons, but does not alter the main pattern.

Taken together, these results suggest that the labor market is primarily a source of *downside* inflation risk. A slack labor market raises the probability that inflation, and core inflation in particular, falls well below target, while the risk of inflation overshooting depends mainly on inflation dynamics and expectations. This asymmetry is not visible in conditional-mean Phillips curves, which summarize the relationship between tightness and inflation with a single slope. It also has an implication for policy: labor market slack is informative about the risk of undershooting the inflation target, but provides limited warning of upside inflation risk at the one-year horizon.

Connections to the literature. The paper contributes to three strands of literature. First, it adds to the Inflation-at-Risk literature (López-Salido and Loria, 2024; Banerjee et al., 2024) by moving from the question of how much inflation risk there is to the question of how much of it can be attributed to a specific source, and by proposing a measure that does so in real time. Second, it contributes to the literature on the Phillips curve by documenting that the relationship between labor market tightness and inflation differs across the inflation distribution, and that it is concentrated in its lower part. Third, it shows how the CoVaR approach of Adrian and Brunnermeier (2016), developed for financial stability, can be carried over to macroeconomic risk measurement, where conditioning on a stressed state of one variable provides a transparent way of quantifying its contribution to the tail risk of another.

Outline of the paper. The rest of the paper is organized as follows. Section 2 introduces CoIaR and describes the estimation and inference strategy. Section 3 presents the data and the results. Section 4 concludes.

2 Conditional Inflation-at-Risk

To measure how much of inflation tail risk is associated with a specific source of risk, I build on the conditional value-at-risk (CoVaR) methodology of Adrian and Brunnermeier (2016). CoVaR was developed to quantify how distress at one financial institution affects the tails of the financial system as a whole. Here, the same logic is applied to inflation. The question is how the tails of future inflation change when a given economic variable is under stress. Throughout, all distributional quantiles are conditional on the information set $\mathcal{I}_t$, and I suppress this dependence for notational simplicity.

The starting point is the standard notion of inflation-at-risk.

Definition 1. *Let π_{t+h} denote the year-on-year inflation rate in period $t+h$, that is, inflation over the twelve months ending in $t+h$.* ***Inflation-at-Risk*** *(IaR) at level $\tau \in (0,1)$ is the value $IaR_{\tau,t}^{(h)}$ such that*

$$\Pr\left[\pi_{t+h} \leq IaR_{\tau,t}^{(h)}\right] = \tau$$

IaR is simply the τ-quantile of the distribution of future inflation, with low values of τ describing downside risk and high values describing upside risk. It summarizes the level of inflation risk implied by all available information, but it says nothing

about where that risk comes from. To isolate the contribution of a particular variable z_t, I condition on its state.

Definition 2. *Let $\mathcal{C}(z_t)$ denote an event concerning the value of z_t.* ***Conditional Inflation-at-Risk*** *(CoIaR) at level $\tau \in (0,1)$ is the value $CoIaR_{\tau,t}^{(h)}$ such that*

$$\Pr\left[\pi_{t+h} \leq CoIaR_{\tau,t}^{(h)} \mid \mathcal{C}(z_t)\right] = \tau$$

Whereas IaR answers the question "*what is the τ-quantile of future inflation given current information?*", CoIaR answers "*what would that quantile be if z_t were in a particular state?*". A natural way to describe the state of z_t is through its own distribution, which leads to the third concept.

Definition 3. *Let $\mathcal{C}_q(z_t)$ denote the event that z_t equals its conditional q-quantile, $q \in (0,1)$. The* ***contribution to inflation risk*** *($\Delta CoIaR$) of z_t at level $\tau \in (0,1)$ is the difference between CoIaR when z_t is in its q-quantile state and CoIaR when z_t is at its median:*

$$\Delta CoIaR_{\tau,q,t}^{(h)} = CoIaR_{\tau,t}^{(h)}\big|_{\mathcal{C}_q(z_t)} - CoIaR_{\tau,t}^{(h)}\big|_{\mathcal{C}_{0.5}(z_t)} \tag{1}$$

In this comparison, neither the inflation quantile τ nor the date t changes; only the state of z_t does. ΔCoIaR therefore captures the part of inflation tail risk that comoves with z_t. If z_t were to move into its tail today, by how much would the corresponding tail of inflation shift over the next h periods? As for CoVaR, is statistical. It measures tail comovement between z_t and future inflation, not a causal effect of z_t. Nothing in this definition is specific to a particular variable. One could condition on financial conditions, on energy prices or on supply disruptions in exactly the same way.

In this paper, I condition on the labor market. The choice follows from the Phillips curve. Labor market slack is the main domestic driver of inflation, and when the labor market tightens, firms compete for workers, wages rise and price pressures build up. I measure the state of the labor market by tightness, $\theta_t = v_t/u_t$, the ratio of the vacancy rate to the unemployment rate, which has proven to be a better predictor of inflation than the unemployment rate or the output gap (Barnichon and Shapiro, 2022) and played a central role in the post-pandemic inflation surge (Ball et al., 2022; Benigno and Eggertsson, 2023). Since tightness rises with labor market pressure, states with $q > 0.5$ correspond to a tighter-than-usual labor market, and states with $q < 0.5$ to a slacker-than-usual one.

Conditioning on the labor market makes the question that motivates the paper measurable. *How would the tails of future inflation change if the labor market were unusually slack or unusually tight today?* The contribution to inflation risk answers it by

comparing a given quantile of future inflation when the labor market is in one of its tails with the same quantile when the labor market is at its median. As shown below, the answer combines two elements, namely how far the labor market can plausibly be from its median given current conditions, and how sensitive the relevant quantile of inflation is to the state of the labor market. Note also that the labor market need not affect all parts of the inflation distribution equally. If tightness shifted every quantile of future inflation by the same amount, the conditional mean would tell the whole story. If instead it also affects the dispersion or the asymmetry of inflation, for example because shocks pass through to prices more strongly when the labor market is tight, then the tails will respond differently from the center. The framework allows for this possibility without imposing it.

Estimation. The measure can be estimated with two quantile regressions. The first describes the distribution of labor market tightness,

$$\ln \theta_t = \kappa(q) + \psi(q)' x_{t-} + \varepsilon_t^{\theta}(q), \tag{2}$$

and the second the distribution of future inflation,

$$\pi_{t+h} = \alpha(\tau) + \beta(\tau) \ln \theta_t + \phi(\tau)' x_{t-} + \varepsilon_{t+h}^{\pi}(\tau) \tag{3}$$

where x_{t-} is a vector of predetermined state variables. I use the same controls in both equations, so that the two steps are consistent with each other and the resulting measure captures the association between the labor market and inflation tail risk holding the broader macroeconomic environment fixed. I interpret this association as conditional tail comovement rather than a causal effect. Equation (3) can be read as a quantile Phillips curve, in which $\beta(\tau)$ is the slope of the Phillips curve at the τ-quantile of the inflation distribution. A constant $\beta(\tau)$ means that tightness shifts the whole distribution in parallel; a $\beta(\tau)$ that varies with τ means that tightness also changes its shape. Both equations are estimated with the smoothed quantile regression estimator of Fernandes et al. (2021), which replaces the standard quantile regression objective with a smoothed counterpart and improves finite-sample accuracy and numerical stability, particularly with a relatively rich set of covariates.

The fitted values of (2) give the conditional quantiles of labor market tightness, which I refer to as *Tightness-at-Risk* (TaR):

$$\text{TaR}_{q,t} = \hat{\kappa}(q) + \hat{\psi}(q)' x_{t-} \tag{4}$$

Evaluated at the observed level of tightness, the fitted values of (3) give inflation-at-risk. Evaluated instead at $\ln \theta_t = \text{TaR}_{q,t}$, they give the level of inflation-at-risk that would prevail if the labor market were in its q-quantile state, that is, conditional inflation-at-risk:

$$\text{CoIaR}^{(h)}_{\tau,q,t} = \hat{\alpha}(\tau) + \hat{\beta}(\tau)\text{TaR}_{q,t} + \hat{\phi}(\tau)' x_{t-} \tag{5}$$

and comparing CoIaR across labor market states gives the labor market contribution to inflation risk:

$$\begin{aligned} \Delta\text{CoIaR}^{(h)}_{\tau,q,t} &= \text{CoIaR}^{(h)}_{\tau,t}\Big|_{\ln\theta_t=\text{TaR}_{q,t}} - \text{CoIaR}^{(h)}_{\tau,t}\Big|_{\ln\theta_t=\text{TaR}_{0.5,t}} \\ &= \underbrace{\hat{\beta}(\tau)}_{\substack{\text{tail}\\ \text{sensitivity}}} \times \underbrace{\left(\text{TaR}_{q,t} - \text{TaR}_{0.5,t}\right)}_{\substack{\text{labor market}\\ \text{risk}}}. \end{aligned} \tag{6}$$

Expression (6) is the empirical counterpart of the question above. The second term is labor market risk, the distance between the q-quantile of tightness and its median given current conditions, which measures how far the labor market can plausibly be from its normal state. The first term measures how strongly the τ-quantile of future inflation responds to that distance. Labor market conditions therefore matter for inflation tail risk when the distribution of tightness is wide or skewed, when the relevant quantile of inflation is sensitive to tightness, or both. Since $\hat{\beta}(\tau)$ does not change over time, all time variation in ΔCoIaR comes from labor market risk,

$$\text{TaR}_{q,t} - \text{TaR}_{0.5,t} = \left[\hat{\kappa}(q) - \hat{\kappa}(0.5)\right] + \left[\hat{\psi}(q) - \hat{\psi}(0.5)\right]' x_{t-} \tag{7}$$

so a state variable moves the measure only if its coefficient differs between the q-quantile and the median of tightness, that is, only if it changes the dispersion or asymmetry of the labor market distribution rather than its location.

Since ΔCoIaR combines estimates from both regressions, its sampling uncertainty reflects estimation error in $\widehat{\beta}(\tau)$ as well as in $\text{TaR}_{q,t}$. I construct confidence bands with a circular block bootstrap (Politis and Romano, 1992). Blocks of consecutive months of $(\ln\theta_t, x_t, \pi_{t+h})$ are resampled, both regressions are re-estimated on each bootstrap sample, and $\Delta\text{CoIaR}^{(h)}_{\tau,q,t}$ is recomputed at the observed values of x_t. Resampling blocks rather than single observations preserves the serial dependence in the data, including the dependence induced by the overlap of π_{t+h} across consecutive months. Pointwise bands are given by the percentiles of the bootstrap distribution at each date.

3 Results

I estimate (2) and (3) and construct (6) using monthly U.S. data for both headline and core CPI inflation. The dependent variable is year-on-year inflation h months ahead. Labor market tightness is the ratio of the vacancy rate to the unemployment rate. The vacancy rate is the number of job openings divided by the labor force. Until December 2000, I use the series of Barnichon (2010), who converts his composite help-wanted index into units of the job openings rate. From January 2001, I use the series of the Automated Business Cycle Dashboard of Michaillat and Saez (2024), which is based on job openings from JOLTS.[1] As control variables I use the past average inflation, defined as the average of the monthly inflation rates over the previous 12 months, one-year-ahead inflation expectations from the University of Michigan survey, and a monthly version of the supply shocks measure of Benigno and Eggertsson (2023), whose construction is desribed in Appendix B. The sample starts in January 1978, when the Michigan expectations series becomes available, and the baseline estimates end in December 2019 so that the results are not driven by the pandemic period.[2]

Table 1 reports the estimates of the inflation quantile regression (5) at the 10th, 50th and 90th percentiles, for horizons of 3, 6 and 12 months. For headline inflation, the coefficient on labor market tightness is positive and significant at the 1 percent level in the lower tail and at the median at horizons of 3 and 6 months. At the 10th percentile it equals 0.70 at $h = 3$ and 0.72 at $h = 6$. In the upper tail, by contrast, the coefficient is small and not significant. At 12 months, the coefficient on tightness remains significant at the median but only marginally so in the tails, and the distribution of headline inflation is dominated by inflation expectations, whose coefficient at the 90th percentile reaches 1.52. Supply shocks also matter for headline inflation, with positive and mostly significant coefficients at horizons of 3 and 6 months. Core inflation displays the same asymmetry, and more sharply. The coefficient on tightness is significant at the 1 percent level in the lower tail and at the median at every horizon, and it increases with the horizon. At the 10th percentile it rises from 0.35 at $h = 3$ to 0.72 at $h = 12$. In the upper tail it is small and never significant. Tightness therefore shifts the lower part of the distribution of future core inflation but has little effect on its upper tail, which is instead driven by the persistence of past inflation and by expected inflation, both of which have their largest coefficients at the 90th percentile. Supply

[1] The dashboard is available at https://pascalmichaillat.org/dashboard/.

[2] Table A1 lists all data sources.

shocks, unlike for headline inflation, play little role for core inflation, consistent with the fact that the supply-shock index captures mainly movements in food, energy and import prices. Overall, the evidence points to an asymmetric role of the labor market, which is informative about downside risks to inflation, and particularly to core inflation, while upside risks are governed by inflation persistence and expectations.

Table 1: Inflation quantile regressions across horizons

	$h=3$			$h=6$			$h=12$		
	$\tau=0.1$	$\tau=0.5$	$\tau=0.9$	$\tau=0.1$	$\tau=0.5$	$\tau=0.9$	$\tau=0.1$	$\tau=0.5$	$\tau=0.9$
Panel A: Headline CPI									
$\ln(v_t/u_t)$	0.703^{***}	0.370^{***}	-0.142	0.717^{***}	0.596^{***}	0.119	0.407^{*}	0.576^{***}	0.309^{*}
	(0.089)	(0.085)	(0.128)	(0.162)	(0.113)	(0.148)	(0.208)	(0.153)	(0.184)
Past average inflation	0.744^{***}	0.711^{***}	0.569^{***}	0.531^{***}	0.465^{***}	0.294^{***}	0.318^{**}	0.165^{**}	0.133
	(0.044)	(0.042)	(0.081)	(0.089)	(0.046)	(0.062)	(0.134)	(0.073)	(0.119)
Supply shock	0.188^{*}	0.143^{***}	0.176^{**}	0.245^{*}	0.216^{***}	0.119^{*}	-0.035	0.047	-0.132^{**}
	(0.110)	(0.037)	(0.087)	(0.143)	(0.055)	(0.072)	(0.057)	(0.069)	(0.067)
Expected inflation	0.267^{***}	0.509^{***}	0.794^{***}	0.483^{***}	0.814^{***}	1.170^{***}	0.247	0.949^{***}	1.518^{***}
	(0.097)	(0.050)	(0.117)	(0.150)	(0.083)	(0.078)	(0.362)	(0.154)	(0.161)
Constant	-0.549^{***}	-0.498^{***}	-0.388^{**}	-0.896^{***}	-0.607^{***}	-0.346^{*}	-0.165	-0.162	-0.310
	(0.210)	(0.093)	(0.184)	(0.265)	(0.179)	(0.187)	(0.670)	(0.276)	(0.357)
Observations		501			498			492	
Panel B: Core CPI									
$\ln(v_t/u_t)$	0.347^{***}	0.264^{***}	0.005	0.506^{***}	0.367^{***}	0.112	0.717^{***}	0.523^{***}	-0.149
	(0.042)	(0.032)	(0.058)	(0.059)	(0.062)	(0.078)	(0.057)	(0.071)	(0.123)
Past average inflation	0.800^{***}	0.883^{***}	0.949^{***}	0.661^{***}	0.763^{***}	0.836^{***}	0.374^{***}	0.611^{***}	0.815^{***}
	(0.034)	(0.017)	(0.030)	(0.082)	(0.035)	(0.031)	(0.055)	(0.044)	(0.057)
Supply shock	-0.005	-0.032^{*}	-0.041	-0.014	-0.010	-0.056^{**}	-0.031	-0.020	0.011
	(0.014)	(0.018)	(0.033)	(0.021)	(0.031)	(0.027)	(0.027)	(0.038)	(0.033)
Expected inflation	0.138^{**}	0.270^{***}	0.335^{***}	0.240^{**}	0.421^{***}	0.572^{***}	0.312^{*}	0.493^{***}	0.733^{***}
	(0.059)	(0.036)	(0.036)	(0.119)	(0.070)	(0.041)	(0.161)	(0.078)	(0.071)
Constant	-0.049	-0.343^{***}	-0.457^{***}	-0.046	-0.429^{***}	-0.669^{***}	0.343	-0.194	-0.968^{***}
	(0.143)	(0.090)	(0.066)	(0.167)	(0.166)	(0.083)	(0.353)	(0.179)	(0.118)
Observations		501			498			492	

Notes: Smoothed quantile regressions (Fernandes et al., 2021) of year-on-year inflation at $t+h$ on log labor market tightness and three controls: average inflation over the previous 12 months (same price index), the supply-shock index at t (Appendix B) and 1-year expected inflation at t from the University of Michigan survey. Sample: 1978:01–2018:12 for $h=12$ (the end date varies with h so that no data after December 2019 are used). Asymptotic standard errors of the smoothed quantile regression estimator in parentheses. ***, **, * denote significance at the 1, 5 and 10 percent levels.

Figures 1 and 2 plot the labor market contribution to inflation risk over time for headline and core inflation. The left panels show downside risk, the change in the 10th percentile of future inflation when the labor market moves from its median state to its slack tail. The right panels show upside risk, the change in the 90th percentile

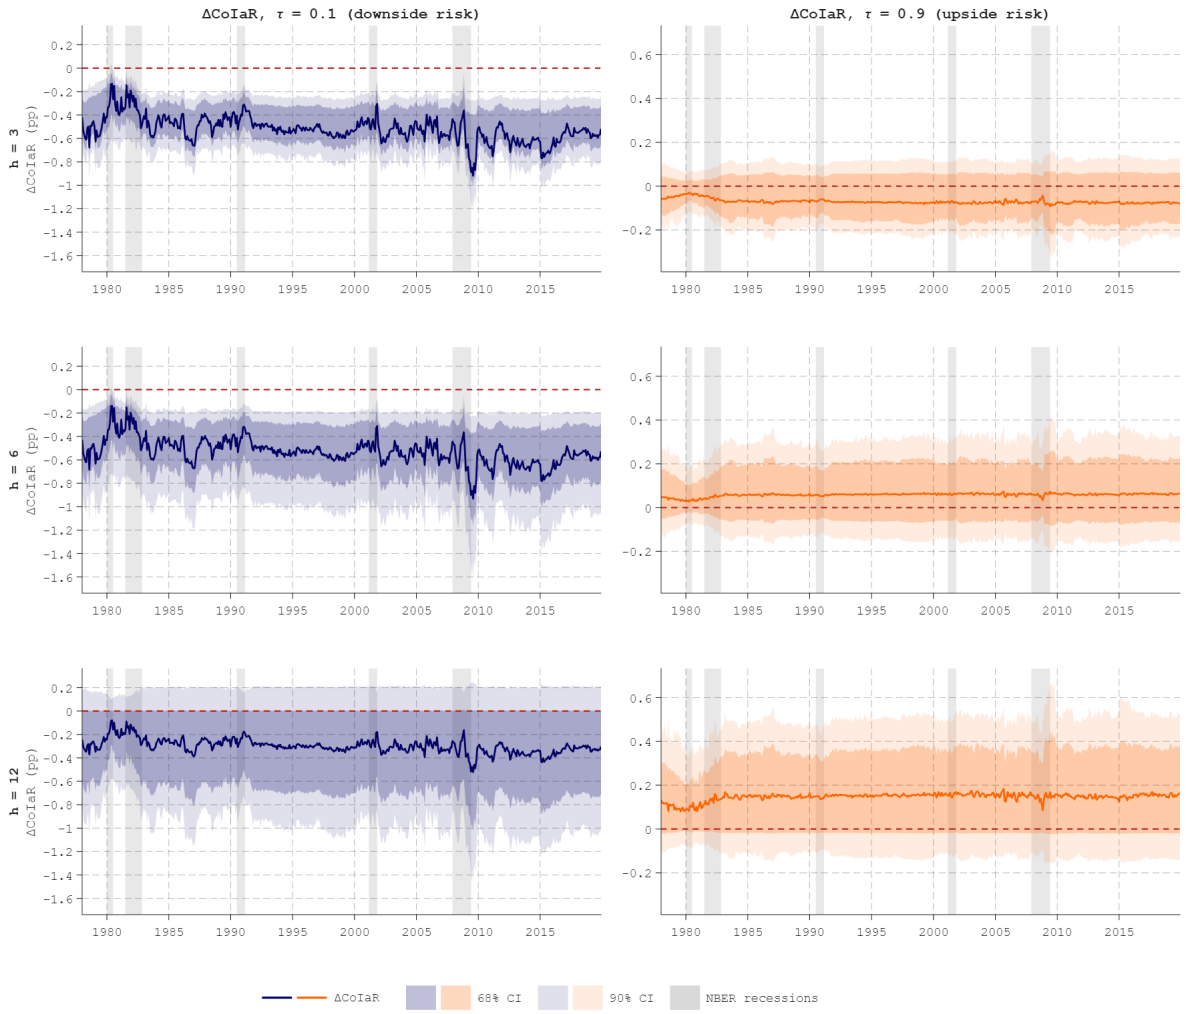

Figure 1: Labor market contribution to headline inflation risk

Notes: $\Delta\text{CoIaR}^{(h)}_{\tau,q,t}$ for year-on-year headline CPI inflation at horizons $h = 3, 6, 12$ months. Left panels: downside risk, the change in the 10th percentile of future inflation when labor market tightness moves from its median to its 10th percentile ($\tau = q = 0.1$). Right panels: upside risk, the change in the 90th percentile when tightness moves to its 90th percentile ($\tau = q = 0.9$). Shaded areas are pointwise 68 and 90 percent confidence bands from a circular block bootstrap with blocks of 12 months and 500 replications. Grey bars denote NBER recessions. Sample: 1978:01–2019:12.

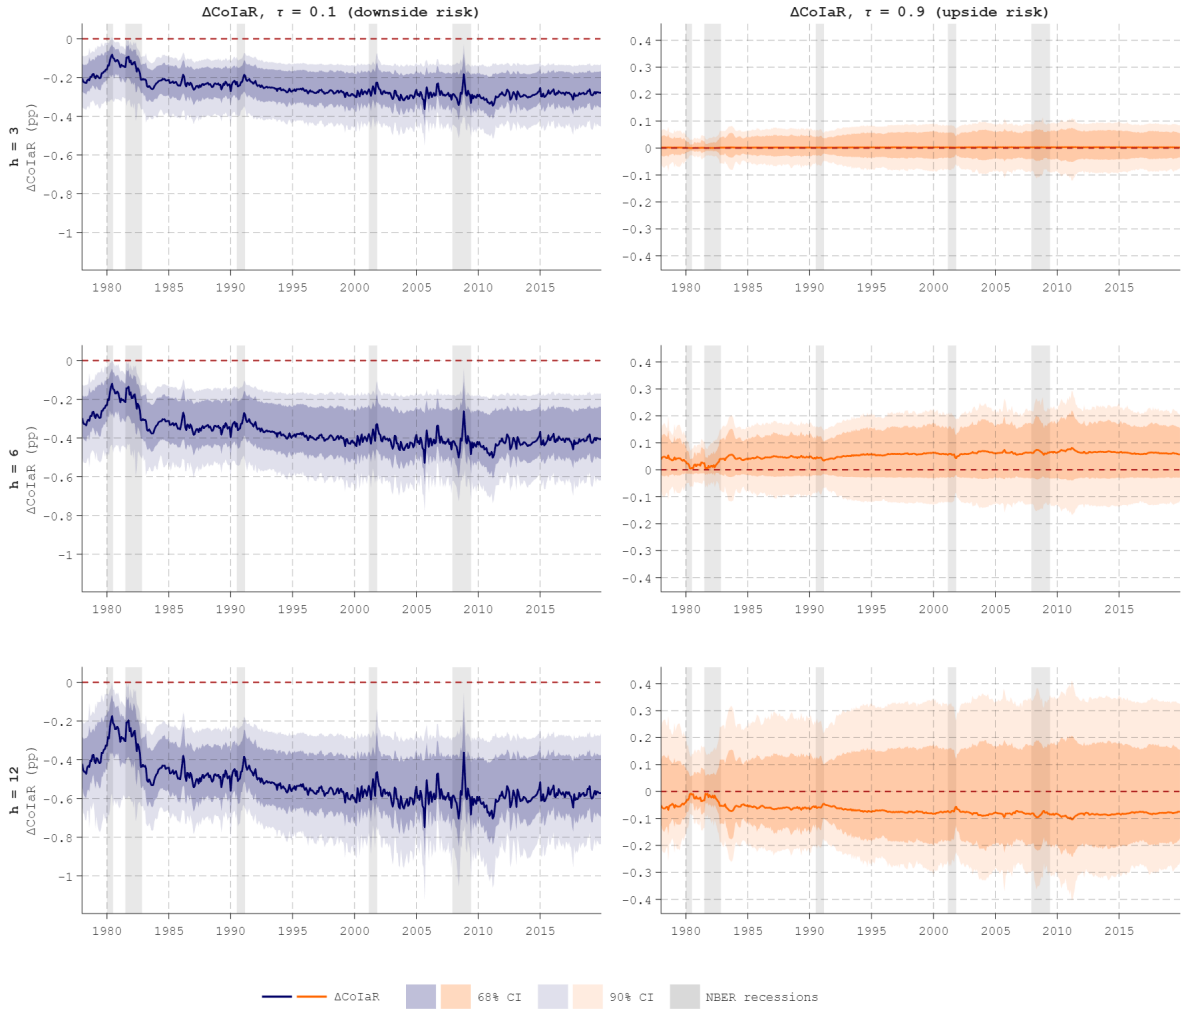

Figure 2: Labor market contribution to core inflation risk

Notes: $\Delta\text{CoIaR}^{(h)}_{\tau,q,t}$ for year-on-year core CPI inflation at horizons $h = 3, 6, 12$ months. Left panels: downside risk, the change in the 10th percentile of future inflation when labor market tightness moves from its median to its 10th percentile ($\tau = q = 0.1$). Right panels: upside risk, the change in the 90th percentile when tightness moves to its 90th percentile ($\tau = q = 0.9$). Shaded areas are pointwise 68 and 90 percent confidence bands from a circular block bootstrap with blocks of 12 months and 500 replications. Grey bars denote NBER recessions. Sample: 1978:01–2019:12.

when the labor market moves to its tight tail. Each row corresponds to a different horizon.

The figures confirm the asymmetry documented in Table 1. For headline inflation, downside risk associated with the labor market is sizable at short horizons. A slack

labor market lowers the 10th percentile of future headline inflation by about 0.5 percentage points at $h = 3$ and $h = 6$, and the 90 percent confidence bands exclude zero over almost the entire sample. At $h = 12$, downside risk falls to about 0.3 percentage points and is less precisely estimated, with the 68 percent bands excluding zero but the 90 percent bands including it. Upside risk is small and not significant at horizons of 3 and 6 months. At $h = 12$, a tight labor market raises the 90th percentile of future headline inflation by about 0.15 percentage points, but the effect is at most marginally significant, since the lower edge of the 68 percent band lies close to zero and the 90 percent band includes it throughout. For core inflation, downside risk is smaller at short horizons but more precisely estimated, and it grows with the horizon. A slack labor market lowers the 10th percentile of future core inflation by around 0.25 percentage points at $h = 3$, 0.4 percentage points at $h = 6$ and 0.55 percentage points at $h = 12$, and the confidence bands exclude zero over almost the entire sample. Upside risk, by contrast, is close to zero at every horizon, and the bands always include zero. A tight labor market does not raise the 90th percentile of future core inflation.

The labor market contribution to inflation risk also varies over time. Since the tail sensitivity is constant, this variation comes entirely from tightness risk, the distance between a tail and the median of the conditional distribution of tightness, which I report in Appendix C. Because this distribution is conditioned on the same state variables as the inflation regression, tightness risk moves with inflation conditions, and mostly with past inflation. When recent inflation is high, as in the early 1980s, the slack tail of tightness lies closer to the median. Downside tightness risk is then small, and so is the labor market contribution to downside inflation risk. When recent inflation falls, the slack tail moves away from the median, downside tightness risk widens, and the contribution becomes larger in absolute value. For headline inflation, this is visible around the Great Recession and in 2015, when the collapse of energy prices lowered past inflation. For core inflation, whose past inflation has been more stable, downside tightness risk, and hence the contribution, varies less after the mid-1980s.

4 Conclusion

This paper has introduced Conditional Inflation-at-Risk, a measure of how the tails of the future inflation distribution change when a given economic variable is in a stressed state. By adapting the CoVaR methodology of Adrian and Brunnermeier (2016) to inflation, the measure moves the analysis of inflation risk from the question of how

much risk there is to the question of where it comes from. It holds the inflation quantile fixed and varies only the state of the conditioning variable, and it decomposes the resulting contribution to inflation risk into two interpretable components, the sensitivity of a given quantile of inflation to the conditioning variable and the risk in that variable itself.

Applied to the labor market, the framework delivers a clear answer to the question that motivates the paper. An unusually slack labor market is associated with a sizable shift in the lower tail of future inflation, while an unusually tight labor market leaves the upper tail largely unchanged. For headline inflation, a slack labor market lowers the 10th percentile of future inflation by about half a percentage point at horizons of 3 and 6 months, but the effect weakens at one year, when the tails of headline inflation are shaped mainly by inflation expectations. The contribution of a tight labor market to upside headline inflation risk is small, and at most marginally significant at the one-year horizon. For core inflation, the labor market contribution to downside inflation risk is smaller at short horizons but more precisely estimated, and it strengthens with the horizon, reaching about half a percentage point at one year. The upper tail of core inflation, by contrast, is unaffected by labor market conditions and is driven instead by inflation persistence and expectations. Extending the sample to include the pandemic does not alter this pattern, although it suggests that very tight labor markets can also raise upside risks at intermediate horizons.

The labor market contribution to inflation risk also varies over time. Because the tail sensitivity is constant, this variation reflects tightness risk, which in turn depends on inflation conditions. When past inflation is low, the slack tail of tightness lies further from the median, and the labor market contributes more to downside inflation risk.

These findings have implications for how central banks read the labor market when assessing the balance of risks. Measures of labor market slack are informative about downside risks to inflation, and can help gauge how far inflation might fall below target in a weakening economy. They are a less reliable signal of upside risk, which calls for close monitoring of inflation dynamics and inflation expectations. More generally, the asymmetry documented here is invisible in conditional-mean Phillips curves, which summarize the relationship between labor market conditions and inflation with a single slope, and highlights the value of a distributional approach.

Several extensions are left for future research. The framework is general, and the

same approach can be used to quantify the contribution of other sources of inflation risk, such as financial conditions, energy prices or global supply disruptions, and to compare them within a common metric. The conditional distribution of labor market tightness is modeled here as a function of the same state variables that enter the inflation regression, which ensures consistency between the two steps. Richer models of tightness risk, for instance including the dynamics of tightness itself, would allow the measure to capture labor market developments more fully. The relationship between tightness and inflation may also be state dependent, as suggested by recent evidence of a steeper Phillips curve when the labor market is very tight (Benigno and Eggertsson, 2023). As more observations from such periods become available, the framework can be extended to allow the tail sensitivity to depend on the state of the labor market. Finally, applying the measure to other economies would show whether the asymmetric role of the labor market is a general feature of inflation risk or specific to the United States.

References

Adrian, T., Boyarchenko, N., and Giannone, D. (2019). Vulnerable growth. *American Economic Review*, 109(4):1263–1289.

Adrian, T. and Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7):1705–1741.

Ball, L., Leigh, D., and Mishra, P. (2022). Understanding US inflation during the COVID era. *Brookings Papers on Economic Activity*, pages 1–54.

Banerjee, R., Contreras, J., Mehrotra, A., and Zampolli, F. (2024). Inflation at risk in advanced and emerging market economies. *Journal of International Money and Finance*, 142:103025.

Barnichon, R. (2010). Building a composite Help-Wanted Index. *Economics Letters*, 109(3):175–178.

Barnichon, R. and Shapiro, A. H. (2022). What's the best measure of economic slack? Economic Letter 2022-04, Federal Reserve Bank of San Francisco.

Benigno, P. and Eggertsson, G. (2023). It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve. Technical report, National Bureau of Economic Research.

Fernandes, M., Guerre, E., and Horta, E. (2021). Smoothing quantile regressions. *Journal of Business and Economic Statistics*, 39(1):338–357.

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

Michaillat, P. and Saez, E. (2024). $u^* = \sqrt{uv}$: The full-employment rate of unemployment in the united states. *Brookings Papers on Economic Activity*, pages 323–390. Fall 2024.

Politis, D. N. and Romano, J. P. (1992). A circular block-resampling procedure for stationary data. In LePage, R. and Billard, L., editors, *Exploring the Limits of Bootstrap*, pages 263–270. Wiley, New York.

Online Appendix

Conditional Inflation-at-Risk

Markeljan Blacëri
University of Padova

November 2026

Note to Readers: This online appendix is intended to complement the main paper and is not self-contained. Readers are encouraged to consult the main text for full context, methodological framework, and empirical results.

Contents

A Data **17**

B A monthly measure of supply shocks **18**

C Tightness risk **22**

D Results including the pandemic period **24**

A Data

Table A1: Data sources

Variable	Description and transformation	Source	Link
Headline CPI	Consumer Price Index for All Urban Consumers, all items, U.S. city average, index 1982–84 = 100, seasonally adjusted, monthly; inflation between t and $t+h$ computed as annualized log change	BLS via FRED	Link
Core CPI	Consumer Price Index for All Urban Consumers, all items less food and energy, index 1982–84 = 100, seasonally adjusted, monthly; annualized log change	BLS via FRED	Link
Headline PCE	Personal consumption expenditures chain-type price index, index 2017 = 100, seasonally adjusted, monthly; annualized log change	BEA via FRED	Link
Core PCE	Personal consumption expenditures excluding food and energy, chain-type price index, index 2017 = 100, seasonally adjusted, monthly; annualized log change	BEA via FRED	Link
Unemployment rate	Civilian unemployment rate, seasonally adjusted, percent, monthly	BLS via FRED	Link
Vacancy rate	Job vacancies as a percent of the labor force, monthly. Composite help-wanted index of Barnichon (2010) until December 2000; JOLTS job openings divided by the labor force (employed plus unemployed), shifted forward one month to align with the CPS, thereafter	Barnichon (2010); Michaillat and Saez (2024)	Link
Labor market tightness	Ratio of the vacancy rate to the unemployment rate, $\theta_t = v_t/u_t$; enters the regressions in logs, $\ln \theta_t$	Own calculation	–

Table A1 continued

Variable	Description and transformation	Source	Link
Expected inflation, Michigan	University of Michigan survey of consumers, median expected price change over the next 12 months, percent, monthly, from January 1978	University of Michigan via FRED	Link
Import deflator	Imports of goods and services, implicit price deflator, index 2017 = 100, seasonally adjusted, quarterly; annualized log change	BEA via FRED	Link
GDP deflator	Gross domestic product, implicit price deflator, index 2017 = 100, seasonally adjusted, quarterly; annualized log change	BEA via FRED	Link
Supply-shock index	Common factor of the CPI gap (headline minus core CPI inflation), the PCE gap (headline minus core PCE inflation) and the import gap (import deflator minus GDP deflator inflation), estimated with a mixed-frequency dynamic factor model; filtered estimate, standardized. See Appendix B	Own calculation	–

B A monthly measure of supply shocks

This appendix describes the construction of the supply-shock index used as a control in the estimation of CoIaR. The index follows Benigno and Eggertsson (2023), who work with quarterly data and summarize supply disturbances by the first principal component of three proxies. Two of these proxies, based on consumer price indices, are also available at monthly frequency. The third is based on national-accounts deflators, which are published only quarterly. Since the rest of my analysis is monthly, I extract the common component of the three proxies at monthly frequency with a mixed-frequency dynamic factor model in the spirit of Mariano and Murasawa (2002), which combines the two monthly proxies with the quarterly one.

Let Δ denote the first difference and let inflation rates be annualized log changes. The three proxies are:

1. the *CPI gap*, the difference between headline and core CPI inflation,

$$g_t^{CPI} = 1200 \left(\Delta \ln CPI_t^{headline} - \Delta \ln CPI_t^{core} \right)$$

   observed monthly,
2. the *PCE gap*, defined analogously for the PCE price index,

$$g_t^{PCE} = 1200 \left(\Delta \ln PCE_t^{headline} - \Delta \ln PCE_t^{core} \right)$$

   observed monthly,
3. the *import gap*, the difference between the inflation rates of the implicit deflator of imports of goods and services and of the GDP deflator,

$$g_s^{M} = 400 \left(\Delta \ln P_s^{M} - \Delta \ln P_s^{GDP} \right)$$

   observed quarterly, where s indexes quarters.

The first two proxies capture shocks to food and energy prices relative to the rest of the consumption basket while the third captures shocks to the relative price of imported goods and services. The sample runs from January 1960 to December 2025. Headline and core CPI are not available for October 2025, when data collection was suspended during the federal government shutdown. Because the CPI gap is a monthly log change, the missing price level removes two observations, g_t^{CPI} for October and for November 2025, which are treated as missing. The state-space framework handles this without imputation: in those two months the factor is updated using the PCE gap, which is available throughout, and, in October, the quarterly import gap for 2025Q3. Each proxy is standardized over the sample before estimation.

Let $\mathcal{S}_t$ denote a latent monthly supply factor and $y_{i,t}$ the standardized proxies. The two monthly proxies load directly on the factor,

$$y_{i,t} = \lambda_i \mathcal{S}_t + e_{i,t}, \quad i \in \{CPI, PCE\} \tag{B1}$$

The quarterly proxy requires a temporal aggregation rule. Quarterly deflators are averages of the underlying monthly price levels, so that, approximating the log of an average by the average of logs, the log price level in a quarter ending in month t is $\ln \overline{P}_t \approx \frac{1}{3} (\ln P_t + \ln P_{t-1} + \ln P_{t-2})$. Its quarter-on-quarter change is therefore a triangular combination of the monthly log changes,

$$\ln \overline{P}_t - \ln \overline{P}_{t-3} \approx \frac{1}{3} (\Delta \ln P_t + 2\Delta \ln P_{t-1} + 3\Delta \ln P_{t-2} + 2\Delta \ln P_{t-3} + \Delta \ln P_{t-4})$$

Expressed in annualized rates, the quarterly growth rate is a weighted average of the annualized monthly growth rates with weights $w = (1, 2, 3, 2, 1)/9$. Following Mariano and Murasawa (2002), I apply the same aggregation to the factor:

$$y_{M,t} = \lambda_M \sum_{j=0}^{4} w_j \mathcal{S}_{t-\ell-j} + e_{M,t} \tag{B2}$$

The observation $y_{M,t}$ is available only once per quarter and is missing in the remaining months. The parameter ℓ is a publication lag. I assign the value for a quarter to the first month after the quarter ends ($\ell = 1$), approximately when the advance estimate of GDP is released, so that the quarterly proxy enters the index only in months in which it would have been available.[3] The factor follows an autoregressive process,

$$\mathcal{S}_t = \rho \mathcal{S}_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, 1), \tag{B3}$$

and the idiosyncratic components $e_{i,t} \sim \mathcal{N}(0, \sigma_i^2)$ are mutually independent and independent of η_t.

The scale of the factor is fixed by setting the variance of η_t to one, and its sign by requiring a positive loading on the CPI gap, so that a positive value of the index corresponds to an adverse (inflationary) supply shock. In addition, I restrict the idiosyncratic variances to be equal across the three standardized proxies, $\sigma_{CPI} = \sigma_{PCE} = \sigma$. This restriction serves two purposes. First, it ties the model to the principal-component approach of Benigno and Eggertsson (2023): in a static factor model with standardized data and isotropic noise, the maximum likelihood estimate of the factor coincides with the first principal component (Tipping and Bishop, 1999). The model in (B1)-(B3) is therefore the mixed-frequency, dynamic analogue of their measure. Second, with only three indicators, two of which are highly correlated[4], the unrestricted model produces a boundary solution in which the idiosyncratic variance of the PCE gap is zero and the factor collapses onto that single series. The restriction rules out this degenerate solution.

Equations (B1)-(B3) form a linear Gaussian state-space model with state vector $\alpha_t = (\mathcal{S}_t, \mathcal{S}_{t-1}, \ldots, \mathcal{S}_{t-5})'$. The parameters $(\rho, \lambda_{CPI}, \lambda_{PCE}, \lambda_M, \sigma)$ are estimated by maximum likelihood, evaluating the likelihood with the Kalman filter (Durbin and Koopman, 2012). Missing observations, namely the quarterly proxy in the months in which it is not released, the start of the PCE series in 1959 and the missing CPI release, are

[3] The index is built from the latest data vintage, so it does not account for data revisions.

[4] The correlation between the monthly CPI and PCE gaps is 0.88.

handled by dropping the corresponding rows of the measurement equation at each date. The filter is initialized at the unconditional distribution of the state.

Two estimates of the factor are available. The smoothed estimate $\mathcal{S}_{t|T}$ uses the full sample and is the natural object to describe the history of supply shocks. The filtered estimate $\mathcal{S}_{t|t}$ uses only observations available up to month t.

Table B2 reports the estimates. The factor is moderately persistent at monthly frequency ($\hat{\rho} = 0.37$), consistent with supply shocks that are largely transitory month to month. All three proxies load positively on the factor. The larger loading on the import gap reflects temporal aggregation rather than a stronger association: the quarterly proxy depends on a weighted average of five monthly realizations of the factor, whose variance is smaller than that of the factor itself, so a larger loading is needed to match the variance of the standardized proxy. Taking this into account, the factor accounts for a similar share of the variance of each proxy: 82% for the CPI gap and the PCE gap, and 80% for the import gap.

Table B2: Mixed-frequency factor model: parameter estimates

Parameter	Estimate
Persistence, ρ	0.369
Loading, CPI gap, λ_{CPI}	0.844
Loading, PCE gap, λ_{PCE}	0.857
Loading, import gap, λ_M	1.210
Idiosyncratic s.d., σ	0.426
Log-likelihood	-1946.2
Months	792

Notes: Maximum likelihood estimates of (B1)–(B3), January 1960 to December 2025. Proxies are standardized; the variance of the factor innovation is normalized to one and idiosyncratic variances are restricted to be equal.

To verify that the monthly index reproduces the measure of Benigno and Eggertsson (2023), I construct their quarterly index directly. I compute the three proxies at quarterly frequency (aggregating the monthly price indices to quarterly averages), standardize them and extract the first principal component, which explains 87% of their joint variance. I then aggregate the monthly factor to the same quarterly concept using the weights w in (B2). The correlation between the quarterly principal component and the aggregated factor is 0.998 for the smoothed estimate and 0.975 for the

filtered estimate. The monthly index therefore preserves the information in the quarterly measure of Benigno and Eggertsson (2023), while adding variation within the quarter and respecting the timing of data releases. The index captures the main supply disturbances of the period, including the oil shocks of 1973–74 and 1979–80, the oil price collapse of 1986, the 1990 Gulf War, the energy price swings of 2005 and 2008, and the post-pandemic episode of 2021–2022.

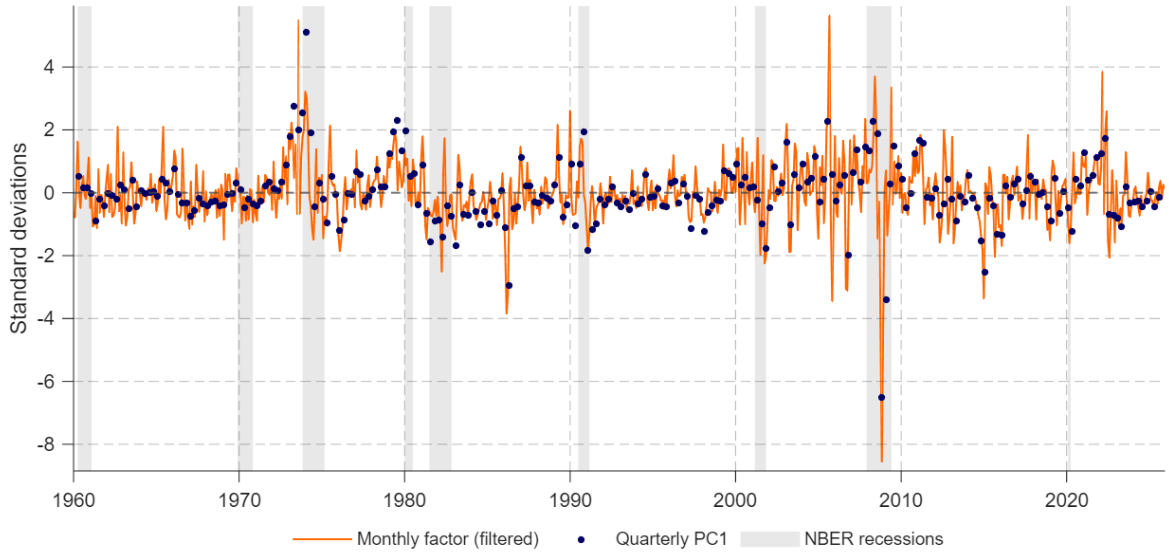

Figure B1: Monthly supply-shock index

Notes: The red line is the filtered estimate of the monthly supply factor $f_{t|t}$ from the mixed-frequency dynamic factor model in (B1)–(B3), standardized to zero mean and unit variance. The blue dots are the first principal component of the three proxies computed at quarterly frequency, as in Benigno and Eggertsson (2023), standardized and plotted in the middle month of each quarter. Positive values correspond to adverse (inflationary) supply shocks. Shaded areas denote NBER recessions. Sample: January 1960 to December 2025.

C Tightness risk

Figure C2 plots the two measures of labor market risk that enter the contribution to inflation risk in (6). The first is downside tightness risk, $\text{TaR}_{0.1,t} - \text{TaR}_{0.5,t}$, the distance between the slack tail and the median of the conditional distribution of tightness. The second is upside tightness risk, $\text{TaR}_{0.9,t} - \text{TaR}_{0.5,t}$, the distance between the tight tail and the median. Since the first step includes past inflation of the relevant price index among the controls, the measures differ slightly between headline and core inflation.

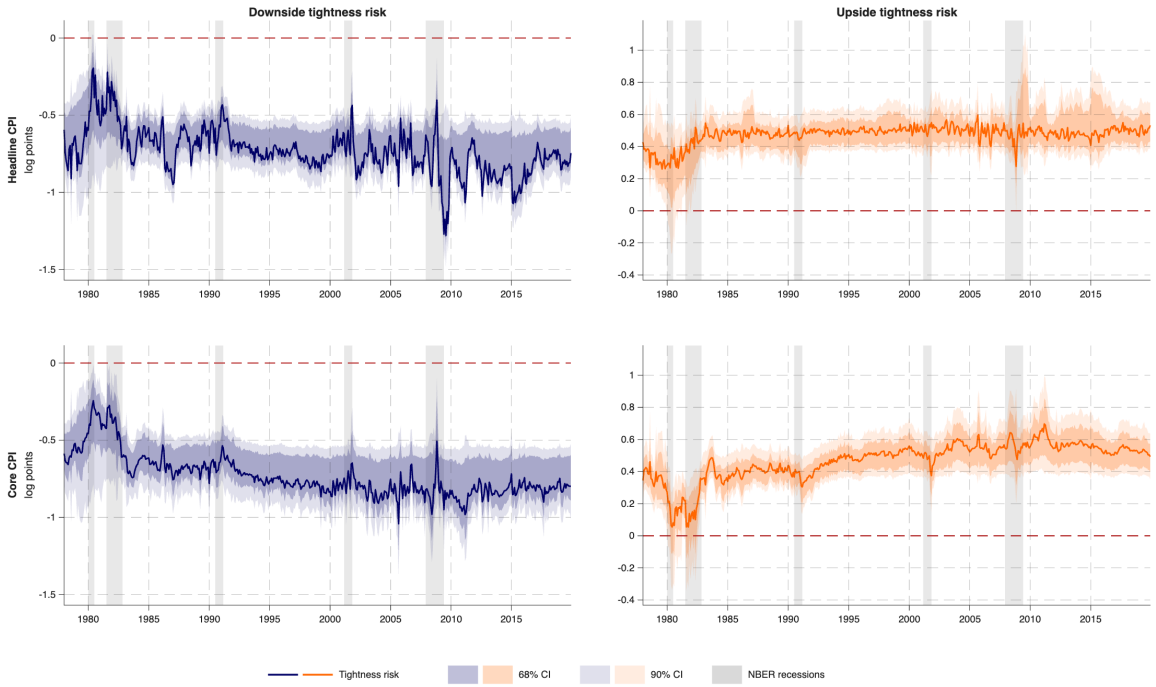

Figure C2: Downside and upside tightness risk

Notes: Downside tightness risk is $\text{TaR}_{0.1,t} - \text{TaR}_{0.5,t}$ and upside tightness risk is $\text{TaR}_{0.9,t} - \text{TaR}_{0.5,t}$, where $\text{TaR}_{q,t}$ is the fitted q-quantile of log labor market tightness from (4), in log points. The controls are average inflation over the previous 12 months of the indicated price index, the supply-shock index and expected inflation. Shaded areas are pointwise 68 and 90 percent confidence bands from a circular block bootstrap of (4) with blocks of 12 months and 500 replications; they reflect estimation uncertainty in the quantile regression coefficients. Grey bars denote NBER recessions. Sample: 1978:01–2019:12.

The conditional distribution of tightness is skewed towards slack. On average, the slack tail lies about 0.7 to 0.8 log points below the median, whereas the tight tail lies only 0.4 to 0.6 log points above it. Both tails are closest to the median in the early 1980s, when inflation and inflation expectations were high, which explains why the contribution of the labor market to downside inflation risk is smallest in that period. Later in the sample, downside tightness risk for headline inflation widens sharply around the Great Recession and again in 2015, when the collapse of energy prices lowered past headline inflation. For core inflation, it remains more stable after the mid-1980s. Upside tightness risk drifts up over the sample, particularly for core inflation. This increase does not translate into upside inflation risk, because the tail sensitivity at the 90th percentile is close to zero. The contribution to inflation risk is the product of the

two terms, and a sizable labor market risk is associated with inflation risk only where inflation is sensitive to tightness.

D Results including the pandemic period

Table D3 reports the estimates of Table 1 on the sample extended to December 2025, which includes the pandemic and the inflation surge of 2021–2023. For core inflation (Panel B), the main pattern is preserved: the coefficient on labor market tightness is positive and significant at the 1 percent level in the lower tail and at the median at every horizon, and it is largest at the 10th percentile at $h = 12$. The coefficients are smaller than in the pre-pandemic sample, however; at the 10th percentile and $h = 12$, for example, the coefficient falls from 0.72 to 0.44. The upper tail behaves differently at intermediate horizons: at $h = 6$, tightness is now also significant at the 90th percentile, a result driven by the post-pandemic period, in which a very tight labor market coincided with a sharp rise in inflation. At $h = 12$, the coefficient in the upper tail remains small and insignificant. For headline inflation (Panel A), tightness remains significant in the lower tail and at the median at horizons of 3 and 6 months, and at $h = 6$ it is also significant at the 90th percentile, while at $h = 12$ none of the coefficients on tightness is significant. Including the pandemic strengthens the role of supply shocks and inflation expectations in the distribution of headline inflation: supply shocks are significant at all three quantiles at horizons of 3 and 6 months, and the coefficient on expected inflation at the 90th percentile and $h = 12$ rises to 1.89. Overall, the pandemic period adds some evidence that tight labor markets are associated with upside inflation risk at intermediate horizons, but it does not overturn the main result that labor market conditions are most informative about the lower part of the inflation distribution.

Figures D3 and D4 report the labor market contribution to inflation risk on the extended sample, and should be compared with Figures 1 and 2. For core inflation, downside risk associated with the labor market remains negative at every horizon and is significant at short horizons, but it is smaller than in the pre-pandemic sample: at $h = 12$, a slack labor market lowers the 10th percentile of future core inflation by about 0.35 percentage points, against 0.55 before the pandemic, and the 90 percent bands now approach zero. Upside risk for core inflation remains small and not significant at every horizon. For headline inflation, downside risk is smaller and less precisely estimated than in the baseline, and it is significant only at $h = 3$. The main difference

Table D3: Inflation quantile regressions across horizons: sample including the pandemic

	$h=3$			$h=6$			$h=12$		
	$\tau=0.1$	$\tau=0.5$	$\tau=0.9$	$\tau=0.1$	$\tau=0.5$	$\tau=0.9$	$\tau=0.1$	$\tau=0.5$	$\tau=0.9$
Panel A: Headline CPI									
$\ln(v_t/u_t)$	0.379^{***}	0.236^{***}	0.040	0.249^{**}	0.308^{***}	0.570^{**}	0.144	0.170	0.396
	(0.090)	(0.081)	(0.144)	(0.111)	(0.119)	(0.229)	(0.111)	(0.146)	(0.260)
Past average inflation	0.700^{***}	0.711^{***}	0.617^{***}	0.502^{***}	0.455^{***}	0.236^{***}	0.282^{**}	0.081	-0.187
	(0.032)	(0.039)	(0.089)	(0.042)	(0.060)	(0.072)	(0.112)	(0.092)	(0.120)
Supply shock	0.245^{***}	0.193^{***}	0.308^{***}	0.279^{***}	0.263^{***}	0.260^{***}	-0.028	0.063	-0.126
	(0.053)	(0.042)	(0.063)	(0.095)	(0.055)	(0.088)	(0.051)	(0.070)	(0.080)
Expected inflation	0.295^{***}	0.496^{***}	0.731^{***}	0.432^{***}	0.794^{***}	1.200^{***}	0.237	1.073^{***}	1.890^{***}
	(0.053)	(0.048)	(0.124)	(0.083)	(0.134)	(0.099)	(0.305)	(0.182)	(0.184)
Constant	-0.791^{***}	-0.550^{***}	-0.160	-1.029^{***}	-0.721^{***}	0.163	-0.202	-0.597^{*}	-0.321
	(0.152)	(0.092)	(0.190)	(0.254)	(0.272)	(0.302)	(0.607)	(0.306)	(0.446)
Observations		573			570			564	
Panel B: Core CPI									
$\ln(v_t/u_t)$	0.230^{***}	0.160^{***}	0.037	0.315^{***}	0.220^{***}	0.214^{**}	0.442^{***}	0.275^{***}	-0.177
	(0.041)	(0.038)	(0.063)	(0.069)	(0.045)	(0.104)	(0.055)	(0.078)	(0.200)
Past average inflation	0.776^{***}	0.875^{***}	0.917^{***}	0.666^{***}	0.738^{***}	0.772^{***}	0.343^{***}	0.550^{***}	0.614^{***}
	(0.030)	(0.018)	(0.030)	(0.063)	(0.029)	(0.039)	(0.052)	(0.040)	(0.106)
Supply shock	0.002	-0.012	0.001	0.007	0.010	-0.045	-0.019	-0.010	-0.014
	(0.014)	(0.023)	(0.028)	(0.026)	(0.026)	(0.036)	(0.028)	(0.038)	(0.063)
Expected inflation	0.131^{***}	0.261^{***}	0.372^{***}	0.198^{**}	0.427^{***}	0.636^{***}	0.345^{**}	0.544^{***}	0.939^{***}
	(0.039)	(0.037)	(0.037)	(0.078)	(0.063)	(0.045)	(0.137)	(0.071)	(0.117)
Constant	-0.083	-0.366^{***}	-0.455^{***}	-0.104	-0.489^{***}	-0.586^{***}	0.082	-0.353^{**}	-0.888^{***}
	(0.100)	(0.089)	(0.081)	(0.115)	(0.132)	(0.115)	(0.300)	(0.151)	(0.222)
Observations		573			570			564	

Notes: Same specification as Table 1, estimated on the sample extended to include the pandemic period. Smoothed quantile regressions (Fernandes et al., 2021) of year-on-year inflation at $t+h$ on log labor market tightness and three controls: average inflation over the previous 12 months (same price index), the supply-shock index at t (Appendix B) and 1-year expected inflation at t from the University of Michigan survey. Sample: 1978:01–2024:12 for $h=12$ (the end date varies with h so that no data after December 2025 are used). Asymptotic standard errors of the smoothed quantile regression estimator in parentheses. ***, **, * denote significance at the 1, 5 and 10 percent levels.

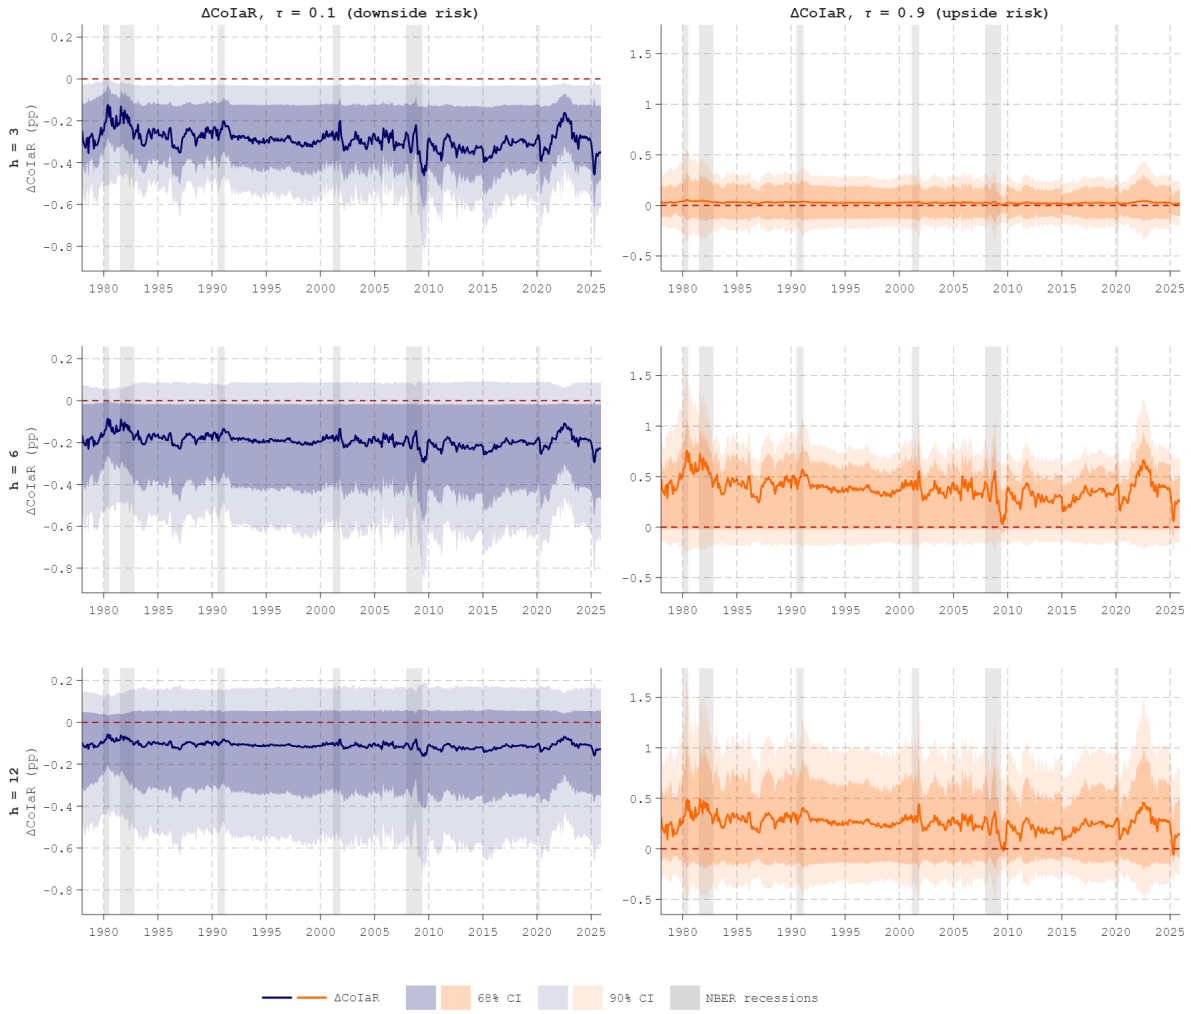

Figure D3: Labor market contribution to headline inflation risk including the pandemic

Notes: Same as Figure 1, estimated on the sample extended to include the pandemic period. Shaded areas are pointwise 68 and 90 percent confidence bands from a circular block bootstrap with blocks of 12 months and 500 replications. Grey bars denote NBER recessions. Sample: 1978:01–2025:12.

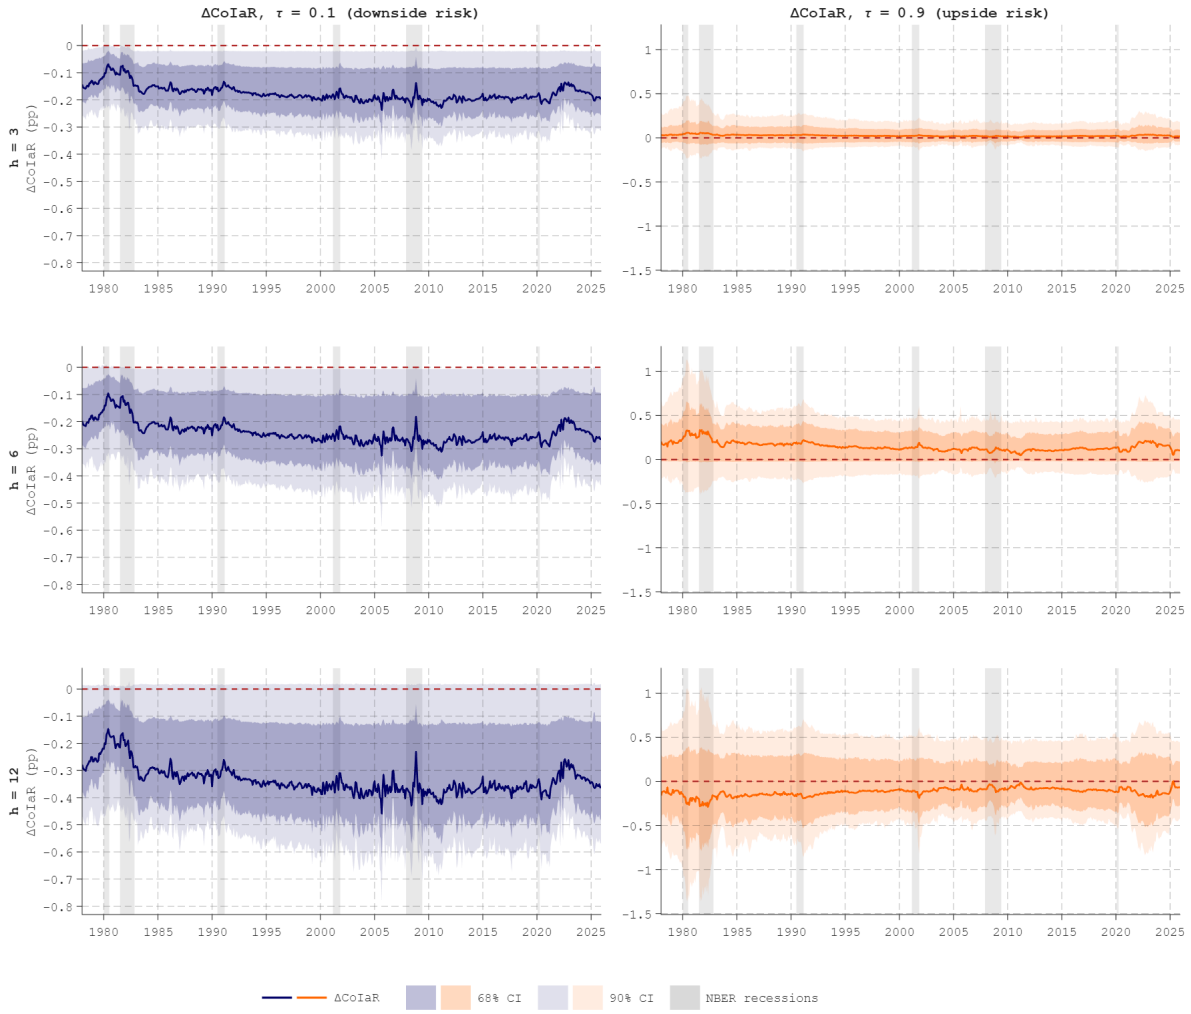

Figure D4: Labor market contribution to core inflation risk including the pandemic

Notes: Same as Figure 2, estimated on the sample extended to include the pandemic period. Shaded areas are pointwise 68 and 90 percent confidence bands from a circular block bootstrap with blocks of 12 months and 500 replications. Grey bars denote NBER recessions. Sample: 1978:01–2025:12.

concerns upside risk at intermediate horizons. At $h = 6$, a tight labor market raises the 90th percentile of future headline inflation by about 0.4 percentage points over most of the sample, with peaks in the early 1980s and in 2022, although the 90 percent bands remain close to zero. This is consistent with the evidence in Table D3 that very tight labor markets were associated with upside inflation risk during the post-pandemic surge. At $h = 12$, upside risk is again not significant.

References

Barnichon, R. (2010). Building a composite Help-Wanted Index. *Economics Letters*, 109(3):175–178.

Benigno, P. and Eggertsson, G. (2023). It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve. Technical report, National Bureau of Economic Research.

Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*. Oxford University Press, 2nd edition.

Fernandes, M., Guerre, E., and Horta, E. (2021). Smoothing quantile regressions. *Journal of Business and Economic Statistics*, 39(1):338–357.

Mariano, R. S. and Murasawa, Y. (2002). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics*, 18(4):427–443.

Michaillat, P. and Saez, E. (2024). $u^* = \sqrt{uv}$: The full-employment rate of unemployment in the united states. *Brookings Papers on Economic Activity*, pages 323–390. Fall 2024.

Tipping, M. E. and Bishop, C. M. (1999). Probabilistic principal component analysis. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, 61(3):611–622.