

INTERNATIONAL MONETARY FUND

Inflation and Labor Markets:

A Bottom-Up View

Sophia Chen, Deniz Igan, Do Lee, Prachi Mishra

WP/24/220

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate.

The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**2024
OCT**



WORKING PAPER

IMF Working Paper
Research Department

Inflation and Labor Markets: A Bottom-Up View
Prepared by Sophia Chen, Deniz Igan, Do Lee, and Prachi Mishra*

Authorized for distribution by Maria Soledad Martinez Peria
October 2024

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: U.S. inflation surged in 2021-22 and has since declined, driven largely by a sharp drop in goods inflation, though services inflation remains elevated. This paper zooms into services inflation, using proprietary microdata on wages to examine its relationship with service sector wage growth at the Metropolitan Statistical Area (MSA) level. We estimate the wage-price pass-through with a local projection instrumental variable model that exploits variation in labor market tightness across MSAs. Our findings reveal a positive and significant relationship between wages and price growth, with a lag. This suggests that the effects of tight labor markets are persistent and may influence the pace of progression toward the inflation target.

RECOMMENDED CITATION: Sophia Chen, Deniz Igan, Do Lee, and Prachi Mishra. 2024. "Inflation and Labor Markets: A Bottom-Up View". IMF Working Paper No. 2024/220.

JEL Classification Numbers:	E24, E31
Keywords:	Inflation; Wages; Labor Market Conditions
Author's E-Mail Address:	ychen2@imf.org , deniz.igan@bis.org , dql204@nyu.edu , prachi.mishra@ashoka.edu.in

* The authors are grateful to Philip Barrett, Nigel Chalk, Era Dabla-Norris, Davide Furceri, Gita Gopinath, Pierre Olivier Gourinchas, Daniel Leigh, Bo Li, Sole Martinez Peria, Papa N'Diaye, and Rodrigo Valdez for valuable discussions and insightful comments and suggestions.

WORKING PAPERS

Inflation and Labor Markets:

A Bottom-Up View

Prepared by Sophia Chen, Deniz Igan, Do Lee, and Prachi Mishra¹

¹ The authors are grateful to Philip Barrett, Nigel Chalk, Era Dabla-Norris, Davide Furceri, Gita Gopinath, Pierre Olivier Gourinchas, Daniel Leigh, Bo Li, Sole Martinez Peria, Papa N'Diaye, and Rodrigo Valdez for valuable discussions and insightful comments and suggestions.

1 Introduction

U.S. inflation rose sharply in 2021-22, with year-on-year CPI inflation peaking at 9.0 percent in June 2022 (Figure 1). By August 2024, the 12-month inflation rate had dropped to 2.5 percent, with indications it may continue falling in the coming months. Both the rise and decline have been equally dramatic.

As dramatic as the movements in headline inflation have been, the divergence between price trends in goods and services has been equally striking. The initial inflation surge in 2021 was driven by a spike in goods prices (Figure 2), resulting from supply-side disruptions in global supply chains and energy prices, along with a pandemic-induced shift in demand from services to goods. Services inflation began rising later in the year as lockdowns eased and demand shifted back from goods to services. While goods inflation peaked in early 2022 and its sharp decline has largely driven the fall in overall inflation, services inflation peaked only in early 2023 and remains elevated. Thus, understanding the drivers of services inflation is crucial to decoding the persistence of inflation.

In addition to the dynamics of goods and services inflation, the labor market has been central to the inflation debate. For instance, Bernanke and Blanchard (2024b) argue that, contrary to initial concerns that inflation would be driven by overly tight labor markets at the onset of the pandemic, the surge in inflation beginning in 2021 was largely due to price shocks—such as sharp increases in commodity prices and specific goods shortages—rather than wage pressures. In contrast, Ball et al. (2022) attribute a significant portion of the inflation rise to labor market slack. Notably, the labor market remains tight, with the vacancy-to-unemployment ratio (V/U) at 1.1 in July 2024 (with a 12-month average of 1.3), compared to “normal times,” when the V/U was well below one. Although the labor market’s contribution to headline CPI inflation has diminished from its peak, it remains the largest estimated contributor to inflation above target, which could have negative implications for the outlook on headline inflation.¹

This paper contributes to the literature by offering one of the first U.S. microdata-based estimates on the impact of service sector wage growth on services inflation through the local labor market tightness channel. We utilize proprietary microdata from Homebase, a payroll service provider for small businesses in the services sector, which covers detailed information on hours and wages for nearly 9 million workers across 1 million U.S. firms. These data are complemented with price data from the Bureau of Labor Statistics (BLS).

We estimate the wage-price pass-through at the Metropolitan Statistical Area (MSA) level, exploring the local labor market tightness channel using a local projection instrumental variable (LP-IV) approach. Our methodology draws on key insights from recent literature on the regional Phillips curve. By using regional data, we avoid the challenge of shifting long-run inflation expectations, which can distort estimates of the Phillips curve slope. It also allows us to differentiate between demand and supply shocks, since central banks cannot counter regional demand shocks with a single national monetary policy instrument.² To exploit variation in regional demand shocks, we instrument wage growth with a shift-share instrument based on local labor market tightness. Similar to the tradable-

¹This is based on the decomposition Ball et al. (2022), extended to the most recent period.

²See Hazell et al. (2022) for a review of the literature and theoretical foundations of the regional Phillips curve.

demand instrument used in Hazell et al. (2022), this instrument captures the idea that national variation in the demand for specific tradable goods will have varying effects on the local demand for non-tradable sectors, depending on the local exposure to the impacted tradable sectors.

We use the vacancy-to-unemployment ratio as our measure of labor market tightness, following Ball et al. (2022) and Blanchard et al. (2022). Before the pandemic, the unemployment rate was the most common measure of labor market tightness, valued for its simplicity and availability. However, several researchers specializing in inflation dynamics, including Furman and Wilson (2021), Barnichon and Shapiro (2022), and Domash and Summers (2022), have argued that the vacancy-to-unemployment ratio better reflects labor market tightness in the post-pandemic economy, where labor force participation has fluctuated significantly.

An first look at the data reveals an interesting pattern between services inflation and service sector wage growth at the aggregate level (Figure 3). Since mid-2021, these two series have shown a strong correlation with a lag, closely aligning with the timing of the rise in headline inflation.³

Our results indicate an important role of local labor market tightness in driving services inflation through service sector wage growth. The first stage shows a strong pass-through from local labor market tightness to wage growth. Specifically, a one-log-point increase in the log vacancy-to-unemployment ratio corresponds to a peak effect of a 27 percentage point increase in year-on-year service wage growth over a 10-month horizon. The second stage indicates that a one-percentage-point increase in service sector wage growth corresponds to a 0.32-percentage-point increase in year-on-year services inflation (excluding housing). Combined, these results suggest that a one-log-point increase in the log vacancy-to-unemployment ratio is associated with an 8.8-percentage-point increase in services inflation (excluding housing) over a 10-month horizon. Moreover, the estimated effects are non-linear, with a higher wage-price pass-through, when labor markets are tighter, or the initial vacancy-to-unemployment ratio is higher. These findings are robust to alternative specifications and are broadly consistent with evidence from more aggregated and lower-frequency data.

These results carry important economic implications. The estimates indicate that local labor market tightness was a key driver of inflation between Q3 2022 and Q1 2023, accounting for an average of 68.7 percent of services inflation (excluding housing) across MSAs during this period. Our findings also reinforce prior evidence from aggregate data, showing that the price Phillips curve—or the relationship between wages and labor market slack—is steeper under tighter labor market conditions.

These findings have important implications for the ongoing policy debate. Even before the post-pandemic inflation surge, services inflation was traditionally the main driver of overall inflation in the United States. Services price growth tends to be more persistent than core goods prices, largely due to the higher labor intensity of services, making them more sensitive to wage growth. Wage growth, in turn, is one of the most persistent input costs,

³The spikes in the wage growth series in early 2020 and early 2021 are consistent with aggregate data from the Current Employment Statistics (CES) (Stewart, 2022). They reflect changes in the composition of the workforce. In 2020, low-paid workers were laid off, pushing the average hourly wage artificially high. In 2021, as the economy reopened, employers filled many of these positions again.

reflecting the frequency of contract resets and other labor market dynamics. The pandemic’s disproportionate impact on the services sector, particularly the difficulty in rehiring workers laid off during lockdowns, has contributed to overall labor market tightness. This highlights the critical role of wage trends in services and their impact on prices. While wage growth accelerated sharply after the pandemic’s acute phase and has recently moderated, ongoing wage pressures are likely as purchasing power has yet to return to pre-pandemic levels.

Overall, our results suggest that a persistently tight labor market will exert persistent pressure on inflation. The effect of the pass-through from wages to service prices takes almost a full year to peak. This pressure may complicate the inflation outlook, especially in an environment prone to supply-side shocks due to geopolitical tensions.

The rest of the paper is organized as follows. Section 2 reviews the rapidly-expanding related literature. Section 3 provides an overview of the data and measurements. Section 4 lays out the empirical strategy. Section 5 presents the findings. Section 6 concludes.

2 Related literature

Our paper is related to the growing literature that seeks to explain inflation dynamics in the post-pandemic period. This body of work has emphasized the importance of shocks to import prices and supply chains, with a natural focus on the goods sector due to its higher tradeability. For instance, Amiti et al. (2022) show that changes in import competition, along with a reduced ability to substitute between labor and intermediate inputs, contributed to the recent surge in inflation. Similarly, LaBelle and Santacreu (2022) highlight the significant impact of global supply chain disruptions on the U.S. Producer Price Index (PPI).

A separate strand of literature highlights the impact of tight domestic labor markets on price pressures. Ball et al. (2022) find that the high vacancy-to-unemployment ratios observed in 2021-2022 can explain a significant portion of the rise in monthly core inflation. They note that “the contribution of V/U to the rise in 12-month inflation is 2.0 percentage points, nearly a third of the total inflation increase. However, the rise in V/U explains more—nearly one-half—of the rise in core inflation and, the effect of V/U is rising over time.”

Similarly, Benigno and Eggertsson (2023) show that the slope of the Phillips curve steepens when the vacancy-to-unemployment ratio is high. They conclude that the recent inflation surge was primarily driven by a labor shortage. Dao et al. (2024) find that tight labor markets continue to contribute significantly to inflation in the U.S., making it an exception among a large sample of advanced and emerging market economies, where inflationary pressure from relative price shocks has subsided, resulting in lower inflation. Bernanke and Blanchard (2024a) find an even larger quantitative role for V/U in explaining price pressures in the US compared to Dao et al. (2024). According to Dao et al. (2024), this difference stems from their inclusion of V/U in their equation for core inflation, while the approach of Bernanke and Blanchard (2024a) focuses solely on the relationship between V/U and wage inflation.

Unlike the aforementioned papers which use the vacancy-to-unemployment ratio, Crump et al. (2022) rely on the unemployment rate as their measure of labor market tightness and complement it with multiple measures of labor compensation. They project underlying

inflation to remain high due to tight labor markets and strong wage growth.⁴ Consistent with these papers, we find a strong role of the labor market in driving recent inflation. A key difference between our approach and these existing studies is that they examine time-series variation in aggregate inflation, while we explore cross-sectional variation across geographical areas. We provide a detailed comparison of the estimates in Section 5.

Our paper is also related to studies that estimate the Phillips curve using cross-sectional data (e.g., Beraja et al., 2019; McLeay and Tenreyro, 2019; Hooper et al., 2020). Hazell et al. (2022) show that using regional data helps overcome the issue of shifting long-run inflation expectations, which can confound the effect of labor market slack. It also allows for a clear distinction between demand and supply shocks. They show that the slope of the regional Phillips curve is steeper than that of an aggregate Phillips curve when labor market conditions are persistent. Most estimations in this literature focus on the pre-pandemic period. For example, Hazell et al. (2022) estimate the slope of the regional Phillips curve using data across U.S. states spanning 1978-2018. An exception is Barnichon and Shapiro (2022), who evaluate the performance of various slack measures in predicting and explaining inflation using MSA-level data from 1982 to 2022. They find that the vacancy-to-unemployment ratio and vacancy filling cost proxies outperform other labor market slack measures, such as the unemployment rate. However, their focus is Phillips curve’s predictive performance and does not examine the wage-price pass-through or estimate the contribution of labor market slack to inflation.

This paper’s main contribution to the literature is the use of proprietary microdata to explore the relationship between services inflation and labor market tightness. This focus is crucial, as services inflation remains elevated, keeping headline inflation well above the Federal Reserve’s 2 percent target. The current debate centers on whether inflation has plateaued at levels above 2 percent, posing a potential challenge for the Fed. As Bernanke and Blanchard (2024b) argue, the effects of overheated labor markets can be persistent, with labor market factors increasingly driving inflation as the influence of goods prices wanes. Labor costs represent a larger share of total costs in services compared to goods, and staggered wage setting further amplifies the role of labor markets in sustaining inflation. By using high-frequency microdata and a cross-sectional Phillips curve framework, we can better identify the key drivers of persistent services inflation.

3 Data and measurements

Our wage data come from a proprietary dataset provided by Homebase, a software company that offers scheduling, payroll reporting, and related services to businesses, primarily in the retail, hospitality, and other service sectors. The dataset is based on timecard records, offering detailed information on work hours and wages from over 80,000 businesses and more than 1 million employees across the United States. It includes granular daily employee-level data on hours worked, wages, job types, and links to the corresponding establishments and parent firms. Additionally, establishment-level details include location (via 5-digit zip codes) and industry classification (via 6-digit NAICS codes). Our sample period spans from

⁴Underlying inflation represents the inflation component that solely depends on the long-run trend and the sequence of current and future unemployment gaps.

January 2019 to December 2022.⁵ We calculate average wages at the monthly frequency for each MSA.

This dataset is notable for its high-frequency, granular coverage of work hours and wage information, setting it apart from other U.S. labor market datasets.⁶ The Homebase dataset’s extensive coverage of low-wage workers and in-person services is especially valuable for studying recent wage dynamics, as this segment experienced particularly strong post-pandemic wage growth (Autor et al., 2023; Chen and Lee, 2024). However, the dataset has limitations, such as the exclusion of certain sectors and the lack of data on tips, benefits, and overtime payments. Despite these limitations, employment trends in the Homebase sample closely correlate with official statistics from the CPS and CES. Additionally, changes in employment and earnings align well with CES data at the month-state level (Dvorkin and Isaacson, 2022; Chen and Lee, 2024).⁷ We discuss the external validity of the Homebase data in detail in Section 5.

For price data, we use the Consumer Price Index (CPI) provided by the Bureau of Labor Statistics (BLS), which offers monthly or bi-monthly data at the MSA level. For MSAs with bi-monthly CPI data, we interpolate linearly to obtain a monthly frequency. Our primary focus is on services excluding housing (i.e., CPI item ‘services less rent of shelter’). We construct year-on-year inflation series at the MSA-month level. We emphasize services inflation because nearly all CPI services are non-tradable (Johnson, 2017).⁸ This focus is important because prices set at the national level—typical for tradable goods—result in a flatter regional Phillips curve. We exclude housing inflation from our analysis because its drivers likely differ from those affecting non-housing services. Consistent with this, the literature has found that housing inflation exhibits a substantially different slope in the Phillips curve compared to non-housing services (Hazell et al., 2022; Stock and Watson, 2020), and that remote work has been a significant driver of post-pandemic housing prices (Howard et al., 2023).

4 Empirical approach

We analyze the data using an LP-IV approach, as used in Jordà et al. (2015) and Jordà et al. (2020). Local projections are a flexible and convenient method for estimating impulse responses (Jordà, 2005), requiring minimal assumptions about the functional form of the responses. The LP-IV approach estimates these impulse responses to shocks using two-stage

⁵The sample on prices and wages starts a year earlier to allow for the calculation of annual changes.

⁶For instance, Compustat lacks data on private firms, the Census lacks detailed wage information, and both the Current Population Survey (CPS) and Current Employment Statistics (CES) lack sufficient geographic variation for cross-sectional analyses. The Quarterly Census of Employment and Wages (QCEW) provides tabulated data by geographic area and industry but does not offer microdata.

⁷The CPS, co-sponsored by the Census Bureau and the BLS, surveys about 60,000 U.S. households and serves as the primary source for official unemployment statistics. The CES, sponsored by the BLS, surveys approximately 145,000 U.S. businesses and government agencies and provides official employment and wage statistics.

⁸According to Johnson (2017), the only three CPI item codes for services classified as tradable are RA04 (Video cassettes, discs, and other media, including rentals), TF09 (unsampled motor vehicle fees), and TG01 (airline fare). Hazell et al. (2022) construct a non-tradable inflation series using BLS microdata, defining non-tradables similarly to the BLS service aggregation but with two exceptions: they classify Food Away from Home as non-tradable and exclude some transportation items (e.g., airline fare).

least squares. Specifically, we estimate:

$$y_{i,t}^h - y_{i,t}^0 = \alpha_t^h + \eta_i^h + \beta^h w_{i,t} + \sum_{k=1}^K \gamma_k^h w_{i,t-k} + \sum_{k=0}^K \delta_k^h y_{i,t-k} + \gamma^h X_{i,s,t} + u_{i,t}^h \quad (1)$$

where i , s , and t denote MSA, state, and month, respectively. $h = 1, \dots, H$ denotes the estimation horizon. We control for lags up to $K = 3$. α_t^h is a time fixed effect. η_i^h is an MSA fixed effect. We cluster the standard errors $u_{i,t}^h$ at the MSA level. Our inflation measure, $y_{i,t}$, is year-on-year services inflation (excluding housing) in MSA i in month t calculated as the logarithmic difference in the prices for services excluding housing, $\ln(prices)_{i,t} - \ln(prices)_{i,t-12}$. Hence, the dependent variable $y_{i,t}^h - y_{i,t}^0$ captures the change in inflation over a horizon of h months from time t to time $t+h$. $w_{i,t}$ is our independent variable of interest. It is year-on-year wage growth in the services sector from Homebase calculated as the logarithmic difference, $\ln(wage)_{i,t} - \ln(wage)_{i,t-12}$.

We include a vector of controls $X_{i,s,t}$ at the state-year level. This includes labor productivity (from BLS) as a proxy for local labor demand, headline inflation shocks defined as the difference between year-on-year headline inflation and core inflation excluding food and energy, and residuals from regressions of previous horizons (i.e., from 1 to $h-1$). According to Teulings and Zubanov (2014), these residuals account for information between time t and $t+h$ that are not fully controlled by the other independent variables (which are as of time t).

The coefficient β^h captures the impulse response of inflation on wage growth over the horizon h . We refer to this coefficient as the wage-price pass-through. The causal interpretation of this coefficient rests on several grounds. Empirically, wages are typically less flexible than prices, making it unlikely that they will quickly adjust to changes in inflation. This is also evident in our sample period. As Figure 3 shows, the rise in wages precedes the rise in inflation.⁹

Moreover, we instrument wage growth with a Bartik shock to local labor market conditions, following Bartik (1991).¹⁰ This allows us to isolate the inflationary component of wage growth that results from labor market competition. To achieve this, we construct a Bartik shock based on local labor market conditions. We measure labor market conditions using the vacancy-to-unemployment ratio (i.e., the number of vacancies divided by the number of unemployed workers), following recent studies (e.g., Ball et al., 2022; Blanchard et al., 2022; Benigno and Eggertsson, 2023). This approach is preferable to using the traditional unemployment gap measure, given the upward shift in the Beveridge curve post-pandemic. The Beveridge curve, which shows the number of vacancies per unemployed worker, has shifted upward, suggesting a similar shift in the traditional unemployment-based Phillips curve, which implies higher inflation at any given unemployment rate. Therefore, using the vacancy-to-unemployment ratio better captures this dynamic.

Formally, the Bartik shock in MSA c in month t is defined as the projected log vacancy-to-unemployment ratio:

$$Shock_{i,t} = \widehat{\ln(V)}_{i,t} - \widehat{\ln(U)}_{i,t}, \quad (2)$$

⁹The lagged response of inflation to wage growth is a common identification assumption in structural vector autoregression (SVAR) estimations. See, for example, Bernanke and Blanchard (2024b).

¹⁰See Soh et al. (2022) and Chen and Lee, 2024 for applications of this method in the context of the pandemic.

where $\widehat{\ln(V)}_{i,t}$ is calculated as

$$\widehat{\ln(V)}_{i,t} = \sum_{k=1}^K \phi_{i,k,2015-2016} * \ln(V_{k,t}). \quad (3)$$

$\ln(V_{k,t})$ is the log number of vacancies for 3-digit NAICS industry k in time t . We include all industries excluding public administration. $\phi_{c,k,2015-2016}$ is the average employment share of industry k in MSA i in 2015-2016. We define $\ln(U_{k,t})$ similarly.

We use industry-level data on unemployment from the CPS and vacancy data from a proprietary dataset provided by Indeed. Indeed is a global search engine for job listings that collects job postings from various sources, including job listing sites, employer career sites, and applicant tracking systems. Duplicated listings are removed, so each job is shown only once, even if it appears on multiple platforms. Our final dataset includes 142 million job postings, covering 421 occupations based on ISCO-08 classifications across 2.9 million companies and 576 counties in the US.

The advantage of using Indeed, as opposed to the Job Openings and Labor Turnover Survey (JOLTS) from the BLS, is its more granular industry classifications, which enable us to capture shocks at the 3-digit NAICS level. A limitation of the Indeed dataset is that it may not capture all job vacancies, as some are not posted online. To address this, we rescale the vacancies from Indeed into nationally representative units, and we describe this adjustment procedure in Appendix A.3. Additionally, Indeed allows us to compute the vacancy-to-unemployment ratio at the MSA-month level—a measure not available from JOLTS. We use this measure to classify local labor markets with labor shortages when investigating the non-linearity of price-wage pass-through based on local labor market conditions.

The identifying assumption underlying our Bartik shock, as Borusyak et al. (2022) show, is that the industry demand shocks, $\ln(V_{k,t})$ and $\ln(U_{k,t})$, are quasi-randomly assigned. This means that the shocks are uncorrelated with relevant unobservables in expectation and that a shock-level law of large numbers applies, that is, the instrument incorporates many sufficiently independent shocks, each with a sufficiently small average exposure. We calculate the shocks at the 3-digit NAICS industry level, rather than at more aggregated levels, in line with the law of large numbers assumption.¹¹ Identification may be threatened if, for example, cost-push shocks are more prevalent in MSAs with industry mixes that would make them more or less susceptible to industry demand shocks. To alleviate this concern, we control for headline inflation shocks at the MSA-month level defined as the difference between year-on-year headline and core excluding food and energy inflation. Ball et al. (2022) show that headlines shocks during our sample period were mainly driven by energy price inflation, disruptions in supply chains, and large price fluctuations in auto-related industries. We also control for labor productivity at the state-year level since prices should reflect unit labor costs rather than wages alone. In Section 5, we discuss the robustness of results with additional controls and other specifications.

¹¹Borusyak et al. (2022) also show that under the additional assumption that the cost-push shocks are spatially uncorrelated, whether to calculate the shocks as leave-one-out averages (i.e., leaving out observations from MSA i when calculating $\widehat{\ln(V)}_{i,t}$ and $\widehat{\ln(U)}_{i,t}$) are unimportant.

5 Findings

Baseline Table 1 presents the summary statistics of the sample. Figure 4 illustrates the baseline results for equation (1). The first stage estimates are positive and statistically significant across all horizons, indicating a positive pass-through from tight labor markets to wages. Similarly, the second stage estimates are positive and statistically significant for horizons ranging from 2 to 12 months, reflecting a positive pass-through from wages to prices.

The results are economically significant. Interpreting the logarithmic difference as a growth rate, the first stage estimate of 26.2 in the 3-month horizon implies that a one-log-point increase in the vacancy-to-unemployment ratio is associated with a 26.2-percentage-point increase in the year-on-year wage growth rate. This indicates a pass-through of 26.2 percent from the vacancy-to-unemployment ratio to wage growth. The effect remains stable over time, peaking at 27.3 percent in the 10-month horizon.

For the second stage, the estimate of 17.3 percent over the 3-month horizon suggests that a one-percentage-point increase in year-on-year wage growth is associated with a 0.173-percentage-point increase in year-on-year inflation. This effect also increases over time, reaching 32.1 percent in the 10-month horizon.

Combining the two stages, a one-log-point increase in the vacancy-to-unemployment ratio corresponds to an increase in year-on-year inflation of 4.5 percentage points over the 3-month horizon, rising to 8.8 percentage points over the 10-month horizon.

We conduct a decomposition exercise for inflation drivers using the baseline estimation results. Figure 5 illustrates the decomposition results based on the 3-month estimations averaged across the MSAs. We find that local labor market tightness, captured by the log vacancy-to-unemployment ratio, emerged as the key driver of inflation since mid-2021. It accounted for, on average, 68.7 percent of the inflation between Q3 2022 and Q1 2023. This result contrasts with the pre-pandemic period and the post-pandemic period up to mid-2021, when common shocks—captured by the time fixed effects—largely offset the effects of local labor market tightness.

These results hold across all MSAs in the sample. Figure 6 presents the decomposition exercise at the MSA level. While the timing and magnitude of the inflation surge vary across MSAs, the model accurately predicts actual inflation in each MSA. Local labor market tightness remains a key driver of inflation between Q3 2022 and Q1 2023 in all MSAs. The share of inflation explained by local labor market tightness ranges from 48.1 to 80.3 percent, with an average of 63.5 percent (see Figure 7).

Robustness Recent theoretical and empirical work has emphasized the importance of non-linearity in Phillips curves (Ball et al., 2022; Benigno and Eggertsson, 2023; Dao et al., 2024). To explore potential non-linearity in our sample, we first examine whether the estimated coefficients vary with the extent of labor shortages. As defined in Benigno and Eggertsson (2023), a labor shortage occurs when the vacancy-to-unemployment ratio exceeds 1, indicating more job vacancies than workers seeking employment.

Figure 8 shows the results when we split the sample into two subsamples based on whether the MSA faced a labor shortage (“High V/U”) or not (“Low V/U”). We find that the first stage results are similar across both subsamples. However, the second stage

results are larger in the presence of a labor shortage, particularly over the 6- to 12-month horizon. In results not reported here, we also explore non-linearity by adding quadratic and cubic terms of the log vacancy-to-unemployment ratio to the baseline regression. While the estimated quadratic term is negative and the cubic term is positive, neither is statistically significant.

The relation between wage growth and inflation we uncover can potentially reflect channels that run beyond the labor market. As discussed in Section 4, the Bartik shock and controls in our baseline specification alleviate the concerns that cost-push shocks, such as food and energy prices, are driving the results. In Figure 9, we additionally control for household income and housing wealth at the MSA level. These controls address the potential concern that the rise in household income and housing wealth—possibly reflecting stimulus spending and rising housing prices—may simultaneously affect local labor markets and inflation. Our results are robust to these controls.

One remaining threat is the potential spatial correlation of the cost-push shocks. To address this threat, we control for time-varying economic conditions at the regional level using region \times time fixed effects. These fixed effects also mitigate spatial correlation due to spillovers. As shown in Figure 10, our results are robust to this control.

Comparison with recent literature How do our results compare to recent estimations that use aggregate data? Benigno and Eggertsson (2023) estimate the pass-through from the log vacancy-to-unemployment ratio to core inflation (measured by all items excluding food and energy) over a 1-quarter horizon. They find a pass-through of 4.7 percentage points with labor shortage (i.e., when the vacancy-to-unemployment ratio is less than 1) and 0.5 percentage points without labor shortage in their 2008-2022 sample. In comparison, our results suggest a pass-through of 5.2 percentage points with labor shortage and 3.6 percentage points without labor shortage.

Ball et al. (2022) estimate the pass-through from the log vacancy-to-unemployment ratio to core inflation (measured by median inflation) and uncover non-linearity with respect to the vacancy-to-unemployment ratio. For example, their results imply that an increase in the vacancy-to-unemployment ratio from 0.5 to 1.5 leads to a 3.5 percentage point rise in core inflation. In contrast, our results indicate a 5.0 percentage point increase in services inflation (excluding housing) for a similar change in the vacancy-to-unemployment ratio.

Overall, our results are comparable to, although slightly larger than, recent estimates based on aggregate data. Differences in the estimates can be attributed to several factors, including variations in data and methods. We expect the pass-through of labor market tightness to be stronger for services inflation compared to core inflation, as services are typically more labor-intensive and less tradable than goods. Additionally, as Hazell et al. (2022) demonstrate in a theoretical model, the slope of the regional Phillips curve is steeper than that of the aggregate Phillips curve when labor market conditions are persistent.¹² As we discuss below, this difference between cross-sectional and aggregate Phillips curves is also observed empirically.

¹²The estimation of an aggregate Phillips curve is also influenced by how long-run inflation expectations are measured. In the regional Phillips curve within a monetary union, long-run inflation expectations are captured by time fixed effects.

External validity Our analysis relies on wage data from Homebase, which provides more granular information than official statistics. This granularity is crucial because official wage data often lack the spatial variation needed to estimate the Phillips curve at the MSA level (e.g., CPS and CES) or the high frequency required to capture rapid movements in post-pandemic wages (e.g., Quarterly Census of Employment and Wages, QCEW). To ensure the robustness of our results, we conduct a series of tests to evaluate the external validity of our results.

First, we assess the external validity of the wage data from Homebase. As discussed above, the national wage trends in the Homebase sample closely align with those from the Current Population Survey (CPS) and the Current Employment Statistics survey (CES) (Dvorkin and Isaacson, 2022; Chen and Lee, 2024). At the state level, we find that monthly wage changes from Homebase and the CES are highly correlated (Chen and Lee, 2024).

We then assess the external validity of our results on wage-price pass-through by re-estimating our baseline specification using official data. Since no official wage data are available at the MSA-month level, we use MSA-quarter level data from the Quarterly Census of Employment and Wages (QCEW).¹³ The sample period for this estimation is from Q1 2019 to Q4 2022 as the baseline.

As shown in Figure 11, the first stage coefficient ranges from 26.3 to 38.2 over the 1- to 4-quarter horizon. The second stage coefficient increases over the same horizon, ranging from 0.13 to 0.26. The combined estimations from both stages suggest a pass-through that peaks at 7.8 percentage points in the 3-quarter horizon. While these magnitudes are similar to those in the baseline, the coefficients are less precisely estimated with the QCEW sample. Figure 12 decomposes the estimation results at the 1-quarter horizon. We again find that the model closely predicts actual inflation and that local labor market tightness has emerged as the key driver since mid-2021. Overall, our findings on wage-price pass-through from the Homebase data are broadly consistent with evidence from official wage data, despite the limitations of the latter in terms of aggregation and frequency.

Finally, we examine the external validity of our findings regarding the role of labor market tightness. This exercise aims to place our results within the context of a longer time series and to compare them with the literature on the aggregate Phillips curve. To achieve this, we adopt a standard form of the aggregate Phillips curve with labor market tightness measured by the log vacancy-to-unemployment ratio. Details of the model and data are discussed in Appendix A.2, and key findings are summarized here.

As shown in Figure 14, we find a significantly positive coefficient for the log vacancy-to-unemployment ratio, indicating an upward-sloping Phillips curve. The coefficient peaks between 12 and 24 months before declining. Specifically, it is 0.6 at the 3-month horizon and increases to between 1.5 and 2.0 at the 12-month horizon and beyond. Using rolling regressions with a 3-year window, we observe substantial time variation in the coefficient.¹⁴

Another way to highlight the important role of labor market conditions in post-pandemic services inflation (excluding housing) is to compare the model’s fitted values with actual inflation outcomes. Figure 15 demonstrates that the fitted values from the rolling window

¹³The QCEW program by the BLS publishes a quarterly count of employment and wages reported by employers, covering more than 95 percent of U.S. jobs.

¹⁴Benigno and Eggertsson (2023) report a coefficient of 0.6 in the full sample and 4.5 for the post-2008 sample during periods of labor shortage.

regression (“Fitted” line) closely align with actual inflation outcomes, effectively tracking the large surge and subsequent decline in inflation during 2021-2023. In contrast, a model that does not incorporate labor market dynamics (by imposing that V/U equals 1, as shown by the “Fitted w/o $\ln(V/U)$ ” line) predicts an earlier inflation surge and a quicker fall in inflation. Under this model, peak inflation would have been 1 percentage point lower, and inflation would have fallen below 2 percent by early 2023—contrary to the observed data.

6 Conclusion

We use a proprietary micro dataset on wages to estimate wage-price pass-through in the United States, exploiting variation in labor market tightness across MSAs. Our findings reveal that service sector wage growth plays an important role in services inflation through the local labor market tightness channel. The pass-through from a tight labor market to service sector wage growth, and from service sector wage growth to services inflation, are both strong. Taken together, local labor market tightness emerged as a key driver of inflation between Q3 2022 and Q1 2023. These results suggest that the effects of overheated labor markets can be persistent when labor market tightness endures, and that substantial wage growth may impede efforts to curb inflation.

	Obs	Mean	Std. Dev.	Min	Median	Max
Services inflation (excl. housing)	906	3.30	2.41	-3.18	2.87	10.03
Wage growth	906	0.95	17.22	-66.33	3.40	63.33
Bartik shock	906	0.41	0.60	-1.27	0.52	1.23
Vacancy-to-unemployment ratio	906	1.07	0.54	0.15	1.04	3.53
Labor productivity	906	1.94	3.10	-4.50	2.10	7.70
Headline shock	906	0.54	1.03	-1.81	0.35	3.89
Log household income per capita	906	11.15	0.19	10.65	11.13	11.73
Log housing wealth	906	4.97	0.23	4.40	4.97	5.55

Table 1: Summary statistics

Notes: This table reports the summary statistics. The sample is a monthly panel of MSAs over January 2019 to December 2022. Observations are weighted by the average size of the labor force in each MSA at 2019. We measure inflation as year-on-year logarithmic difference in the price level, $\ln(prices)_{i,t} - \ln(prices)_{i,t-12}$ for MSA i and month t . Wage growth denotes year-on-year wage growth in the service sector from Homebase measured by logarithmic difference, $\ln(wage)_{i,t} - \ln(wage)_{i,t-12}$. We instrument wage growth using the Bartik shock for the log vacancy-to-unemployment ratio, which is defined in equation (2). Labor productivity (from BLS) proxies for local labor demand. Headline inflation shocks are defined as the difference between year-on-year headline and core excluding food and energy inflation. Housing wealth is defined as the product between Census homeownership rates and the Freddie Mac House Price Index. Source: BLS, Homebase, and authors' calculations.




Figure 1: Headline inflation

Notes: This figure plots the year-on-year U.S. headline inflation. Source: BLS and authors' calculations.




Figure 2: Services and goods inflation

Notes: This figure plots the year-on-year CPI inflation for all items, services, and goods. Source: BLS and authors' calculations.



Figure 3: Price and wage inflation

Notes: This figure plots year-on-year average hourly wage growth from Homebase (left axis) and year-on-year services (excluding housing) inflation (right axis) from the BLS. Source: BLS, Homebase, and authors' calculations.




Figure 4: Services inflation (excluding housing) response, MSA level baseline model

Notes: This figure plots first and second stage estimates from the LP-IV model in equation (1):

$$y_{i,t}^h - y_{i,t}^0 = \alpha_t^h + \eta_i^h + \beta^h w_{i,t} + \sum_{k=1}^K \gamma_k^h w_{i,t-k} + \sum_{k=0}^K \delta_k^h y_{i,t-k} + \gamma^h X_{i,s,t} + u_{i,t}^h$$

where i , s , and t denote MSA, state, and month respectively. $h = 1, \dots, H$ denotes the estimation horizon. We control for lags up to $K = 3$. $y_{i,t}$ is year-on-year inflation of MSA i in month t measured by the logarithmic difference in the prices for services excluding housing, $\ln(\text{prices})_{i,t} - \ln(\text{prices})_{i,t-12}$. The dependent variable $y_{i,t}^h - y_{i,t}^0$ captures the difference in inflation over a horizon of h months from time t to time $t + h$. α_t^h is a time fixed effect. η_i^h is an MSA fixed effect. $w_{i,t}$ is our independent variable of interest, year-on-year wage growth in the service sector from Homebase measured by logarithmic difference, $\ln(\text{wage})_{i,t} - \ln(\text{wage})_{i,t-12}$. We instrument wage growth using the Bartik shock for the log vacancy-to-unemployment ratio, which is defined in equation (2). We also include a vector of controls $X_{i,s,t}$ at the state-year level, including labor productivity (from BLS) to proxy for local labor demand, headline inflation shocks defined as the difference between year-on-year headline and core excluding food and energy inflation, and residuals from the regressions for previous horizons. The sample is a monthly panel of MSAs over January 2019 to December 2022. Observations are weighted by the average size of the labor force in each MSA at 2019. We cluster the standard errors $u_{i,t}^h$ at the MSA level. The solid line shows the point estimates and the dashed lines show the 90 percent confidence interval. Source: BLS, Homebase, and authors' calculations.




Figure 5: Predictions for services inflation (excluding housing), MSA level baseline model

Notes: This figure plots fitted values from the LP-IV model estimated at the MSA level (Figure 4):

$$y_{i,t}^h = \alpha_{2,t}^h + \eta_{2,i}^h + \beta_2^h \widehat{w}_{i,t} + \sum_{k=1}^K \gamma_{2,k}^h w_{i,t-k} + \sum_{k=0}^K \delta_{2,k}^h y_{i,t-k} + \gamma_2^h X_{i,s,t} + u_{2,i,t}^h$$

$$w_{i,t} = \alpha_{1,t}^h + \eta_{1,i}^h + \beta_1^h Shock_{i,t} + \sum_{k=1}^K \gamma_{1,k}^h w_{i,t-k} + \sum_{k=0}^K \delta_{1,k}^h y_{i,t-k} + \gamma_1^h X_{i,s,t} + u_{1,i,t}^h$$

where i , s , and t denote MSA, state, and month respectively. $h = 3$ denotes the estimation horizon. All values are unweighted averages across MSAs in the sample. The solid line plots actual inflation, and the dashed line plots fitted values from the LP-IV model. The bars show the contribution of each independent variable in the LP-IV model, where the contribution captures the combined effect of each variable from the first and second stages. “ln(V/U)” plots the fitted value from the Bartik shock $Shock_{i,t}$ for the log vacancy-to-unemployment ratio defined in equation (2). “Lag Wage Growth” contains fitted values from $w_{i,t-k}$, “Lag Inflation” from $y_{i,t-k}$, “MSA FE” from $\alpha_{1,t}^h$ and $\alpha_{2,t}^h$, “Time FE” from $\eta_{1,t}^h$ and $\eta_{2,t}^h$, and “Headline Shock” and “Productivity” from the controls in $X_{i,s,t}$. Source: BLS, Homebase, and authors’ calculations.

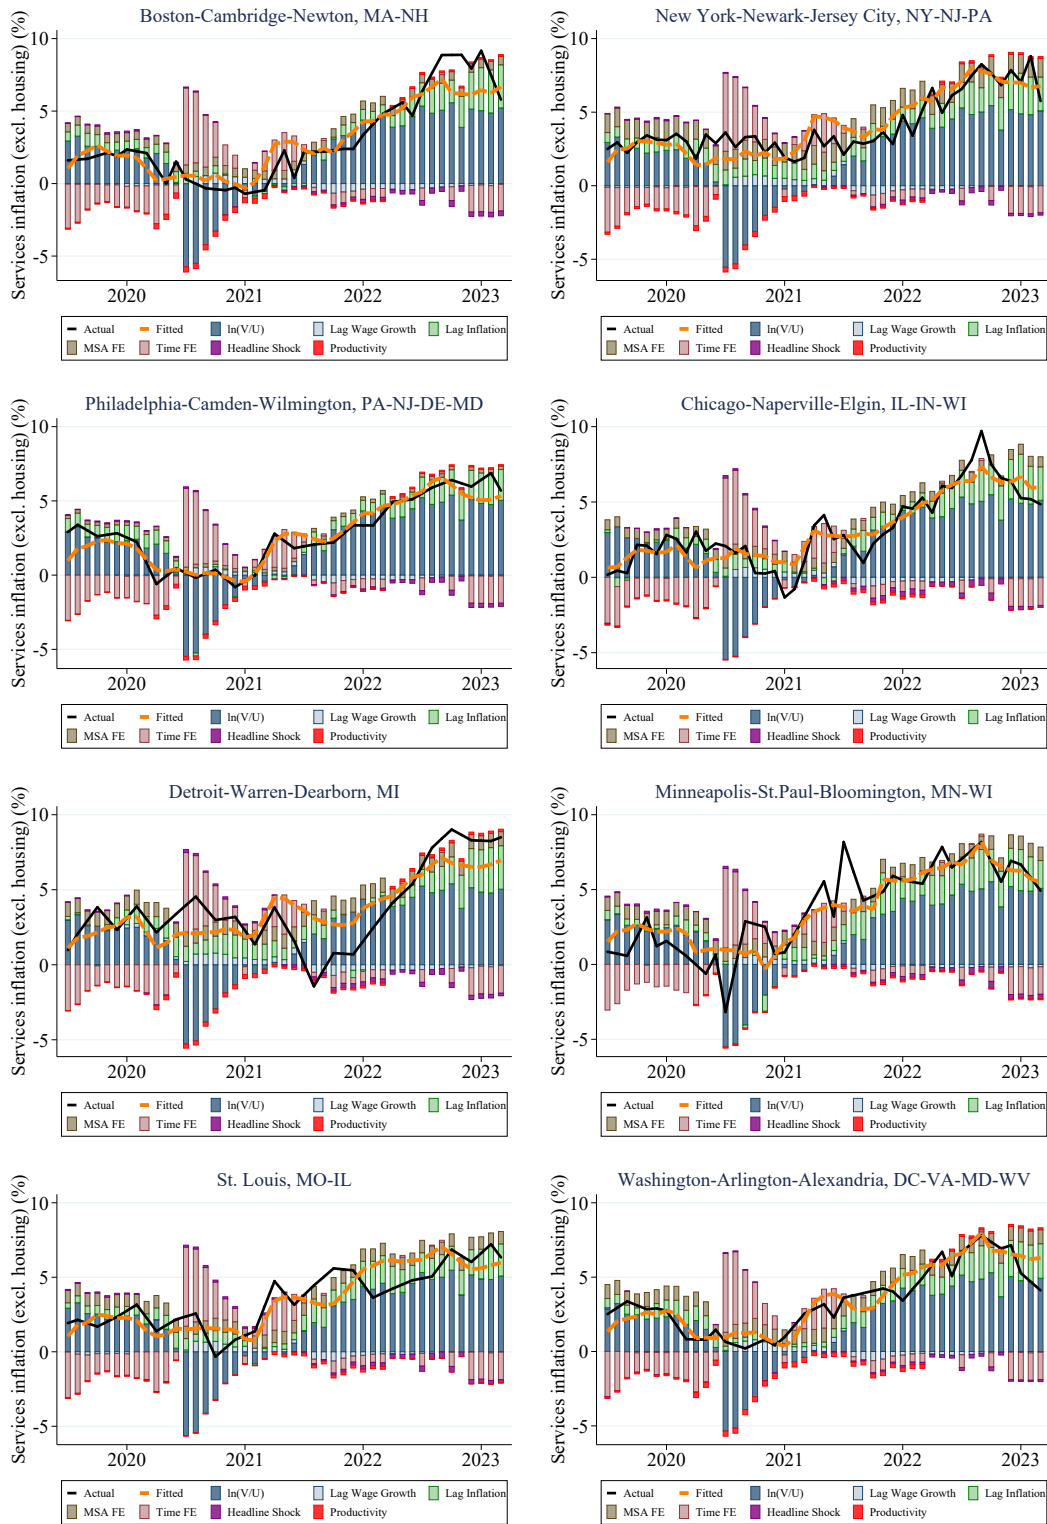


Figure 6: Predictions for services inflation (excluding housing) by MSA

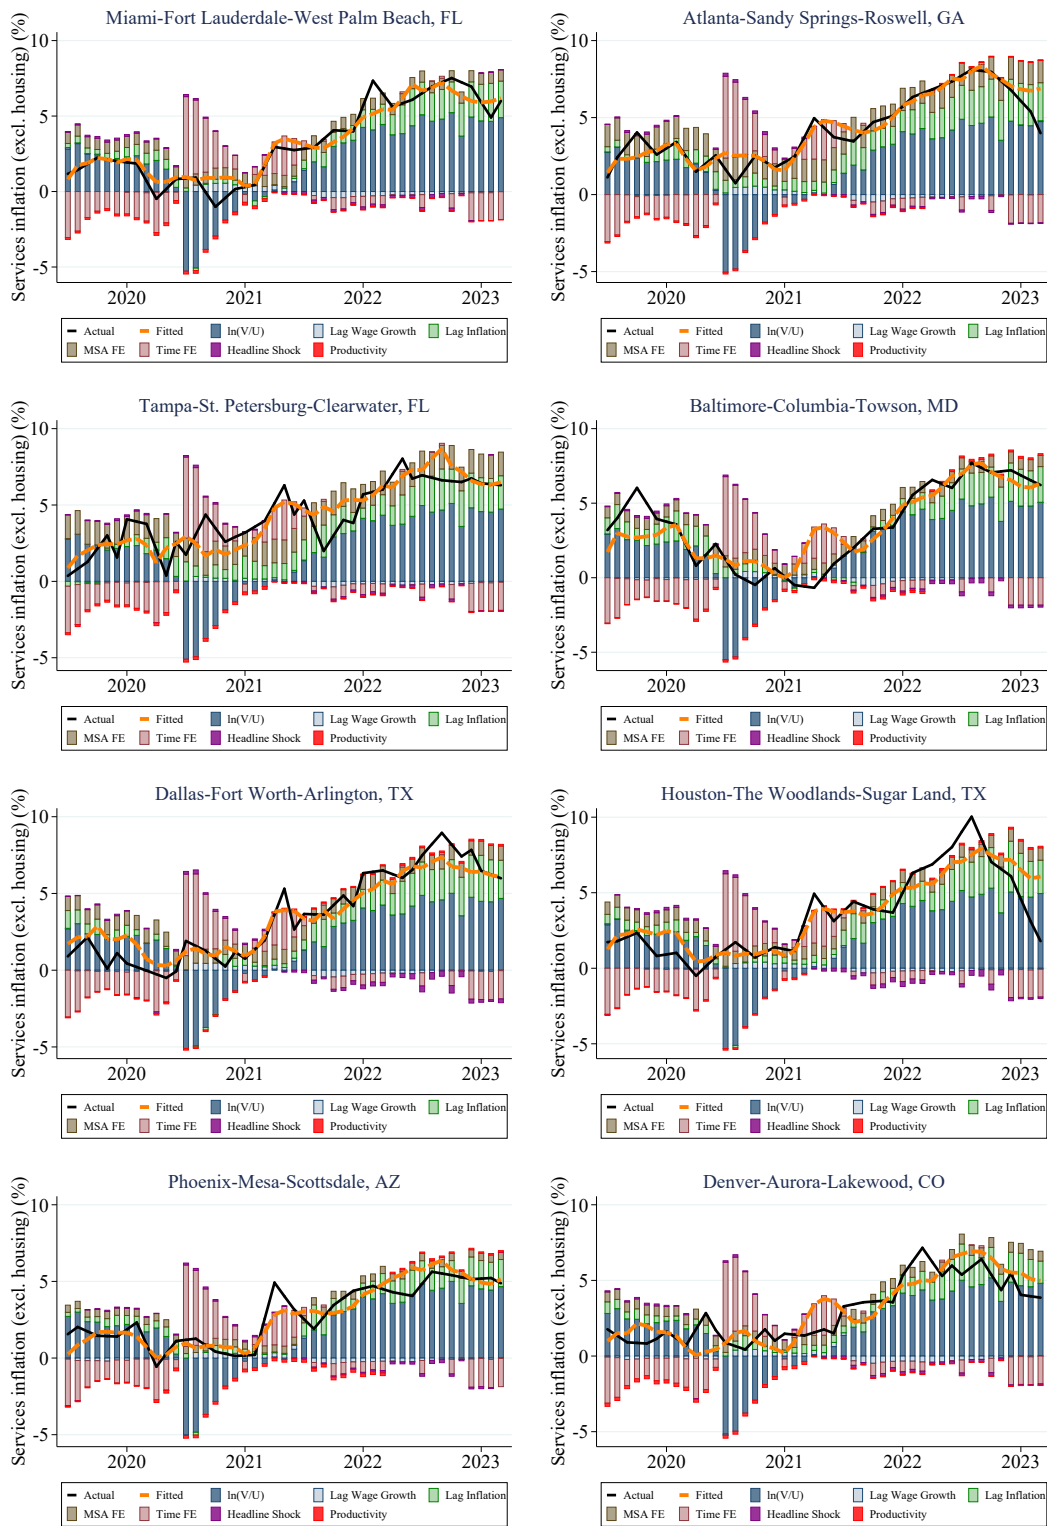


Figure 6: Predictions for services inflation (excluding housing) by MSA (Continued)

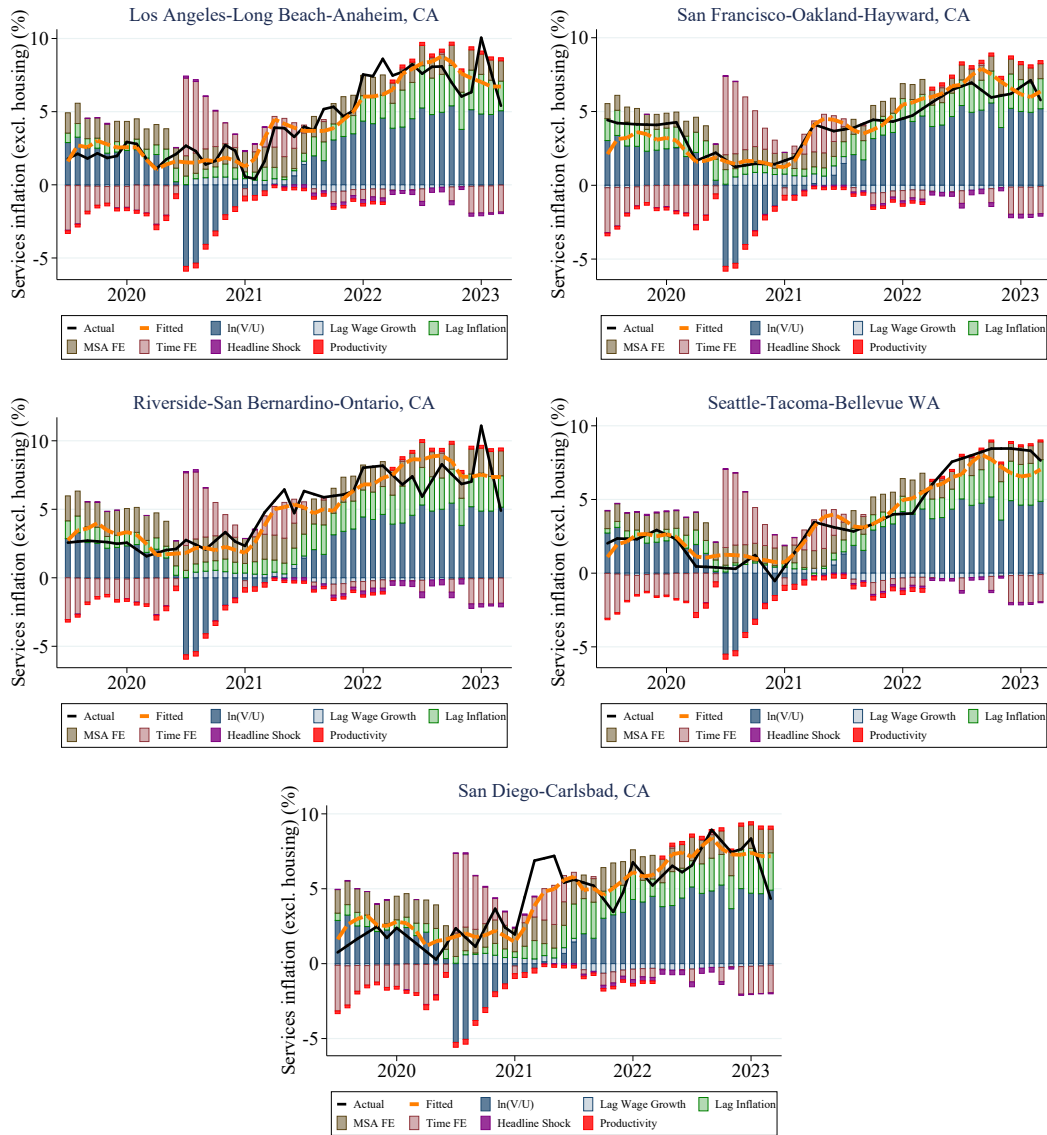


Figure 6: Predictions for services inflation (excluding housing) by MSA (Continued)

Notes: This figure plots fitted values from the LP-IV model estimated at the MSA level (Figure 4). The solid line plots actual inflation, and the dashed line plots fitted values from the LP-IV model. The bars show the contribution of each independent variable in the LP-IV model, where the contribution captures the combined effect of each variable from the first and second stages. “ $\ln(V/U)$ ” plots the fitted value from the Bartik shock $Shock_{i,t}$ for the log vacancy-to-unemployment ratio defined in equation (2). “Lag Wage Growth” contains fitted values from $w_{i,t-k}$, “Lag Inflation” from $y_{i,t-k}$, “MSA FE” from $\alpha_{1,t}^h$ and $\alpha_{2,t}^h$, “Time FE” from $\eta_{1,t}^h$ and $\eta_{2,t}^h$, and “Headline Shock” and “Productivity” from the controls in $X_{i,s,t}$. Source: BLS, Homebase, and authors’ calculations.




Figure 7: Share of services inflation (excluding housing) explained by $\ln(V/U)$

Notes: This figure plots the average share of inflation explained by $\ln(V/U)$ based on fitted values of the LP-IV model estimated at the MSA level at the 3-month horizon (Figure 4):

$$y_{i,t}^h = \alpha_{2,t}^h + \eta_{2,i}^h + \beta_2^h \widehat{w}_{i,t} + \sum_{k=1}^K \gamma_{2,k}^h w_{i,t-k} + \sum_{k=0}^K \delta_{2,k}^h y_{i,t-k} + \gamma_2^h X_{i,s,t} + u_{2,i,t}^h$$

$$w_{i,t} = \alpha_{1,t}^h + \eta_{1,i}^h + \beta_1^h Shock_{i,t} + \sum_{k=1}^K \gamma_{1,k}^h w_{i,t-k} + \sum_{k=0}^K \delta_{1,k}^h y_{i,t-k} + \gamma_1^h X_{i,s,t} + u_{1,i,t}^h$$

where i , s , and t denote MSA, state, and month respectively. $h = 3$ denotes the estimation horizon. Share of services inflation (excluding housing) explained by $\ln(V/U)$ is calculated at the MSA level as the average fitted value for $\ln(V/U)$ over January 2021 to March 2023, as a share of the average actual services inflation (excluding housing) over the same period. Black dashed line shows the average across MSAs. The fitted value for $\ln(V/U)$ refers to the Bartik shock $Shock_{i,t}$ for the log vacancy-to-unemployment ratio defined in equation (2). Source: BLS, Homebase, and authors' calculations.

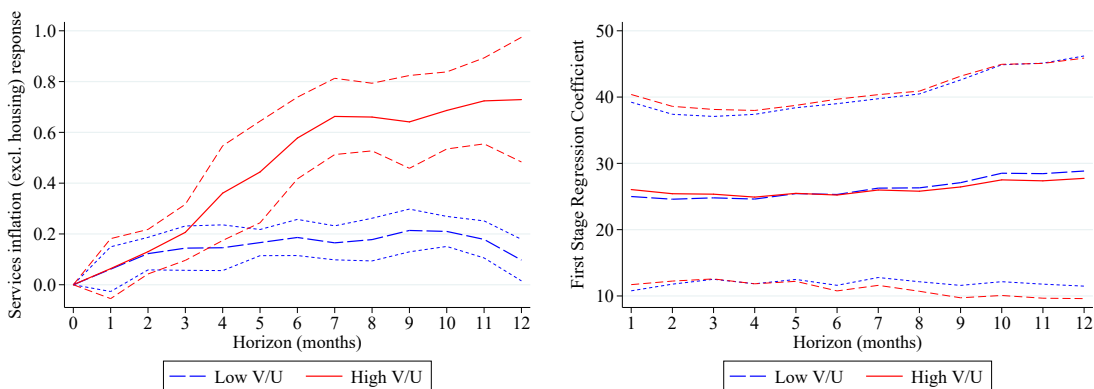


Figure 8: Services inflation (excluding housing) response, labor shortage subsample

Notes: This figure plots first and second stage estimates from the LP-IV model in equation (1) estimated over subsamples split by whether the MSA faced a labor shortage (“High V/U”) or otherwise (“Low V/U”):

$$y_{i,t}^h - y_{i,t}^0 = \alpha_t^h + \eta_i^h + \beta^h w_{i,t} + \sum_{k=1}^K \gamma_k^h w_{i,t-k} + \sum_{k=0}^K \delta_k^h y_{i,t-k} + \gamma^h X_{i,s,t} + u_{i,t}^h$$

where i , s , and t denote MSA, state, and month respectively. A labor shortage refers to a situation where the vacancy-to-unemployment ratio $V_{i,t}/U_{i,t}$ is greater than 1. $h = 1, \dots, H$ denotes the estimation horizon. We control for lags up to $K = 3$. $y_{i,t}$ is year-on-year inflation of MSA i in month t measured by the logarithmic difference in the prices for services excluding housing, $\ln(\text{prices})_{i,t} - \ln(\text{prices})_{i,t-12}$. The dependent variable $y_{i,t}^h - y_{i,t}^0$ captures the difference in inflation over a horizon of h months from time t to time $t+h$. α_t^h is a time fixed effect. η_i^h is an MSA fixed effect. $w_{i,t}$ is our independent variable of interest, year-on-year wage growth in the service sector from Homebase measured by logarithmic difference, $\ln(\text{wage})_{i,t} - \ln(\text{wage})_{i,t-12}$. We instrument wage growth using the Bartik shock for the log vacancy-to-unemployment ratio, which is defined in equation (2). We also include a vector of controls $X_{i,s,t}$ at the state-year level, including labor productivity (from BLS) to proxy for local labor demand, headline inflation shocks defined as the difference between year-on-year headline and core excluding food and energy inflation, and residuals from the regressions for previous horizons. The sample is a monthly panel of MSAs over January 2019 to December 2022. Observations are weighted by the average size of the labor force in each MSA at 2019. We cluster the standard errors $u_{i,t}^h$ at the MSA level. The solid line shows the point estimates and the dashed lines show the 90 percent confidence interval. Source: BLS, Homebase, and authors’ calculations.




Figure 9: Services inflation (excluding housing) response, MSA level model with region-time fixed effects, household income, and housing wealth

Notes: This figure plots first and second stage estimates from the following LP-IV model:

$$y_{i,t}^h - y_{i,t}^0 = \alpha_{r,t}^h + \eta_i^h + \beta^h w_{i,t} + \sum_{k=1}^K \gamma_k^h w_{i,t-k} + \sum_{k=0}^K \delta_k^h y_{i,t-k} + \gamma^h X_{i,s,t} + u_{i,t}^h$$

where i , s , r , and t denote MSA, state, region, and month respectively. $h = 1, \dots, H$ denotes the estimation horizon. We control for lags up to $K = 3$. $y_{i,t}$ is year-on-year inflation of MSA i in month t measured by the logarithmic difference in the prices for services excluding housing, $\ln(\text{prices})_{i,t} - \ln(\text{prices})_{i,t-12}$. The dependent variable $y_{i,t}^h - y_{i,t}^0$ captures the difference in inflation over a horizon of h months from time t to time $t + h$. $\alpha_{r,t}^h$ is a region-time fixed effect. η_i^h is an MSA fixed effect. $w_{i,t}$ is our independent variable of interest, year-on-year wage growth in the service sector from Homebase measured by logarithmic difference, $\ln(\text{wage})_{i,t} - \ln(\text{wage})_{i,t-12}$. We instrument wage growth using the Bartik shock for the log vacancy-to-unemployment ratio, which is defined in equation (2). We also include a vector of controls $X_{i,s,t}$ at the state-year level, including labor productivity (from BLS) to proxy for local labor demand, headline inflation shocks defined as the difference between year-on-year headline and core excluding food and energy inflation, log household income per capita, log housing wealth defined as the product between Census homeownership rates and Freddie Mac House Price Index, and residuals from the regressions for previous horizons. The sample is a monthly panel of MSAs over January 2019 to December 2022. MSAs are mapped to states based on the principal state on which the MSA belongs to. States are categorized into four Census Bureau designated regions. Observations are weighted by the average size of the labor force in each MSA at 2019. We cluster the standard errors $u_{i,t}^h$ at the MSA level. The solid line shows the point estimates and the dashed lines show the 90 percent confidence interval. Source: BLS, Homebase, and authors' calculations.




Figure 10: Services inflation (excluding housing) response, MSA level model with region-time fixed effects

Notes: This figure plots first and second stage estimates from the following LP-IV model:

$$y_{i,t}^h - y_{i,t}^0 = \alpha_{r,t}^h + \eta_i^h + \beta^h w_{i,t} + \sum_{k=1}^K \gamma_k^h w_{i,t-k} + \sum_{k=0}^K \delta_k^h y_{i,t-k} + \gamma^h X_{i,s,t} + u_{i,t}^h$$

where i , s , r , and t denote MSA, state, region, and month respectively. $h = 1, \dots, H$ denotes the estimation horizon. We control for lags up to $K = 3$. $y_{i,t}$ is year-on-year inflation of MSA i in month t measured by the logarithmic difference in the prices for services excluding housing, $\ln(\text{prices})_{i,t} - \ln(\text{prices})_{i,t-12}$. The dependent variable $y_{i,t}^h - y_{i,t}^0$ captures the difference in inflation over a horizon of h months from time t to time $t + h$. $\alpha_{r,t}^h$ is a region-time fixed effect. η_i^h is an MSA fixed effect. $w_{i,t}$ is our independent variable of interest, year-on-year wage growth in the service sector from Homebase measured by logarithmic difference, $\ln(\text{wage})_{i,t} - \ln(\text{wage})_{i,t-12}$. We instrument wage growth using the Bartik shock for the log vacancy-to-unemployment ratio, which is defined in equation (2). We also include a vector of controls $X_{i,s,t}$ at the state-year level, including labor productivity (from BLS) to proxy for local labor demand, headline inflation shocks defined as the difference between year-on-year headline and core excluding food and energy inflation, and residuals from the regressions for previous horizons. The sample is a monthly panel of MSAs over January 2019 to December 2022. MSAs are mapped to states based on the principal state on which the MSA belongs to. States are categorized into four Census Bureau designated regions. Observations are weighted by the average size of the labor force in each MSA at 2019. We cluster the standard errors $u_{i,t}^h$ at the MSA level. The solid line shows the point estimates and the dashed lines show the 90 percent confidence interval. Source: BLS, Homebase, and authors' calculations.

References

- Amiti, Mary, Sebastian Heise, Fatih Karahan, and Aysegül Sahin**, “Inflation Strikes Back: The Return of Wage to Price Pass-Through,” NBER Working Papers 31211, NBER 2022.
- Autor, D., A. Dube, and A. McGrew**, “The Unexpected Compression: Competition at Work in the Low Wage Labor Market,” NBER Working Papers 31010, NBER 2023.
- Ball, L., D. Leigh, and P. Mishra**, “Understanding US Inflation during the COVID Era,” NBER Working Papers 30613, NBER 2022.
- Barnichon, Regis**, “Building a composite Help-Wanted Index,” *Economics Letters*, 2010, 109 (3), 175–178.
- and **Adam Hale Shapiro**, “What’s the Best Measure of Economic Slack?,” FRBSF Economic Letter 2022-04, Federal Reserve Bank of San Francisco 2022.
- Bartik, Timothy J.**, *Who Benefits from State and Local Economic Development Policies?* number wbsle. In ‘Books from Upjohn Press.’, W.E. Upjohn Institute for Employment Research, November 1991.
- Benigno, Pierpaolo and Gauti B Eggertsson**, “It’s Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve,” NBER Working Papers 31197, NBER April 2023.
- Beraja, Martin, Erik Hurst, and Juan Ospina**, “The Aggregate Implications of Regional Business Cycles,” *Econometrica*, 2019, 87 (6), 1789–1833.
- Bernanke, Ben S. and Olivier J Blanchard**, “An Analysis of Pandemic-Era Inflation in 11 Economies,” Working Paper Series WP24-11, Peterson Institute for International Economics May 2024.
- and —, “What Caused the US Pandemic-Era Inflation?,” *American Economic Journal: Macroeconomics*, 2024.
- Blanchard, Olivier, Alex Domash, and Lawrence Summers**, “Bad News for the Fed from the Beveridge Space,” *Peterson Institute for International Economics Policy Brief*, 2022, 22 (7).
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-Experimental Shift-Share Research Designs: Sampling-based vs. Design-based Uncertainty in Regression Analysis,” *The Review of Economic Studies*, 2022, 89 (1), 181–213.
- Chen, Sophia and Do Lee**, “Divergence in Post-Pandemic Earnings Growth: Evidence from Micro Data,” *IMF Working Paper*, 2024.
- Crump, Richard K, Stefano Eusepi, Marc Giannoni, and Aysegül Şahin**, “The Unemployment-Inflation Trade-off Revisited: The Phillips Curve in COVID Times,” NBER Working Papers 29785, NBER February 2022.
- Dao, Mai Chi, Pierre-Olivier Gourinchas, Prachi Mishra, and Daniel Leigh**, “Understanding the International Rise and Fall of Inflation Since 2020,” *Journal of Monetary Economics*, 2024.
- Domash, Alex and Lawrence H. Summers**, “How Tight are U.S. Labor Markets?,” NBER Working Papers 29739, NBER 2022.
- Dvorkin, M. A. and M. Isaacson**, “Tracking Wage Inflation in Real Time,” *Federal Reserve Bank of St Louis, On the Economy Blog*, 2022.
- Furman, Jason and Powell Wilson III**, “What is the Best Measure of Labor Market Tightness?,” Peterson Institute for International Economics blog November 22, Peterson Institute for International Economics 2021.
- Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson**, “The Slope of the Phillips Curve: Evidence from U.S. States,” *The Quarterly Journal of Economics*, 2022, 123 (3), 1299–1344.

- Hooper, Peter, Frederic S. Mishkin, and Amir Sufi**, “Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or is it Just Hibernating?,” *Research in Economics*, 2020, *74* (1), 26–62.
- Howard, Greg, Jack Liebersohn, and Adam Ozimek**, “The Short- and Long-run Effects of Remote Work on U.S. Housing Markets,” *Journal of Financial Economics*, 2023, *150* (1), 166–184.
- Johnson, Noah N.**, “Tradable and Nontradable Inflation Indexes: Replicating New Zealand’s Tradable Indexes with BLS CPI Data,” *Monthly Labor Review*, U.S. Bureau of Labor Statistics 2017.
- Jordà, Òscar**, “Estimation and Inference of Impulse Responses by Local Projections,” *American economic review*, 2005, *95* (1), 161–182.
- LaBelle, Jesse and Ana Maria Santacreu**, “Global Supply Chain Disruptions and Inflation During the COVID-19 Pandemic,” *Review*, 2022, *104* (2), 78–91.
- McLeay, Michael and Silvana Tenreyro**, “Optimal Inflation and the Identification of the Phillips Curve,” *NBER Macroeconomics Annual*, 2019, *34*.
- Soh, Jiaming, Myrto Oikonomou, Carlo Pizzinelli, Ipppei Shibata, and Marina Mendes Tavares**, “Did the COVID-19 Recession Increase the Demand for Digital Occupations in the United States? Evidence from Employment and Vacancies Data,” *IMF Working Papers 2022/195*, International Monetary Fund 2022.
- Stewart, Jay**, “Why was Labor Productivity Growth So High during the COVID-19 Pandemic? The Role of Labor Composition,” *BLS Working Papers 545*, U.S. Bureau of Labor Statistics 2022.
- Stock, James H. and Mark W. Watson**, “Slack and Cyclically Sensitive Inflation,” *Journal of Money, Credit and Banking*, December 2020, *52* (S2), 393–428.
- Teulings, Coen N. and Nikolay Zubanov**, “Is Economic Recovery a Myth? Robust Estimation of Impulse Responses,” *Journal of Applied Econometrics*, 2014, *29* (3), 497–514.
- Òscar Jordà, Moritz Schularick, and Alan M. Taylor**, “Betting the House,” *Journal of International Economics*, 2015, *96* (Supplement 1), S2–S18.
- , —, and —, “The Effects of Quasi-Random Monetary Experiments,” *Journal of Monetary Economics*, 2020, *112*, 381–419.

A Appendix

A.1 Cross-sectional evidence from official statistics

We estimate our LP-IV model using wage data from the QCEW at the MSA quarterly level:

$$y_{i,t}^h - y_{i,t}^0 = \alpha_t^h + \eta_i^h + \beta^h w_{i,t} + \sum_{k=1}^K \gamma_k^h w_{i,t-k} + \sum_{k=0}^K \delta_k^h y_{i,t-k} + \gamma^h X_{i,s,t} + u_{i,t}^h \quad (4)$$

where i , s , and t denote MSA, state, and quarter, respectively. $h = 1, \dots, H$ denotes the estimation horizon. We control for lags up to $K = 1$. $y_{i,t}$ is year-on-year inflation of MSA i in quarter t measured by the logarithmic difference in the prices for services excluding housing, $\ln(prices)_{i,t} - \ln(prices)_{i,t-4}$. The dependent variable $y_{i,t}^h - y_{i,t}^0$ captures the difference in inflation over a horizon of h quarters starting at time t . α_t^h is a time fixed effect. η_i^h is an MSA fixed effect. $w_{i,t}$ is our independent variable of interest, year-on-year wage growth in the service sector from the QCEW measured by logarithmic difference, $\ln(wage)_{i,t} - \ln(wage)_{i,t-4}$. We instrument wage growth using the Bartik shock for the log vacancy-to-unemployment ratio, which is defined in equation (2), where the monthly shock is converted to quarterly frequencies by taking the average within each calendar quarter. We also include a vector of controls $X_{i,s,t}$ at the state-year level, including labor productivity (from BLS) to proxy for local labor demand, headline inflation shocks defined as the difference between year-on-year headline and core excluding food and energy inflation, and residuals from the regressions for previous horizons. Observations are weighted by the average size of the labor force in each MSA at 2019. We cluster the standard errors $u_{i,t}^h$ at the MSA level. The sample period is from Q1 2019 to Q4 2022.¹⁵ Figure 11 plots the estimated coefficients and Figure 12 decomposes estimation results at the 1-quarter horizon.

A.2 Time-series evidence from official statistics

We use official statistics from BLS on vacancy, unemployment, and prices at the national level to estimate a dynamic Phillips curve using the following LP specification

$$y_t^h = \alpha^h + \beta^h \ln \theta_t + \sum_{k=1}^K \gamma_{\theta,k}^h \ln \theta_{t-k} + \sum_{k=1}^K \gamma_{y,k}^h y_{t-k} + \gamma_x^h X_t + u_t^h$$

for monthly time t and horizon $h = 0, 1, \dots, 36$ months. y_t is year-on-year services (excluding housing) inflation in month t , computed as logarithmic changes. The dependent variable y_t^h denotes inflation h months ahead of time t . $\ln \theta_t$ is our independent variable of interest, the log vacancy-to-unemployment ratio from Barnichon (2010) and updated from JOLTS. X_t is a vector of control variables including labor productivity (from BLS), headline inflation shocks, and inflation expectations. Headline inflation shock is defined as the difference between year-on-year headline and core excluding food and energy inflation. Inflation expectations are measured using the 2-year inflation expectations from the Federal Reserve Bank of Cleveland, combined with 12-month inflation expectations from the Livingston survey before 1982. We include $K = 3$ lags of the dependent variable and each independent variable. The sample period is January 1986 to December 2023. Figure 13 plots the estimated coefficients.

We explore the time-variation in this parameter by keeping the horizon fixed at $h = 12$ months (the peak of the estimation above) and running rolling regressions with a 3-year window. Figure 14 plots the estimated coefficients. We find substantial time-variation in the coefficient. The estimate for the 2021-2023 period is around 4.0. This is quantitatively very similar to the estimates in Benigno and Eggertsson (2023) during periods with a high vacancy-to-unemployment ratio.

To illustrate the important role of labor market tightness for services (excluding housing) inflation, Figure 15 plots the fitted values (the ‘‘Fitted’’ line) from the rolling regression above as well as the fitted value (the ‘‘Fitted w/o $\ln(V/U)$ ’’ line) from the same regression without a tight labor market by imposing that V/U equal 1. It is evident that the ‘‘Fitted’’ line offers a much better fit of actual inflation than the ‘‘Fitted w/o $\ln(V/U)$ ’’ line, indicating an important role of labor market tightness in explaining the recent inflation surge, especially after mid-2021. Without a tight labor market, the services (excluding housing)

¹⁵Wage data start from Q1 1990 but the Bartik shock data only starts in Q1 2019.

inflation would have peaked in early 2021 instead of early 2022 and would have fallen under 2 percent, instead of the actual rate of 4 percent, in early 2023.

A.3 Constructing nationally representative vacancies

Data on vacancies from Indeed may not necessarily be nationally representative because some job openings may not be posted online. To address this limitation, we rescale the Indeed vacancies data by applying the following procedure. Nationally representative data are available for unemployment from BLS's CPS (2- and 3-digit NAICS industry level) and vacancies from BLS's JOLTS (2-digit NAICS industry level). To obtain nationally representative vacancies at the 3-digit NAICS level, we first calculate the log vacancy-to-unemployment ratio $\ln(V_{k^3,t}/U_{k^3,t})$ for each k^3 at the 3-digit NAICS level using data on vacancies from Indeed and data on unemployment from the CPS. We regress the 3-digit NAICS level log vacancy-to-unemployment ratio $\ln(V_{k^3,t}/U_{k^3,t})$ onto the corresponding 2-digit NAICS level log vacancy-to-unemployment ratio $\ln(V_{k^2,t}/U_{k^2,t})$, where the 2-digit NAICS level data use vacancies from Indeed and unemployment from the CPS. We obtain the fitted values of the regression to obtain a nationally representative 3-digit NAICS level log vacancy-to-unemployment ratio $\ln(\widehat{V_{k^3,t}}/\widehat{U_{k^3,t}})$. Finally, we add back the 3-digit log unemployment from the denominator of the ratio and then exponentiate the expression to convert the fitted values of log vacancy-to-unemployment ratio into nationally representative 3-digit vacancies: $\exp(\ln(\widehat{V_{k^3,t}}/\widehat{U_{k^3,t}}) + \ln(U_{k^3,t}))$.




Figure 11: Services inflation (excluding housing) response, MSA level using QCEW wages

Notes: This figure plots first and second stage estimates from the following LP-IV model:

$$y_{i,t}^h - y_{i,t}^0 = \alpha_t^h + \eta_i^h + \beta^h w_{i,t} + \sum_{k=1}^K \gamma_k^h w_{i,t-k} + \sum_{k=0}^K \delta_k^h y_{i,t-k} + \gamma^h X_{i,s,t} + u_{i,t}^h$$

where i , s , and t denote MSA, state, and quarter, respectively. $h = 1, \dots, H$ denotes the estimation horizon. We control for lags up to $K = 1$. $y_{i,t}$ is year-on-year inflation of MSA i in quarter t measured by the logarithmic difference in the prices for services excluding housing, $\ln(\text{prices})_{i,t} - \ln(\text{prices})_{i,t-4}$. The dependent variable $y_{i,t}^h - y_{i,t}^0$ captures the difference in inflation over a horizon of h quarters starting at time t . α_t^h is a time fixed effect. η_i^h is an MSA fixed effect. $w_{i,t}$ is our independent variable of interest, year-on-year wage growth in the service sector from the QCEW measured by logarithmic difference, $\ln(\text{wage})_{i,t} - \ln(\text{wage})_{i,t-4}$. We instrument wage growth using the Bartik shock for the log vacancy-to-unemployment ratio, which is defined in equation (2). We also include a vector of controls $X_{i,s,t}$ at the state-year level, including labor productivity (from BLS) to proxy for local labor demand, headline inflation shocks defined as the difference between year-on-year headline and core excluding food and energy inflation, and residuals from the regressions for previous horizons. The sample is a quarterly panel of MSAs over Q1 2019 to Q4 2022. Observations are weighted by the average size of the labor force in each MSA at 2019. We cluster the standard errors $u_{i,t}^h$ at the MSA level. The solid line shows the point estimates and the dashed lines show the 90 percent confidence interval. Source: BLS, QCEW, and authors' calculations.




Figure 12: Predictions for services inflation (excluding housing), MSA level using QCEW wages

Notes: This figure plots fitted values from the LP-IV model estimated at the MSA level (Figure 11):

$$y_{i,t}^h = \alpha_{2,t}^h + \eta_{2,i}^h + \beta_2^h \widehat{w}_{i,t} + \sum_{k=1}^K \gamma_{2,k}^h w_{i,t-k} + \sum_{k=0}^K \delta_{2,k}^h y_{i,t-k} + \gamma_2^h X_{i,s,t} + u_{2,i,t}^h$$

$$w_{i,t} = \alpha_{1,t}^h + \eta_{1,i}^h + \beta_1^h Shock_{i,t} + \sum_{k=1}^K \gamma_{1,k}^h w_{i,t-k} + \sum_{k=0}^K \delta_{1,k}^h y_{i,t-k} + \gamma_1^h X_{i,s,t} + u_{1,i,t}^h$$

where i , s , and t denote MSA, state, and quarter, respectively. $h = 1$ denotes the estimation horizon. The solid line plots actual inflation, and the dashed line plots fitted values from the LP-IV model. All values are unweighted averages across MSAs in the sample. The bars show the contribution of each independent variable in the LP-IV model, where the contribution captures the combined effect of each variable from the first and second stages. “ln(V/U)” plots the fitted value from the Bartik shock $Shock_{i,t}$ for the log vacancy-to-unemployment ratio defined in equation (2). “Lag Wage Growth” contains fitted values from $w_{i,t-k}$, “Lag Inflation” from $y_{i,t-k}$, “MSA FE” from $\alpha_{1,t}^h$ and $\alpha_{2,t}^h$, “Time FE” from $\eta_{1,t}^h$ and $\eta_{2,t}^h$, and “Headline Shock” and “Productivity” from the controls in $X_{i,s,t}$. Source: BLS, Homebase, and authors’ calculations.




Figure 13: Services inflation (excluding housing) response, national level official statistics

Notes: This figure plots estimates of β^h from the following local projection (LP) model:

$$y_t^h = \alpha^h + \beta^h \ln \theta_t + \sum_{k=1}^K \gamma_{\theta,k}^h \ln \theta_{t-k} + \sum_{k=1}^K \gamma_{y,k}^h y_{t-k} + \gamma_x^h X_t + u_t^h$$

for monthly time t and horizon $h = 0, 1, \dots, 36$ months. y_t is year-on-year CPI services inflation (excluding housing) in month t , computed as logarithmic changes. The dependent variable y_t^h denotes inflation h months ahead of time t . $\ln \theta_t$ is our independent variable of interest, the log vacancy-to-unemployment ratio from Barnichon (2010) and updated from JOLTS. X_t is a vector of control variables including labor productivity (from BLS), headline inflation shocks, and inflation expectations. Headline inflation shock is defined as the difference between year-on-year headline and core excluding food and energy inflation. Inflation expectations are measured using the 2-year inflation expectations from the Federal Reserve Bank of Cleveland, combined with 12-month inflation expectations from the Livingston survey before 1982. We include $K = 3$ lags of the dependent variable and each independent variable. The sample period is January 1986 to December 2023. Standard errors are Newey-West. The solid line shows the point estimates and the dashed lines show the 90 percent confidence interval. Source: BLS, Homebase, and authors' calculations.




Figure 14: 12-month services inflation (excluding housing) response, 3-year rolling sample, national level official statistics

Notes: This figure plots rolling estimates from the following local projection (LP) model:

$$y_t^h = \alpha^h + \beta^h \ln \theta_t + \sum_{k=1}^K \gamma_{\theta,k}^h \ln \theta_{t-k} + \sum_{k=1}^K \gamma_{y,k}^h y_{t-k} + \gamma_x^h X_t + u_t^h$$

for monthly time t and horizon $h = 3$ months. The regression is estimated rolling samples with a backward-looking 3-year window, and the x axis denotes the last month in each rolling sample. y_t is year-on-year CPI services inflation (excluding housing) in month t , computed as logarithmic changes. The dependent variable y_t^h denotes inflation h months ahead of time t . $\ln \theta_t$ is our independent variable of interest, the log vacancy-to-unemployment ratio from Barnichon (2010) and updated from JOLTS. X_t is a vector of control variables including labor productivity (from BLS), headline inflation shocks, and inflation expectations. Headline inflation shock is defined as the difference between year-on-year headline and core excluding food and energy inflation. Inflation expectations are measured using the 2-year inflation expectations from the Federal Reserve Bank of Cleveland, combined with 12-month inflation expectations from the Livingston survey before 1982. We include $K = 3$ lags of the dependent variable and each independent variable. Standard errors are Newey-West. The solid line shows the point estimates and the dashed lines show the 90 percent confidence interval. Source: BLS, Homebase, and authors' calculations.




Figure 15: Predictions for services inflation (excluding housing), national level official statistics

Notes: This figure plots fitted values from the local projection (LP) model estimated over 3-year rolling samples (Figure 13):

$$y_t^h = \alpha^h + \beta^h \ln \theta_t + \sum_{k=1}^K \gamma_{\theta,k}^h \ln \theta_{t-k} + \sum_{k=1}^K \gamma_{y,k}^h y_{t-k} + \gamma_x^h X_t + u_t^h$$

for monthly time t and horizon $h = 3$ months. The regression is estimated rolling samples with a backward-looking 3-year window. “Fitted” line plots predictions from the full model. “Fitted w/o $\ln(V/U)$ ” line plots predictions from the same regression without a tight labor market by imposing V/U equals 1. y_t is year-on-year CPI services inflation (excluding housing) in month t , computed as logarithmic changes. The dependent variable y_t^h denotes inflation h months ahead of time t . $\ln \theta_t$ is our independent variable of interest, the log vacancy-to-unemployment ratio from Barnichon (2010) and updated from JOLTS. X_t is a vector of control variables including labor productivity (from BLS), headline inflation shocks, and inflation expectations. Headline inflation shock is defined as the difference between year-on-year headline and core excluding food and energy inflation. Inflation expectations are measured using the 2-year inflation expectations from the Federal Reserve Bank of Cleveland, combined with 12-month inflation expectations from the Livingston survey before 1982. We include $K = 3$ lags of the dependent variable and each independent variable. Source: BLS, Homebase, and authors’ calculations.



PUBLICATIONS

Inflation and Labor Markets: A Bottom-Up View
Working Paper No. WP/2024/220