The Phillips Curve after COVID: Evidence from US Cities

Andrea Cerrato* Giulia Gitti†

March, 2025

Abstract

We estimate the slope of the Phillips curve before, during, and after COVID. To do so, we exploit panel variation in inflation and unemployment dynamics across US metropolitan statistical areas (MSAs), using a shift-share instrument to isolate demand-driven fluctuations in local unemployment rates. We specify a two-region New-Keynesian model to derive the slope of the aggregate Phillips curve from our MSA-level estimates. We find that the slope of the Phillips curve increased substantially during the post-COVID inflation and disinflation episodes, reaching its highest level since the mid-1970s. A sizeable estimated decrease in the elasticity of labor supply and the evolution of the frequency of price change quantitatively account for a large fraction of the post-COVID Phillips curve steepening.

Keywords: inflation, COVID, Phillips curve

JEL code: E30

*Federal Reserve Bank of Boston. Contact: andrea.cerrato@bos.frb.org

†Collegio Carlo Alberto and University of Turin, Economics Department. Contact: giulia.gitti@carloalberto.org

We thank Hassan Afrouzi, Regis Barnichon, Gabriel Chodorow-Reich, Gauti Eggertsson, Yuriy Gorodnichenko, Amy Handlan, Jonathon Hazell, Juan Herreno, Peter Hull, Pascal Michaillat, Enrico Moretti, Emi Nakamura, Maurice Obstfeld, Christina Romer, Stephanie Schmitt-Grohe, Benjamin Schoefer, Jon Steinsson, Martin Uribe, Daniel Wilson, all participants to the May 2022 GEMS conference, the August 2022 Macro Lunch at UC Berkeley, the October 2022 Macro Breakfast at Brown University, and the Spring 2023 NBER Monetary Economics Meeting for helpful suggestions and constructive discussions. We thank Carola Binder, Rupal Kamdar, and Jane M. Ryngaert for kindly sharing their data. We are grateful to Dhiren Patki and David Cai for having assembled and shared the Nielsen MSA-level price index data. The views expressed in this paper do not necessarily reflect the views of the Federal Reserve Bank of Boston, the principals of the Board of Governors, or the Federal Reserve System.

1 Introduction

In June 2022, the 12-month US inflation rate hit a 40-year high at 9% after averaging 2.2% between 2000 and 2020. At the same time, the US labor market reached exceptionally high levels of tightness (Crump et al. 2022; Michaillat and Saez 2022; Blanchard et al. 2022). The standard New Keynesian model derives a structural relationship between inflation and the unemployment gap, commonly known as the New Keynesian Phillips curve. This relationship formalizes the pattern in which workers ask for higher wages and firms increase prices during demand-driven booms. The effect of the unemployment gap on inflation is captured by the slope of the Phillips curve. The general consensus among macroeconomists is that the Phillips curve was flat in the 30 years before the pandemic (Hooper et al., 2020). If the slope remained small, changes in the unemployment gap occurred during and after COVID would have a negligible impact on inflation.

In this paper, we estimate the slope of the Phillips curve before, during, and after the COVID-19 pandemic to assess whether the sensitivity of inflation to the unemployment gap changed in the aftermath of the pandemic. This exercise is complex, as severe demand and supply shocks occurred contemporaneously and within an extremely narrow time frame, limiting statistical power.[1] To overcome these challenges, we combine the use of panel variation in inflation and unemployment at the US metropolitan area level with an instrumental variable approach. Panel data provide us with a larger sample size for parameter estimation than the time series (Mavroeidis et al., 2014). Our empirical strategy is based on a two-region New-Keynesian model of a monetary union that clarifies the threats to identification. Within the model, we derive the regional Phillips curve and relate it to its aggregate counterpart, showing that the slopes of the two coincide.

To our knowledge, this is the first paper providing quasi-experimental estimates of the slope of the Phillips curve during and after COVID. Our benchmark estimates imply a stark steepening relative to pre-COVID, especially during the post-COVID inflation phase. Considering the estimates provided by the literature for periods prior to 1990,[2] we conclude that the US Phillips curve has recently been steeper than at any time since the late 1970s. We further document that the largest part of the Phillips curve steepening is driven by a sharp decrease in the estimated labor supply elasticity, while a smaller part by an increase in the frequency of price changes (Blanco et al., 2024). The estimated steepening of the Phillips curve implies that fluctuations in the unemployment gap explained up to 55% of the post-COVID

[1] Ball et al. (2022) explicitly state that they "do not present results for the pandemic period alone, which would mean estimating seven parameters with ten quarters of data."

[2] See, for instance, Hazell et al. (2022).

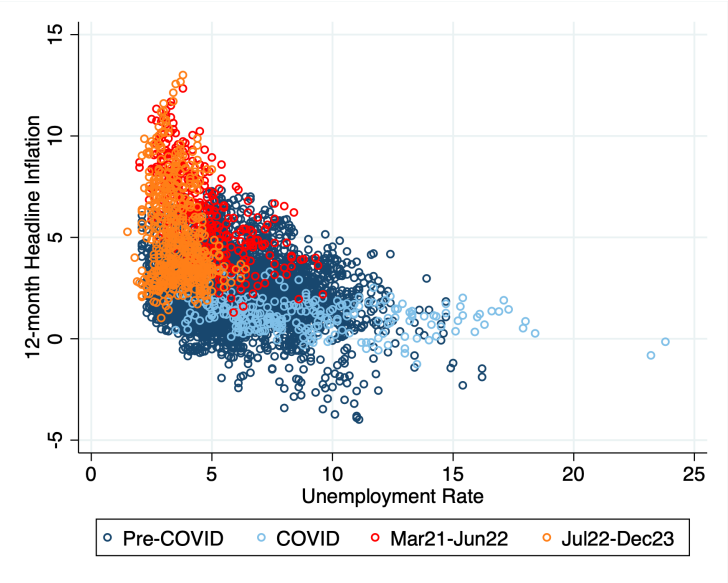
rise in inflation.

To guide our empirical exercise, we rely on a two-region New Keynesian general equilibrium model of a monetary union accounting for the supply-side drivers of COVID and post-COVID inflation dynamics, such as the Great Resignation or semiconductor shortages. We allow for shifts in labor supply preferences and outline a vertically-linked production structure consisting of an international intermediate-input market and local monopolistically competitive final-goods markets. Domestic firms operating in the intermediate-input sector use labor as their factor of production, while final-goods firms employ the intermediate input and labor to produce differentiated consumption goods.

In accordance with the resulting regional Phillips curve equation, local final-goods inflation is driven by short-run inflation expectations and the local unemployment gap, defined as the difference between the actual unemployment rate and the natural rate of unemployment. The natural rate of unemployment depends on three terms. The first denotes the incidence of intermediate-input price shocks on local inflation, capturing the impact of supply chain disruptions. The second represents local shocks to households' disutility of labor, which likely increased during the pandemic causing labor shortages. The third captures local productivity shocks in the final-goods sector. Distinguishing among these three terms of the natural rate of unemployment allows us to address identification concerns stemming from supply-side factors in our empirical estimation.

Figure 1 plots the relationship between 12-month inflation and unemployment rates for 21 MSAs before, during, and after COVID. Raw data clearly point to a steepening after the pandemic. However, the simple correlation shown in Figure 1 could be driven by aggregate and local confounders. At the national level, the Federal Reserve Bank acted promptly to support the economy as it was being hit by COVID and fight inflation in subsequent periods. Endogenous policy responses bias the estimation of the slope of the Phillips curve when using time-series data (Fitzgerald and Nicolini, 2014). In our setting, time fixed effects control for federal policy responses and long-run inflation expectations driven by the monetary policy regime in place, as in Hazell et al. (2022). At the local level, heterogeneous natural unemployment rate dynamics across metropolitan areas bias our estimates because they induce a negative correlation between inflation and the unemployment gap. To address this issue, we follow an empirical strategy based on the combination of an instrumental variable approach and the inclusion of control variables.

Figure 1: The Phillips Correlation Across US Cities



Notes. The scatter plot shows the relationship between the 12-month, all-items inflation rate and the unemployment rate for all observations in our sample. The blue dots denote observations belonging to the pre-COVID period (i.e., Jan 1990-Feb 2020), the green dots denote observations belonging to the COVID period (i.e., Mar 2020-Feb 2021), and the red dots denote observations belonging to the post-COVID period (i.e., Mar 2021-Sep 2022).

Through the lenses of our model, identification requires us to absorb shocks to the relative price of intermediate inputs, local shocks to the disutility of labor, and local productivity shocks in the final-goods sectors. Since relative intermediate-input prices are observable, we directly control for them in our empirical exercise. This term controls, for instance, for the impact of an increase in prices of internationally traded cars on local inflation transmitted through a higher price charged by domestic car dealers. To isolate fluctuations in unemployment rates from local labor supply shocks and productivity shocks in the final-goods sectors, we construct a shift-share instrument proxying for MSA-level labor demand shocks in the tradable intermediate-input sectors (Bartik, 1991). The intuition behind our instrument is that intermediate-input sector labor demand shocks raise employment and wages. Marginal costs in the final-goods sectors consequently increase, driving up prices. This mechanism has a differential impact across cities based on the employment shares of their intermediate-input sectors. For instance, a national labor demand shock in the manufacturing sector affects marginal costs in the final-goods sectors relatively more in manufacturing-intensive cities like Detroit.

We address potential concerns about the validity of our results through several robustness checks. We show that the steepening of the Phillips curve is not mainly driven by the food, energy, and shelter components of the CPI. We also document that our results are robust to controlling for a proxy of regional

long-run inflation expectations, wealth effects on the labor supply induced by the post-pandemic stimulus checks, and alternative measures and instruments for the relative intermediate-input price term. Finally, we leverage retail scanner data to construct MSA-level price indexes for a wider sample of 851 MSAs and verify external validity and robustness of our results to small sample bias and appropriate standard error clustering procedures.

Finally, we evaluate two potential drivers, suggested by our model, of the estimated Phillips curve steepening: changes in the elasticity of the labor supply and the frequency of price adjustments. On the one hand, we leverage variation in unemployment and average monthly earnings, adjusted for inflation and composition, across 385 MSAs, combined with our instrumental variable strategy, to estimate the labor supply elasticity in the pre-COVID and post-COVID periods. We find that the elasticity of labor supply sharply decreased in the aftermath of the pandemic. This result implies that, as the labor supply becomes more inelastic, demand shocks move real wages more than employment, affecting firms' marginal costs and prices. On the other hand, Blanco et al. (2024) document a sizable increase in the frequency of price changes in the post-COVID period. If a larger fraction of firms change prices in each period, demand shocks induce stronger inflation responses. We quantify that the decrease in the labor supply elasticity and the increase in the frequency of price adjustments explain, respectively, 75% and 9% of the steepening of the Phillips curve.

Our paper fills a relevant gap in the literature on post-COVID inflation dynamics: we leverage panel data to identify and estimate the slope of the Phillips curve after COVID, documenting a sizable steepening. A rich and promising theoretical literature has emerged to investigate the causes underlying the steepening of the Phillips curve. Harding et al. (2022) generate a nonlinear Phillips curve through a quasi-kinked demand schedule for goods. Benigno and Eggertsson (2023) and Schmitt-Grohé and Uribe (2022) give rise to the nonlinearity through nonlinear wage Phillips curves. Comin et al. (2023), Bohr (2024), and Fornaro (2024) explore the role of capacity constraints. Blanco et al. (2024), Cavallo et al. (2024), and Karadi et al. (2024) use menu-cost models to endogenize the increase in the frequency of price changes. Erceg et al. (2024) generate a nonlinear price Phillips curve through indexation and learning about the persistency of the shock, while Pfäuti (2023) use consumers' inflation attention thresholds. Our decomposition of the Phillips curve steepening points to the relevance of the nonlinear wage Phillips curve mechanism, and to a lesser extent to the presence of menu-cost types of price rigidities.

A prominent empirical study in this literature is Bernanke and Blanchard (2023), which quantifies the relative contribution of product-market and labor-market shocks to the post-COVID inflation surge,

concluding that product-market shocks played a more important role. However, Bernanke and Blanchard (2023) do not take into account the steepening of the Phillips curve in deriving their results. Some papers have argued about the importance of looking at the vacancy-to-unemployment ratio to measure tightness in the labor market in the post-COVID period (Ball et al., 2022; Bernanke and Blanchard, 2023; Benigno and Eggertsson, 2023; Gitti, 2023; Barnichon and Shapiro, 2024; Beaudry et al., 2025). Following the standard New Keynesian model, we use the unemployment gap as the main independent variable, abstracting from the presence of search-and-matching frictions in the labor market.[3] Other papers use panel variation to study dynamics relevant to the post-COVID inflation surge. Autor et al. (2023) estimate the wage Phillips curve, documenting a steepening of the curve for low-wage workers, in line with our evidence. Smith et al. (2023) estimate breaks in the slope of the Phillips curve by means of Bayesian panel methods, finding no breaks after the pandemic.

Within the large literature on the Phillips curve, our work contributes to the strand that combines panel variation and theory to estimate the slope of the Phillips curve (Fitzgerald and Nicolini 2014, Babb and Detmeister 2017, Beraja et al. 2019, McLeay and Tenreyro 2020, Hooper et al. 2020, Fitzgerald et al. 2020, and Hazell et al. 2022). In particular, our paper most closely relates to Hazell et al. (2022), who rely on state-level panel variation of non-tradable inflation and unemployment to show that the US Phillips curve has been flat since the 1970s. They conclude that pre-COVID inflation dynamics were mostly driven by shifting expectations about long-run monetary policy, as opposed to a time-varying slope of the Phillips curve. With respect to their paper, we make theoretical and empirical contributions. First, we enrich the model to account for the simultaneous occurrence of severe and unprecedented COVID-related shocks. To do so, we feature a distinct, vertically-linked, production structure, and introduce shocks to the disutility of labor. Our model illustrates that it is possible to use headline inflation, as opposed to non-tradable inflation, when estimating the slope of the Phillips curve with regional data. Second, using a similar empirical specification controlling for time fixed effects, we find a novel result: the post-COVID US Phillips curve is substantially steeper.

The remainder of the article is structured as follows. Section 2 describes the model and the derivation of the regional and aggregate Phillips curves. Section 3 discusses data sources and presents summary statistics. Section 4 introduces the empirical strategy and Section 5 shows our main results. Section 6 presents the robustness checks. Section 7 explores the drivers of the steepening of the Phillips curve. Finally, Section 8 concludes.

[3]See Crump et al. (2024) for a thorough discussion of the implications of using the two measures.

2 Vertical Supply Chains and the Phillips Curve

We propose a two-region New-Keynesian model of a monetary union with an internationally traded intermediate-input sector and a local non-tradable final-goods sector in each region. The purpose of the model is to derive the regional Phillips curve in an economic environment featuring two sets of shocks relevant to COVID and post-COVID inflation dynamics (Bernanke and Blanchard, 2023). First, labor supply shocks inducing overheating in the labor market such as the Great Resignation. Second, relative price shocks in the intermediate-input markets capturing fluctuations in global commodity prices and shortages induced by supply chain disruptions. Accounting for these shocks is of paramount importance to identify and estimate the slope of the Phillips curve, as they were pointed out as important drivers of the post-COVID inflation surge and could correlate with the unemployment rate. We show that the slopes of the regional and aggregate Phillips curves coincide and discuss the key assumptions behind this result.

2.1 Model Setup

The economy is made of two regions, i and j, in a monetary union sharing the same preferences, market structure, and firm behavior. Both regions are characterized by a continuum of population of size ζ and $(1-\zeta)$, respectively. Labor is immobile across regions and perfectly mobile across sectors within a region. A common monetary authority sets interest rates following a Taylor rule, featuring a long-run inflation target and a consistent unemployment rate target. In its simplest form, the model abstracts from fiscal policy. In the following sections we describe the economy of region i but the same applies to region j.

2.1.1 Households

The representative household in region i maximizes

$$E_0 \sum_{t=0}^{\infty} \beta^t u(C_{it}, N_{it}), \tag{1}$$

where C_{it} is per-capita consumption of a composite final good, N_{it} is per-capita employment, and β is the household's subjective discount factor. Households have Greenwood–Hercowitz–Huffman (GHH)

preferences over the composite final good and labor as in Greenwood et al. (1988), such that

$$u(C_{it}, N_{it}) = \frac{\left(C_{it} - \chi_{it} \frac{N_{it}^{1+\phi^{-1}}}{1+\phi^{-1}}\right)^{1-\sigma^{-1}}}{1-\sigma^{-1}}, \tag{2}$$

where χ_{it} is an exogenous variable capturing shocks to the disutility of labor, ϕ denotes the household's Frisch elasticity of labor supply, and σ is the household's elasticity of intertemporal substitution. Households have constant elasticity of substitution (CES) preferences over final-goods varieties g such that

$$C_{it} = \Big[\int_0^1 C_{it}(g)^{\frac{\theta-1}{\theta}} \, dg \Big]^{\frac{\theta}{\theta-1}}, \tag{3}$$

where $C_{it}(g)$ is per-capita consumption of final-goods variety g and $\theta > 1$ denotes the elasticity of substitution among different varieties.

The representative household is subject to the following budget constraint:

$$\int_0^1 C_{it}(g) P_{it}(g) \, dg + E_t[M_{it,t+1} B_{i,t+1}] \leq B_{it} + W_{it} N_{it} + \int_0^1 \Pi_{it}(g) \, dg, \tag{4}$$

where $P_{it}(g)$ denotes the price of variety g, $M_{it,t+1}$ is the one-period-ahead stochastic discount factor, B_{it} is a random variable denoting payoffs of the state contingent portfolio held in period t, W_{it} denotes the nominal wage rate, and $\Pi_{it}(g)$ are the profits of the final-goods firm producing variety g. We assume that there is a complete set of financial markets across the two regions and household debt cannot exceed the present value of future income in any state to rule out Ponzi schemes.

GHH preferences rule out wealth effects on the labor supply, implying that the household's optimal labor supply is

$$\chi_{it} N_{it}^{\phi^{-1}} = \frac{W_{it}}{P_{it}}, \tag{5}$$

where P_{it} denotes the price index of the composite final good C_{it} and takes the following form:

$$P_{it} = \left[\int_0^1 P_{it}(g)^{1-\theta} dg \right]^{\frac{1}{1-\theta}}. \tag{6}$$

2.1.2 Firms

The model features an intermediate-input sector and a final-goods sector. The two sectors are vertically linked to capture the incidence of supply chain disruptions on inflation. The representative firm operat-

ing in the intermediate-input sector uses labor to produce a tradable homogeneous intermediate good, according to a decreasing returns to scale production function characterized by region-specific technology. The production function of the representative intermediate-input firm takes the following form:

$$X_{it} = Z_{xit} N_{xit}^{\eta}, \tag{7}$$

where Z_{xit} and N_{xit} denote regional productivity and labor in the intermediate-input sector, respectively. The intermediate input produced by the representative firm, X_{it}, is traded on a perfectly competitive international market, and its price, P_{xt}, is assumed to be exogenous. We discuss the implications of relaxing this assumption in the next section, assuming that P_{xt} is endogenous to domestic economic conditions.

The final-goods sector comprises a continuum of monopolistically competitive firms indexed by g. Each firm specializes in the production of a differentiated good non-tradable across regions. The production function is characterized by constant returns to scale and takes the following form:

$$Y_{it}(g) = Z_{yit} N_{yit}(g)^{\gamma} X_{it}(g)^{1-\gamma}, \tag{8}$$

where Z_{yit} denotes regional productivity in the final-goods sector, and $N_{yit}(g)$ and $X_{it}(g)$ denote the quantity of labor and intermediate good used by firm g, respectively. Final-goods firm g maximizes the present discounted value of current and expected future profits, taking into account that, in each period, it can reset its price only with probability $1-\alpha$, $\alpha \in (0,1)$. The problem takes the following form:

$$\max_{P_{it}(g), N_{yit}(g), X_{it}(g)} \; E_t \sum_{k=0}^{\infty} M_{it,t+k} [P_{it+k}(g) Y_{it+k}(g) - W_{it+k} N_{yit+k}(g) - P_{xt+k} X_{it+k}(g)], \tag{9}$$

subject to the production technology in Equation (8) and the demand for its product

$$Y_{it}(g) = \zeta Y_{it} \left(\frac{P_{it}(g)}{P_{it}} \right)^{-\theta}. \tag{10}$$

The optimal choices of labor and intermediate input imply the factor demand curves

$$W_{it} N_{yit}(g) = \gamma MC_{yit} Y_{it}, \tag{11}$$

$$P_{xt} X_{it}(g) = (1-\gamma) MC_{yit} Y_{it}, \tag{12}$$

where MC_{yit} denotes nominal marginal costs and takes the following form:

$$MC_{yit} = \frac{1}{Z_{yit}} \left(\frac{W_{it}}{\gamma} \right)^{\gamma} \left(\frac{P_{xt}}{1-\gamma} \right)^{1-\gamma}. \tag{13}$$

If firm g is able to re-optimize its price in t, it will set $P_{it}(g)$ to satisfy

$$\sum_{k=0}^{\infty} \alpha^k E_t \left[M_{it,t+k} Y_{it+k}(g) \left(\frac{P_{it}(g)}{P_{it-1}} - \frac{\theta}{\theta - 1} \frac{MC_{yit+k}}{P_{it+k}} \frac{P_{it+k}}{P_{it-1}} \right) \right] = 0. \tag{14}$$

Intuitively, the firm will set its price to be equal to a constant markup, $\frac{\theta}{\theta-1}$, over a weighted average of current and expected future real marginal costs, as with probability α the firm will not be able to change its price in the next period.

2.1.3 Monetary Authority

The monetary authority implements a common monetary policy across the two regions following the Taylor rule:[4]

$$\hat{r}^n = \phi_\pi(\pi_t - \bar{\pi}_{t,t+\infty}) - \phi_u(\hat{u}_t - \bar{u}_{t,t+\infty}) + \varepsilon_t. \tag{15}$$

$\pi_t = \zeta \pi_{it} + (1-\zeta)\pi_{jt}$ denotes economy-wide inflation, where $\pi_{it} = p_{it} - p_{it-1}$ is consumer price inflation in region i and π_{jt} is the counterpart in region j. Within this framework, we define unemployment in i as $U_{it} = 1 - N_{it}$. Then, to a first-order approximation, $\hat{u}_{it} = -\hat{n}_{it}$. Hence, $\hat{u}_t = -\hat{n}_t = -[\zeta \hat{n}_{it} + (1-\zeta)\hat{n}_{jt}]$ denotes the deviation of aggregate unemployment rate from its steady-state value. $\bar{\pi}_{t,t+\infty}$ represents a time-varying long-run inflation target. We further assume that the monetary authority targets an unemployment rate consistent with its long-run inflation target, $\bar{u}_{t,t+\infty} = \frac{(1-\beta)\bar{\pi}_{t,t+\infty}}{\kappa}$.[5] Finally, ϕ_π and ϕ_u ensure a unique locally bounded equilibrium, and ε_t denotes a transitory monetary shock, assumed to follow an AR(1) process. The model in its simplest form abstracts from fiscal policy, as the government does not tax, spend, nor issues debt, and monetary policy has no fiscal implications.

[4] Hatted variables represent deviations from a zero-inflation steady state and lower-case variables are logs of upper-case variables.

[5] κ denotes the slope of the Phillips curve. The relationship between $\bar{\pi}_{t,t+\infty}$ and $\bar{u}_{t,t+\infty}$ is based on the assumption that the monetary authority knows the long-run relationship between inflation and unemployment implied by the aggregate Phillips curve (Equation 21).

2.2 Derivation of Regional and Aggregate Phillips Curve

An equilibrium in this economy is an allocation consistent with households' and firms' optimization, the interest rate rule, and market clearing conditions. Log-linearizing the model around a zero-inflation steady state, we obtain the following expression for regional inflation:

$$\pi_{it} = \beta E_t \pi_{it+1} + \lambda(\hat{mc}_{it} - \hat{p}_{it}), \tag{16}$$

where π_{it} is regional inflation, $E_t\pi_{it+1}$ represents the one-period-ahead inflation expectations and $(\hat{mc}_{it} - \hat{p}_{it})$ denotes real marginal costs. The parameter

$$\lambda = \frac{(1-\alpha\beta)(1-\alpha)}{\alpha} \tag{17}$$

is a function of the discount factor β and the frequency of price change $1-\alpha$, capturing frictions in price setting.

Combining the expressions for the log-linearized marginal costs (Equation 13) and the optimal labor supply (Equation 5), we obtain the regional Phillips curve

$$\pi_{it} = \beta E_t \pi_{it+1} - \frac{\lambda\gamma}{\phi}\hat{u}_{it} + \lambda(1-\gamma)\hat{p}_{xit} + \lambda\gamma\hat{\chi}_{it} - \lambda\hat{z}_{yit}. \tag{18}$$

In addition to short-run inflation expectations, regional inflation π_{it} depends negatively on the unemployment rate $\hat{u}_{it}$ and the productivity shock in the final-goods sector $\hat{z}_{yit}$, and positively on the relative price of the intermediate input $\hat{p}_{xit} = \hat{p}_{xt} - \hat{p}_{it}$ and the labor supply shock $\hat{\chi}_{it}$. The intuition behind the presence of $\hat{p}_{xit}$ in the regional Phillips curve is the following. Given that the price of intermediate input, P_{xt}, is common across regions, an identical absolute intermediate-input price change has a higher (lower) pass-through on regional inflation rates the lower (higher) the baseline regional price level.

According to our model, the natural unemployment rate (i.e., the unemployment rate prevailing in a flexible-price regime) in region i is

$$\hat{u}^n_{it} = \frac{\phi}{\gamma}\Big[(1-\gamma)\hat{p}_{xit} + \gamma\hat{\chi}_{it} - \hat{z}_{yit}\Big], \tag{19}$$

implying that we can re-write the regional Phillips curve as

$$\pi_{it} = \beta E_t \pi_{it+1} - \kappa(\hat{u}_{it} - \hat{u}^n_{it}). \tag{20}$$

The parameter $\kappa = \lambda\gamma\phi^{-1}$ denotes the slope of the regional Phillips curve and governs the response of inflation to fluctuations in the unemployment gap, $\hat{u}_{it} - \hat{u}_{it}^n$. This derivation clarifies that shocks to the relative intermediate-input price, the labor supply, and the productivity of the final-goods sector affect inflation through changes in the natural unemployment rate. Combining the regional Phillips curves in i and j, we obtain the aggregate Phillips curve

$$\pi_t = \beta E_t \pi_{t+1} - \kappa(\hat{u}_t - \hat{u}_t^n). \tag{21}$$

An important implication of Equation (21) is that the slopes of the regional and aggregate Phillips curves coincide and are equal to κ.

This result relies on four key assumptions. First, the vertical production structure in our model assumes that all final goods are non-tradable and trade occurs exclusively in the intermediate-input sector as in Beraja et al. (2019). This implies that overall consumer price inflation depends only on non-tradable final goods at both the regional and aggregate levels. Consequently, the sensitivity of overall consumer price inflation to the unemployment gap (i.e., the slope of the Phillips curve) is the same at the regional and aggregate levels. If we assumed instead that overall consumer price inflation had a tradable component, the sensitivity of inflation to the unemployment gap at the regional level would be lower than at the aggregate level, as tradable goods are priced nationally and do not depend on regional economic fluctuations.[6]

Second, the representative household has GHH preferences over consumption and labor, ruling out wealth effects on labor supply. If the representative household had additively separable preferences over consumption and labor, inflation would additionally depend on consumption at both the regional and aggregate levels. By market clearing, regional consumption equals regional output in the final-goods sector, which in turn depends on the traded intermediate input. Trade in the intermediate-input sector provides risk-sharing between the two regions, which is absent at the aggregate level. This makes the relationship between consumption and unemployment different at the two levels, implying different slopes for the regional versus aggregate Phillips curves.

Third, the intermediate-input price, P_{xt}, is assumed to be exogenously set by international markets at both the regional and aggregate levels. If we assumed instead that the representative final-goods

[6] Hazell et al. (2022) assume that overall consumer price inflation has a tradable component, obtaining that the slope of the non-tradable regional Phillips curve equals the slope of the aggregate Phillips curve. We find that changes in the unemployment gap affect inflation in both final consumption goods – generally considered tradable – and services – generally considered non-tradable. This provides solid empirical ground for our assumption.

firm in each region took P_{xt} as given but P_{xt} was endogenous to national economic conditions, then the sensitivity of inflation to the unemployment gap would be lower at the regional level than at the aggregate level.

Finally, labor is assumed to be immobile across regions as ζ, population in region i, is fixed. This implies that the labor supply elasticity is the same at the regional and aggregate levels. If we allowed for some degree of labor mobility across regions, the regional labor supply elasticity would be larger than the aggregate, resulting in a lower response of wages to labor demand shocks. This in turn would make the regional Phillips curve flatter than the aggregate one.

2.3 Clarifying the Relationship Between κ and ψ

Hazell et al. (2022) formalize how panel variation in inflation and unemployment helps identify the slope of the Phillips curve by differencing out changes in long-run inflation expectations that depend on the monetary policy regime in place. In doing so, they draw a clear distinction between κ – the sensitivity of inflation to current unemployment – and ψ – the sensitivity of inflation to the expected present discounted value of current and future unemployment. To guide the interpretation of our empirical findings, we derive the relationship between κ and ψ in our model.

To do so, we solve the regional Phillips curve forward, assuming that the deviation of the regional unemployment rate from its steady state can be expressed as the sum of a transitory and a permanent component, $\hat{u}_{it} = \tilde{u}_{it} + E_t\bar{u}_{t,t+\infty}$. The permanent component captures private agents' beliefs about the unemployment rate consistent with the monetary authority's long-run inflation target, $\bar{\pi}_{t,t+\infty}$. We obtain

$$\pi_{it} = E_t\bar{\pi}_{t,t+\infty} - \kappa E_t \sum_{j=0}^{\infty} \beta^j (\tilde{u}_{it+j} - \hat{u}^n_{it+j}). \tag{22}$$

According to Equation (22), current regional inflation depends on the expectations about the long-run inflation target, $E_t\bar{\pi}_{t,t+\infty} = -\frac{\kappa}{1-\beta}E_t\bar{u}_{t,t+\infty}$, and the expected present discounted value of all current and future transitory shocks to the regional unemployment rate, $E_t\sum_{j=0}^{\infty}\beta^j\tilde{u}_{it+j}$, and natural rate of unemployment, $E_t\sum_{j=0}^{\infty}\beta^j\hat{u}^n_{it+j}$. Note that the long-run inflation expectations, $E_t\bar{\pi}_{t,t+\infty}$, are common across regions as they depend solely on the monetary policy regime in place.

Assuming for simplicity that $\tilde{u}_{it}$ follows an AR(1) process with autocorrelation coefficient $\rho_{\tilde{u}}$, the regional Phillips curve becomes

$$\pi_{it} = E_t\bar{\pi}_{t,t+\infty} - \psi\tilde{u}_{it} + \omega_{it}, \tag{23}$$

where $\psi = \frac{\kappa}{(1-\beta\rho_{\tilde{u}})}$ and $\omega_{it} = \kappa E_t \sum_{j=0}^{\infty} \beta^j \hat{u}^n_{it+j}$. Note that if unemployment is persistent, then $\psi > \kappa$. To derive Equation (23), we assume full-information rational expectations. However, we obtain a similar result under the weaker assumption that the law of iterated expectations holds. Denoting the agents' subjective forecasts by the superscript F, we derive the following regional Phillips curve:

$$\pi_{it} = E^F_t \bar{\pi}_{t,t+\infty} - \frac{\kappa}{(1-\beta\rho^F_{\tilde{u}})}\tilde{u}_{it} + \kappa E^F_t \sum_{j=0}^{\infty} \beta^j \hat{u}^n_{it+j}. \tag{24}$$

If $\rho^F_{\tilde{u}} < \rho_{\tilde{u}}$, the Phillips curve is less forward-looking than the full-information rational expectations one, resulting in a lower discrepancy between ψ and κ.

3 Data

We collect data covering 21 US metropolitan areas from January 1990 to December 2023. The Bureau of Labor Statistics (BLS) provides monthly or bi-monthly MSA-level CPI data. All prices are collected monthly in the New York, Chicago, and Los Angeles metropolitan areas. In other locations, food and energy prices are collected monthly, and the prices of other items are collected bi-monthly. The starting date of CPI data collection differs among the metropolitan areas included in the sample. For a more detailed description of CPI data, we refer to Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008). We gather data on the following CPI categories: all items, all items excluding food and energy (i.e., core CPI), goods, services, services excluding shelter, and food.

We use these data to construct inflation, our dependent variable, as the 12-month percent difference in the CPI. We use headline and core inflation measures as the dependent variables in our benchmark specification and the other measures of inflation for heterogeneity analysis and robustness checks. We linearly interpolate MSA-specific CPI series collected bi-monthly to fully exploit the variation of MSA-level unemployment rates and instrumental variables in the COVID and post-COVID samples. Since interpolation introduces measurement error in our dependent variable only, our estimates do not suffer from attenuation bias. Appendix B discusses the properties of the inflation interpolation errors in more detail, showing that they are centered around zero and do not correlate with the instrumental variable used in our empirical strategy. We acknowledge that all inflation measures are potentially subject to various types of error, particularly so during a pandemic that dramatically shifted consumers' habits. In this regard, Reinsdorf (2020) and Cavallo (2020) argue that CPI weights reflecting pre-COVID consumption

bundles are likely to bias downward the measure of inflation during COVID.[7]

Monthly MSA-level labor market data are available through the BLS's Local Area Unemployment Statistics (LAUS). The LAUS program uses non-survey methodologies to estimate the number of employed and unemployed individuals for sub-national areas, using the national not-seasonally-adjusted estimates from the Current Population Survey (CPS) as controls. LAUS provide MSA-level estimates of the labor force, employment, unemployment, and unemployment rate. We use the unemployment rate as our main independent variable.

To construct our shift-share instrument and control variables, we need additional data. For the shift components of our instrument, we draw monthly data on national employment by industry, starting from January 1987, from the CPS. For the share components of our instrument, we construct MSA-level industry employment shares from the Census micro-data in 1990, our baseline period in our main specification. Furthermore, we collect several measures of the national producer price index (PPI) from the BLS spanning the January 1990 to December 2023 period to construct the relative intermediate-input price index.[8]

Figure 2: Unemployment and Inflation Across US Cities



Notes. The bar graph shows the June 2022 unemployment rate (red bar) and the June 2022 12-month, all-items inflation rate (blue bar) by US metropolitan area. Monthly unemployment rates at the MSA level come from the LAUS. Monthly CPI data at the MSA level come from the BLS. Inflation rates are computed as 12-month percent differences of MSA-level CPIs.

[7] Inflation in goods increased more than inflation in services and consumers substituted away from services to goods (Figure A.3).

[8] We leverage additional data sources to perform robustness checks. We discuss them in Section 6.

The resulting dataset is a panel of MSA-year-month observations. Figure 2 reveals a remarkable degree of geographical heterogeneity in inflation across MSAs. As of June 2022 – when headline inflation peaked – Phoenix-Mesa-Scottsdale (AZ) was the MSA with the highest inflation rate in the US (12.3%), while New York-Newark-Jersey City (NY-NJ-PA) was the MSA experiencing the lowest inflation rate (6.7%). Unemployment rates instead varied to a lesser extent across MSAs. As of June 2022, Chicago-Naperville-Elgin (IL-IN-WI) experienced the highest unemployment rate (5%), while Minneapolis-St.Paul-Bloomington (MN-WI) displayed the lowest figure in the sample (2.7%).

To further support our evidence, we use the Nielsen Retail Scanner data to construct a retail price index at the MSA level following the methodology developed by Beraja et al. (2019). This endows us with an expanded sample of 851 MSAs useful to provide external validity to our main results, test some key mechanisms behind our findings, and address residual concerns related to potential biases stemming from small sample size or high-frequency variation in confounders and to the clustering of standard errors. For specifications using the MSA-level Nielsen inflation as the dependent variable, we use 2000 County Business Patterns (CBP) data to construct our instrument's industry employment shares for two reasons. First, the representativeness of the Census micro-data for smaller MSAs is limited. Second, CBP and Census industry classifications differ in 1990 but not in 2000.

Finally, we leverage data from the Census Bureau Job-to-Job (J2J) publicly available files. These data are MSA-level aggregates of the Census Longitudinal Employer-Household Dynamics (LEHD) and provide information on labor market transitions and average monthly earnings for job stayers and new hires of different age groups, sex, and education levels. We adjust average nominal monthly earnings for changes in composition and deflate them by the Nielsen price index to obtain a measure of MSA-level real average monthly earnings and use it to examine the mechanisms behind our main result.

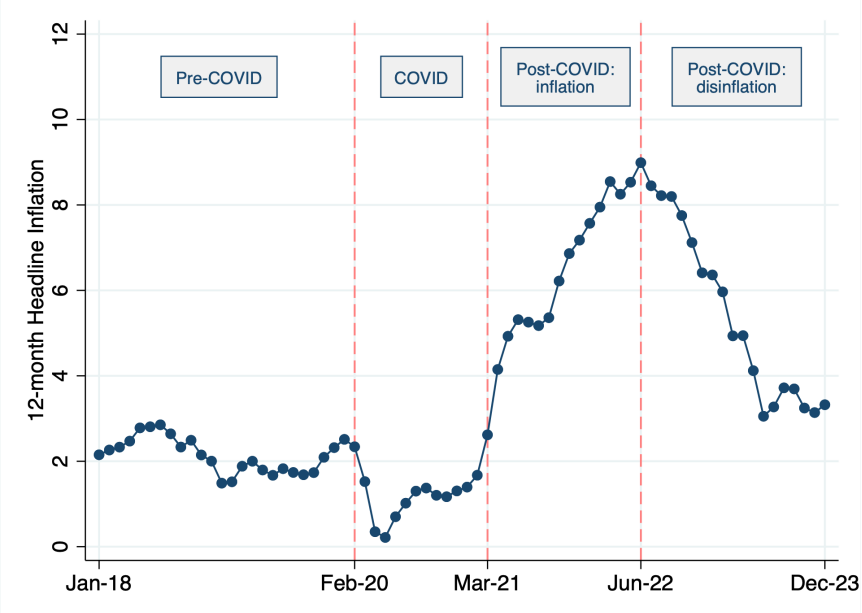
4 Empirical Strategy

The objective of our main empirical exercise is to estimate the parameter ψ in Equation (23) before, during, and after COVID. We define the COVID period as starting in March 2020, when the first COVID cases were reported in the US, and the post-COVID period as starting in March 2021, when real consumption expenditures reverted to their pre-pandemic trend.[9] We further decompose the post-COVID period into two phases capturing the inflation and the disinflation periods. As shown in Figure

[9]Figure A.1 in Appendix A shows the time series of US real consumption expenditures from January 2018 to December 2023, as measured by the Bureau of Economic Analysis (BEA).

3, the inflation phase goes from March 2021 to June 2022, while the disinflation phase spans the July 2022 to December 2023 period.

Figure 3: Breakdown of the COVID and Post-COVID phases in our sample



Notes. The bar graph shows the September 2022 unemployment rate (blue bar) and the September 2022 12-month, all-items inflation rate (red bar) by US metropolitan area. Monthly unemployment rates at the MSA level come from the LAUS. Monthly CPI data at the MSA level come from the BLS. Inflation rates are computed as 12-month percent differences of MSA-level CPIs.

Using the definition of the natural rate of unemployment (Equation 19), the error term in Equation (23) can be expressed as

$$\omega_{it} = \kappa E_t \sum_{k=0}^{\infty} \beta^j \left[(1-\gamma)\hat{p}_{xit+k} + \gamma\hat{\chi}_{it+k} - \hat{z}_{yit+k}\right]. \tag{25}$$

Among the different components of ω_{it}, shocks to the relative price of the intermediate input, $\hat{p}_{xit}$, are observable, while shocks to the labor supply, $\hat{\chi}_{it}$, and the productivity of the final-goods sector, $\hat{z}_{yit}$, are not. Therefore, assuming that $\hat{p}_{xit}$ follows an AR(1) process with an auto-correlation coefficient of ρ_{p_x}, we rewrite Equation (23) as

$$\pi_{it} = E_t\bar{\pi}_{t,t+\infty} - \psi\tilde{u}_{it} + \delta\hat{p}_{xit} + \nu_{it}, \tag{26}$$

where $\delta = \frac{\kappa(1-\gamma)}{1-\beta\rho_{p_x}}$ and $\nu_{it} = \kappa E_t \sum_{k=0}^{\infty} \beta^j \left[\gamma\hat{\chi}_{it+k} - \hat{z}_{yit+k}\right]$.

To test whether the slope of the Phillips curve changes in the aftermath of the pandemic, we estimate Equation (26) separately in the four periods of interest. To do so, we specify the following empirical

model, allowing all parameters to change across periods:

$$\pi_{it} = \eta_i + \xi_t - \psi u_{it} + \delta p_{xit} + \nu_{it}. \tag{27}$$

π_{it} denotes the 12-month headline inflation rate in MSA i and year-month t.[10] η_i represents MSA fixed effects, absorbing time-invariant characteristics of metropolitan areas, such as differences in long-run economic fundamentals across cities. Since these MSA-specific constant terms are allowed to shift across periods, we also control for structural changes at the MSA level brought about by the pandemic. ξ_t denotes year-quarter fixed effects, absorbing aggregate shocks, such as endogenous aggregate fiscal and monetary policies. As we show in Section 2, the inclusion of time fixed effects is also essential to difference out beliefs about the common long-run monetary policy regime in place. Finally, u_{it} is the unemployment rate and p_{xit} is the relative price of the intermediate input, measured as the log difference between the US manufacturing PPI and the MSA-level all-items CPI.

Identifying the parameter ψ in Equation (27) requires u_{it} to be uncorrelated with unobserved local shocks to the labor supply and final-goods sectoral productivity contained in the error term, ν_{it}. To isolate demand-driven fluctuations in MSA-level unemployment rates, we construct a shift-share instrument capturing labor demand shocks in the tradable intermediate-input sectors. Through the lenses of our model, these shocks are driven by changes in intermediate-input sectoral productivity. As a result of a positive intermediate-input sectoral productivity shock, labor demand in the intermediate input sector increases and regional unemployment decreases. Free mobility across sectors implies that wages increase both in the intermediate-input and final-goods sector, causing final-goods firms to increase prices.

Our shift-share instrument takes the following form:

$$z_{xit}^{IV} = \sum_{s=1}^{S} e_{is} \times \Delta_{3Y} \log N_{USst}, \tag{28}$$

where $s = 1, ..., S$ denotes 1990 2-digit tradable intermediate-input Census industries. These industries include agriculture, mining, manufacturing of durable and non-durable goods, wholesale trade, and financial services. $\Delta_{3Y} \log N_{USst}$ denotes the three-year percentage change in national employment by industry, capturing labor demand shocks in the intermediate-input sectors at business cycle frequencies. $e_{is} = \frac{N_{is1990}}{N_{i1990}}$ are MSA-level intermediate-input sectoral employment shares in 1990, the baseline period in our sample, measuring local exposure to national labor demand shocks. Our shift-share instrument

[10] Using the 12-month inflation rate as dependent variable allows us to eliminate seasonality.

captures the idea that a national labor demand shock in the manufacturing sector affects real marginal costs of final-goods firms relatively more in manufacturing-intensive cities like Detroit.

The exogeneity condition for instrument validity requires that z_{xit}^{IV} affects inflation only through the unemployment rate. However, our shift-share instrument can be correlated with both terms in the relative price of the intermediate inputs, $p_{xit} = p_{xt} - p_{it}$. Suppose that the price of intermediate inputs is not exogenous as assumed in the model but dependent on aggregate productivity in the intermediate input sector. Then, national intermediate-input sectoral labor demand shocks driven by changes in sectoral productivity also affect the price of intermediate inputs, p_{xt}, and, ultimately, inflation. Furthermore, as the shift-share instrument is correlated with unemployment, it mechanically affects the relative intermediate-input price through its impact on local prices, p_{it}. We deal with this threat to the exclusion restriction by directly controlling for p_{xit} and instrumenting it with $p_{xit}^{IV} = p_{xt} - p_{it-24}$, where p_{it-24} is the log of the two-year lagged CPI. Note that p_{xit}^{IV} allows us to offset the negative mechanical correlation between the relative intermediate-input price and inflation.

Our identifying assumption is that, conditioning on MSA fixed effects, time fixed effects,[11] and the relative intermediate-input price, local labor demand shocks in the intermediate-input sectors, as proxied by our shift-share instrument, do not correlate with local labor supply and final-goods sectoral productivity shocks. The conditional exogeneity with respect to local labor supply shocks amounts to assuming that national industry-level employment growth rates in the intermediate-input sectors are uncorrelated with industry-level aggregates of local labor supply shocks (Borusyak et al., 2022). The conditional exogeneity with respect to local final-goods sectoral productivity shocks requires these shocks to be uncorrelated with our shift-share instrument.

Table 1 shows the first-stage coefficients and F-statistics for the pre-COVID, COVID, and post-COVID inflation and disinflation periods. In columns (1) to (4), the dependent variable is the unemployment rate, u_{it}. The instrument z_{xit}^{IV} significantly predicts local unemployment, exhibiting negative coefficients across all specifications. The F-statistics indicate that our shift-share instrument is strong in all periods. Column (2) shows that p_{xit}^{IV} is positively correlated with u_{it} during COVID, suggesting that the supply shocks that occurred during the pandemic harmed local labor markets. In columns (5) to (8), we show that the instrument p_{xit}^{IV} strongly predicts the local relative intermediate-input price, p_{xit}.

[11] As recommended by Borusyak et al. (2022), we interact time fixed effects with the sum of employment shares in our shift-share instrument, $\sum_s e_{is}$. This ensures that we absorb heterogeneous trends driven solely by the weight of the intermediate-input industries in the MSA-level industry mix.

Table 1: First-stage coefficients

	u_{it}				p_{xit}			
	(1) Pre	(2) COVID	(3) Post: inflation	(4) Post: disinflation	(5) Pre	(6) COVID	(7) Post: inflation	(8) Post: disinflation
z_{xit}^{IV}	-0.43***	-2.02***	-0.54***	-0.33***	-0.06	0.16*	-0.46***	-0.26***
	(0.04)	(0.26)	(0.09)	(0.04)	(0.05)	(0.09)	(0.15)	(0.08)
p_{xit}^{IV}	0.02	0.25***	-0.00	0.01	0.88***	0.57***	0.63***	0.73***
	(0.01)	(0.09)	(0.02)	(0.01)	(0.02)	(0.04)	(0.02)	(0.03)
Observations	4855	252	336	378	4855	252	336	378
MSA-Period FE	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times$ $\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓	✓	✓
F-stat	51.85	30.37	22.85	34.57	1386.32	123.83	921.98	292.43

Notes. This table presents the first stage regression coefficients for IV estimation of Equation 27. In columns (1) to (4), the dependent variable is the unemployment rate, u_{it}. In columns (5) to (8), the dependent variable is the relative price of intermediate input, p_{xit}. Columns (1) and (5) present the first-stage coefficients for the pre-COVID period (i.e., from January 1990 to February 2020), columns (2) and (6) present the first-stage coefficients for the COVID period (i.e., from March 2020 to February 2021), columns (3) and (7) present the first-stage coefficients for the post-COVID inflation phase (i.e., from March 2021 to June 2022), and columns (4) and (8) present the first-stage coefficients for the post-COVID disinflation phase (i.e., after July 2022). The shift-share instrument constructed with tradable intermediate-input industries, z_{xit}^{IV}, and the ratio between the intermediate-input price and the 24-month lagged regional CPI, p_{xit}^{IV}, denote the main independent variables. All the specifications include MSA fixed effects (allowed to shift across the pre-COVID, COVID, and post-COVID periods) and year-quarter fixed effects. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5 Results

5.1 The Post-COVID Steepening of the Phillips Curve

Table 2 summarizes our main result, revealing a significant steepening of the Phillips curve in both the post-COVID inflation and disinflation periods. Columns (1) to (4) use headline inflation as the dependent variable. In line with the extensive literature documenting the flattening of the Phillips curve from the 1990s, columns (1) and (2) show that the slope of the Phillips curve is small and not statistically different from zero in the pre-COVID and COVID periods. In contrast, our estimate of ψ during the post-COVID inflation phase increases substantially to 1.41 (column 3). The estimated coefficient remains statistically significant in the disinflation phase despite dropping to 0.74 (column 4). We further document that the relative intermediate-input price term significantly affects inflation across the four periods, particularly so starting from the onset of the pandemic. This captures the impact of higher commodity and intermediate-input prices on inflation due to post-pandemic supply chain disruptions and global demand recovery.

Given the volatility of headline inflation, central banks often focus on core inflation as a gauge of more persistent inflation dynamics.[12] Columns (5) to (8) confirm that our results hold when using core inflation as the dependent variable. Notably, our estimate of ψ in the post-COVID inflation period is almost identical to the one in column (3), while it becomes more precise and slightly higher in the disinflation phase. Interestingly, the impact of the relative intermediate-input price term on core inflation is substantially lower than the one on headline inflation during the pandemic and its aftermath. In particular, the effect is small and not statistically different from zero in the inflation surge phase. This result implies that the larger impact of the relative intermediate-input price term on headline inflation since the onset of the pandemic is mostly channeled through its effects on consumer food and energy prices. This is not surprising given that traded intermediate inputs account for a large share of expenditures and price volatility in the retail food and energy sectors.

Table 2: IV estimates of Equation (27)

	π_{it}: headline				π_{it}: core			
	(1) Pre	(2) COVID	(3) Post: inflation	(4) Post: disinflation	(5) Pre	(6) COVID	(7) Post: inflation	(8) Post: disinflation
u_{it}	-0.07	0.02	-1.41***	-0.74*	-0.03	0.00	-1.38***	-0.89***
	(0.09)	(0.06)	(0.28)	(0.38)	(0.10)	(0.05)	(0.27)	(0.32)
p_{xit}	0.09***	0.34***	0.18***	0.41***	0.08***	0.14*	0.06	0.20***
	(0.01)	(0.08)	(0.05)	(0.05)	(0.01)	(0.07)	(0.05)	(0.06)
Observations	4855	252	336	378	4855	252	336	378
MSA-Period FE	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times$ $\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓	✓	✓

Notes. This table presents estimates of ψ from Equation (27) for the pre-COVID (i.e., from January 1990 to February 2020), COVID (i.e., from March 2020 to February 2021), post-COVID inflation (i.e., from March 2021 to June 2022), and post-COVID disinflation (i.e., from July 2022 to December 2023) periods. Columns from (1) to (4) use 12-month headline inflation as the dependent variable, while columns from (5) to (8) use 12-month core inflation as the dependent variable. All columns feature the unemployment rate as the main independent variable. All specifications control for MSA fixed effects (allowed to shift across the pre-COVID, COVID, and post-COVID periods), year-quarter fixed effects, and intermediate-input prices relative to headline CPI. All columns display IV estimates of ψ obtained by instrumenting u_{it} with the shift-share instrument z^{IV}_{xit} and p_{xit} with p^{IV}_{xit}. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Heterogeneity of Phillips Curve Steepening Across Items

Our main result in Table 2 documents that the sensitivity of inflation to the unemployment rate, the slope of the Phillips curve, increased in the post-COVID period. However, post-pandemic US inflation

[12] Core inflation is defined as headline inflation excluding food and energy.

dynamics differed markedly between goods and services. The left panel of Figure A.3 in Appendix A plots the 12-month inflation rates for goods and services from January 2018 to December 2023, showing that the inflation surge is driven by goods more than services. Part of these dynamics could be explained by a more pronounced increase in the sensitivity of goods inflation relative to services inflation to the post-pandemic decrease in unemployment. Therefore, a natural extension of our analysis requires to investigate to what extent the post-COVID increase in the slope of the Phillips curve is driven by goods and services, respectively.

Table 3 shows that the goods Phillips curve steepened more than the services one in the post-COVID inflation period. Contrarily, the goods Phillips curve flattened in the post-COVID disinflation period, while the services Phillips curve steepened further. Columns (1) to (4) document the spike in the slope of the goods Phillips curve from 0.08 to 2.52 occurred in the inflation phase. Columns (5) to (8) show the more gradual and persistent steepening in the services Phillips curve with the slope increasing from 0.08 in the pre-COVID period to 0.85 during the inflation phase and reaching 1.35 in the disinflation phase. Relative price shocks in the intermediate-input markets have a large and volatile impact on goods inflation, while they have a limited and stable effect on services inflation.

Note that local unemployment has a large and significant impact on goods' prices in the post-COVID inflation phase. This is consistent with the failure of the law of one price for consumers' goods and provides empirical ground to the assumption of our model that final goods are non-tradable. This follows from the intuition that, even for final goods with a large share of tradable intermediate inputs, prices are highly responsive to local economic conditions, due to either fluctuations in the local component of real marginal costs or endogenous pricing responses. Interestingly, we estimate a null coefficient on the unemployment rate term and a large and significant coefficient on the relative intermediate-input price term of the goods Phillips curve in the post-COVID disinflation phase. This implies that the share of headline disinflation attributable to disinflation in goods is solely driven by negative shocks to relative intermediate-input prices.

Table 3: IV estimates of Equation (27) for goods and services inflation

	π_{it}: goods				π_{it}: services			
	(1) Pre	(2) COVID	(3) Post: inflation	(4) Post: disinflation	(5) Pre	(6) COVID	(7) Post: inflation	(8) Post: disinflation
u_{it}	-0.08	-0.07	-2.52***	0.11	-0.08	0.06	-0.85***	-1.35***
	(0.13)	(0.08)	(0.49)	(0.58)	(0.11)	(0.05)	(0.27)	(0.39)
p_{xit}	0.09***	0.67***	0.23**	0.81***	0.09***	0.12*	0.15***	0.15**
	(0.02)	(0.15)	(0.09)	(0.07)	(0.02)	(0.07)	(0.05)	(0.06)
Observations	4855	252	336	378	4855	252	336	378
MSA-Period FE	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times$ $\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓	✓	✓

Notes. This table presents estimates of ψ from Equation (27) for the pre-COVID (i.e., from January 1990 to February 2020), COVID (i.e., from March 2020 to February 2021), post-COVID inflation (i.e., from March 2021 to June 2022), and post-COVID disinflation (i.e., from July 2022 to December 2023) periods. Columns from (1) to (4) use 12-month goods inflation as the dependent variable, while columns from (5) to (8) use 12-month services inflation as the dependent variable. All columns feature the unemployment rate as the main independent variable. All specifications control for MSA fixed effects (allowed to shift across the pre-COVID, COVID, and post-COVID periods), year-quarter fixed effects, and intermediate-input prices relative to headline CPI. All columns display IV estimates of ψ obtained by instrumenting u_{it} with the shift-share instrument z^{IV}_{xit} and p_{xit} with p^{IV}_{xit}. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Since the shelter component of services inflation was more persistent than other components (Figure A.3, right panel) and accounts for a large share of services expenses in the consumer's basket, we further show that the steepening in the services Phillips curve is robust to the exclusion of shelter. Table 4 displays the estimated coefficients for services less shelter and shelter inflation, respectively. Both the services less shelter and the shelter Phillips curves mirror the gradual steepening of the services one, although the coefficient on shelter inflation is imprecisely estimated in the post-COVID disinflation period. The coefficient on the unemployment rate term in column (6) captures a weak positive relationship between housing prices and unemployment during COVID. Mondragon and Wieland (2022) show that increases in housing demand due to remote work pushed housing prices up during the pandemic and document that MSA-level exposure to remote work was positively correlated with the rise in unemployment across MSAs during COVID.

Table 4: IV estimates of Equation (27) for services less shelter and shelter inflation

	π_{it}: services less shelter				π_{it}: shelter			
	(1) Pre	(2) COVID	(3) Post: inflation	(4) Post: disinflation	(5) Pre	(6) COVID	(7) Post: inflation	(8) Post: disinflation
u_{it}	-0.10	-0.02	-0.77**	-1.53**	-0.00	0.12*	-0.82***	-1.65
	(0.17)	(0.08)	(0.37)	(0.64)	(0.25)	(0.06)	(0.31)	(1.01)
p_{xit}	-0.01	0.05	0.18***	0.00	0.22***	0.22*	0.12**	0.24**
	(0.02)	(0.11)	(0.06)	(0.09)	(0.03)	(0.12)	(0.06)	(0.11)
Observations	4855	252	336	378	4639	252	336	378
MSA-Period FE	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times$ $\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓	✓	✓

Notes. This table presents estimates of ψ from Equation (27) for the pre-COVID (i.e., from January 1990 to February 2020), COVID (i.e., from March 2020 to February 2021), post-COVID inflation (i.e., from March 2021 to June 2022), and post-COVID disinflation (i.e., from July 2022 to December 2023) periods. Columns from (1) to (4) use 12-month services less shelter inflation as the dependent variable, while columns from (5) to (8) use 12-month shelter inflation as the dependent variable. All columns feature the unemployment rate as the main independent variable. All specifications control for MSA fixed effects (allowed to shift across the pre-COVID, COVID, and post-COVID periods), year-quarter fixed effects, and intermediate-input prices relative to headline CPI. All columns display IV estimates of ψ obtained by instrumenting u_{it} with the shift-share instrument z_{xit}^{IV} and p_{xit} with p_{xit}^{IV}. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Contribution of the Phillips Curve Steepening to Post-COVID Inflation

The model tells us how our estimates for ψ, capturing the sensitivity of inflation to the present discounted value of current and future unemployment, relate to κ, which represents the sensitivity of inflation to current unemployment. Specifically,

$$\psi = \frac{\kappa}{1 - \beta\rho_u},$$

where β is the discount factor and ρ_u denotes the AR(1) autocorrelation coefficient of the unemployment rate. If we know the values of β and ρ_u and how they change across periods, we can evaluate to what extent the steepening of the Phillips curve that we estimate in Table 2 is driven by a rise in the sensitivity of inflation to current unemployment. We calibrate β to be equal to 0.99 and assume it remains constant before and after COVID. We estimate ρ_u in our sample and allow it to vary across periods. We find that the persistency of the unemployment rate drops from 0.91 in the pre-COVID period to 0.77 in the inflation phase of the post-COVID period.

Table 5: Contribution of the unemployment gap for different values of κ to post-COVID inflation

	κ	(1) % contribution of $\hat{u}_t - \hat{u}_t^{n,CBO}$	(2) % contribution of $\hat{u}_t - \hat{u}_t^{n,MS}$
Rotemberg and Woodford (1997)	0.019	0.73	0.49
Galí (2015)	0.085	3.29	2.19
Hazell, Herreno, Nakamura and Steinsson (2022)	0.0062	0.24	0.16
Our pre-COVID slope	0.0069	0.27	0.18
Our post-COVID slope	0.335	12.98	8.61
	ψ		
Our pre-COVID slope	0.07	2.71	1.80
Our post-COVID slope	1.41	54.63	36.25

Notes. This table presents the quantification of the contribution of the unemployment gap dynamics to the post-COVID inflation surge for different values of the slope of the Phillips curve. The first panel reports values of κ taken from Rotemberg and Woodford (1997), Galí (2015), Hazell et al. (2022), and our pre-COVID and post-COVID inflation phase estimates of ψ from Table 2, columns (1) and (3), multiplied by $(1 - \beta\rho_u)$. The slopes from Rotemberg and Woodford (1997) and Galí (2015) are adjusted for the elasticity of output with respect to employment. The second panel reports our pre-COVID and post-COVID inflation phase estimates of ψ from Table 2, columns (1) and (3). In column (1) the natural unemployment rate is equal to the estimates of the Congressional Budget Office (CBO), while in column (2) it is constructed as $u_t^n = \sqrt{u_t v_t}$, where u_t denotes the unemployment rate and v_t denotes the vacancy rate (Michaillat and Saez, 2022). To obtain the values in columns (1) and (2), we proceed as follows. First, we compute the difference between the 12-month headline inflation rate in June 2022 and March 2021 ($\Delta\pi_t$). Second, we compute the difference between the unemployment gap in June 2022 and March 2021 and we multiply it by the slope of the Phillips curve ($\kappa \times \Delta u_t^n$). Finally, we compute the ratio between $\Delta\pi_t$ and $\kappa \times \Delta u_t^n$ and multiply by 100.

Table 5 shows that the substantial increase in ψ that we estimate in the aftermath of the pandemic is driven by a significant rise in κ. The value of κ we obtain in the pre-COVID sample by combining our estimate of ψ from column 1 of Table 2 and the calibration of β and ρ_u is 0.0069. This value is line with the estimate of κ provided by Hazell et al. (2022) over a similar period – their sample goes from 1987 to 2017. κ increases to 0.335 in the inflation phase of the post-COVID period. The increase in κ from before to after COVID is larger than the increase in ψ because the persistency of unemployment, ρ_u, decreased between the two periods. We conclude that the steepening of the Phillips curve documented by ψ is a lower bound of the steepening implied by κ.

Table 5 illustrates another important result: absent the increase in the slope of the Phillips curve, the contribution of the unemployment gap to the post-COVID surge of inflation at the aggregate level would have been close to zero. To make this point, we quantify how much of the 6.4 percentage point increase in US headline inflation from March 2021 to June 2022 is due to the decrease in the unemployment gap

with and without the steepening of the Phillips curve[13]. To measure the unemployment gap, we use two different estimates of the natural rate of unemployment. In column 1, we use the natural rate of unemployment estimated by the Congressional Budget Office (CBO). In column 2, we follow Michaillat and Saez (2022) and measure the natural rate of unemployment as: $\hat{u}_t^n = \sqrt{u_t v_t}$, where u_t and v_t denote, respectively, the unemployment and vacancy rates. During the inflation phase of the post-COVID period, the unemployment gap decreases by 2.74 percentage points in column 1 and 1.64 percentage points in column 2.

Table 5 shows that if the Phillips curve was flat as documented by Hazell et al. (2022) and signaled by our pre-COVID slope, the evolution of the unemployment gap would have contributed to the post-COVID inflation surge only for a share ranging from 0.16% to 0.27%. The size of the contribution does not change significantly if we use estimates of the slope of the Phillips curve obtained through different methodologies[14]. Rotemberg and Woodford (1997) use structural estimation and show that κ is equal to 0.019. Galí (2015) estimate the slope of the Phillips curve to be 0.085 by calibrating a New Keynesian model. As these values of κ are slightly larger, the contribution of the decrease of the unemployment gap to the post-COVID inflation surge increases up to 3.29%. However, only the sizable steepening of the Phillips curve that we document in the inflation phase of the post-COVID period increases significantly the size of the quantified contribution. Table 5 also reports the contribution of the evolution of the unemployment gap to the post-COVID inflation surge using our estimates of ψ from before and after COVID. This quantification includes the effect of future expected unemployment gaps on inflation. The main result is invariant. However, the unemployment gap and the endogenous response of expectations now explain from 36.3% up to 54.6% of the post-COVID inflation surge.

6 Robustness Checks

In this section, we show that our main results are robust to threats to identification stemming from different model assumptions and empirical challenges such as measurement, potential omitted variables, small sample bias, and clustering.

[13] To make this quantification, we rely on the theoretical result that the slope of the regional and aggregate Phillips curve coincide. In the paper, we discuss the theoretical assumptions driving this result and argue that if we relax these assumptions, the aggregate slope would be higher than the regional one. Therefore, the quantification presented in Table 5 computed using our estimates are to be considered a lower bound.

[14] We follow Hazell et al. (2022) and adjust the estimates of κ obtained using these different methodologies by the elasticity of output with respect to employment.

6.1 Inflation Expectations

An important model assumption regarding inflation expectations allows identification of ψ. Regional unemployment depends on a transitory regional component and a permanent aggregate component consistent with the long-run inflation target set by the monetary authority. As a consequence, long-run inflation expectations affect current inflation, are common across regions, and can be differenced out by time fixed effects. We relax this assumption allowing the permanent component of regional unemployment to vary across regions. As a consequence, long-run inflation expectations are region-specific and cannot be differenced out by time fixed effects. Therefore, we estimate ψ in Equation (26) controlling for a proxy of long-run regional inflation expectations.

To perform this exercise, we need to construct a proxy of MSA-level long-run inflation expectations. To do so, we follow Binder et al. (2024) who show that inflation expectations of Republicans and Democrats as measured by the Michigan Survey of Consumers dramatically diverged during and after the pandemic. Therefore, we construct a shift-share variable interacting time-series variation in average 5-year-ahead inflation expectations of Republicans and Democrats with the 2016 MSA-level Republican and Democratic vote shares. Formally:

$$E_{it}\pi_{t+5} = \sum_{p\in\{R,D,I\}} \frac{v_{ip2016}}{\sum_{p\in\{R,D,I\}} v_{ip2016}} E_{pt}\pi_{t+5}, \tag{29}$$

where the subscript p denotes a political party (Republican, R, Democratic, D, or Independent, I), and v_{ip2016} represents the votes in MSA i for party p in the 2016 Presidential Elections. $E_{pt}\pi_{t+5}$ is the average inflation expectations by party affiliation measured by the Michigan Survey of Consumers from 2017 onward.

Table A.1 shows that the increase in ψ occurred in the aftermath of the pandemic is robust to controlling for the proxy of regional long-run inflation expectations proposed by Binder et al. (2024). Notably, $E_{it}\pi_{t+5}$ positively affects π_{it} in the post-COVID inflation period, thus capturing the role of long-run inflation expectations in driving inflation dynamics. Relative to our baseline specification, the size of ψ decreases from 1.41 to 1.18, implying that long-run inflation expectations in the post-COVID inflation period have a local dimension that correlates negatively with the unemployment rate. Interestingly, in the pre-COVID period the estimated coefficient on $E_{it}\pi_{t+5}$ is 0.91 and ψ does not change with respect to our baseline estimate. These results are consistent with the argument developed by Hazell et al. (2022) that before COVID long-run inflation expectations affect current inflation one-to-one and are not correlated

with local economic conditions.

6.2 Wealth Effects on the Labor Supply

In our model, the assumption that households preferences over consumption and labor take a GHH form implies no wealth effects on the optimal labor supply. This assumption is supported by recent empirical evidence provided by Ganong et al. (2024) who show that the large increase in unemployment benefits occurred during the pandemic resulted in higher consumption without affecting employment dynamics. If we relax this assumption and instead assume separable preferences over consumption and labor, we obtain the following expression for the regional Phillips curve:

$$\pi_{it} = E_t \bar{\pi}_{t,t+\infty} - \psi \tilde{u}_{it} + \mu \hat{c}_{it} + \delta \hat{p}_{xit} + \nu_{it}, \tag{30}$$

where $\hat{c}_{it}$ denotes deviations of regional consumption from the steady state and $\mu = \frac{\lambda \sigma^{-1}}{1-\beta\rho_c}$.[15] To the extent that $\hat{c}_{it}$ correlates with $\tilde{u}_{it}$, our estimate of ψ would be biased.

To address this concern, we add as a control to our main specification a proxy for consumption. Specifically, we construct an MSA-level wealth shock measure, τ_i, defined as the cumulative transfers received from the central government by residents of MSA i over the 2019-2021 period divided by MSA-level income in 2019, and interact it with time fixed effects, ξ_t. Transfers include economic impact payments (EIP) to eligible households, expanded child tax credits, and expanded unemployment benefits. MSA-level annual measures of each of these transfers and income are made publicly available by the IRS. The average transfer-induced wealth shock in our sample is 11.3% of MSA-level income with a minimum of 6.6% in San Francisco-Oakland-Hayward (CA) and a maximum 21.6% in Riverside-San Bernardino-Ontario (CA).

Table A.2 shows that the post-COVID Phillips curve steepening is robust to the inclusion of a control for post-pandemic MSA-level wealth shocks, $\tau_i \times \xi_t$. Since unemployment benefits could mechanically correlate with unemployment, we report the estimated coefficients when the wealth shocks τ_i are constructed excluding or including unemployment benefits. Consistently with the findings in Ganong et al. (2024), our baseline estimates are not sensitive to this control, irrespective of the inclusion of unemployment benefits in the wealth shock measure. This result implies that our baseline estimates are not biased and the slopes of the regional and aggregate Phillips curves do not differ because of wealth effects on the

[15] Here we assume that $\hat{c}_{it}$ follows an AR(1) process with an autocorrelation coefficient of ρ_c.

labor supply.

6.3 Alternative Measures and Instruments for Relative Intermediate-Input Prices

In our main specification, we measure the relative price of intermediate inputs as the ratio between the US PPI in the manufacturing sector and the contemporary MSA-level headline CPI and instrument it with the same ratio with the 2-year lagged MSA-level headline CPI in the denominator. Although the price of manufacturing products is the natural proxy for intermediate-input prices, we test the robustness of our results to different measures of the relative intermediate-input price and alternative instrumental variables. We perform several checks. First, we expand the notion of intermediate inputs including all commodities present in the PPI, as opposed to manufacturing commodities only. Second, we capture fluctuations in intermediate-input prices due to trade and global supply chain disruptions by defining p_{xit} as the relative import price index or instrumenting the US manufacturing PPI with the New York FED Global Supply Chain Pressure Index (GSCPI). Finally, since fluctuations in energy prices induced significant volatility in global commodity prices – especially after the Russian invasion of Ukraine – we control for the relative price of energy in the CPI and instrument the US manufacturing PPI with the West Texas Intermediate (WTI) crude oil price.

Table A.3 shows that the post-COVID steepening of the Phillips curve is robust to alternative measures of the relative intermediate-input price. Specifically, in the left panel p_{xit} is defined as the relative US all-commodities PPI, while in the right panel p_{xit} denotes the relative import price index (MPI) as measured by BLS. The estimated coefficients in the left panel are very close to the ones in the left panel of Table 2, displaying similar dynamics for the effects of both the unemployment rate and the relative intermediate-input price on inflation. The right panel displays a significant steepening in the post-COVID inflation phase, while the coefficient on the unemployment rate in the post-COVID disinflation phase becomes smaller and statistically non-significant. Interestingly, the relative price of imports has a large and significant impact on inflation only in the COVID and the post-COVID disinflation periods.

Table A.4 documents that our results do not change when employing alternative instrumental variables for the relative intermediate-input price term. In the left panel, the US manufacturing PPI is instrumented by the WTI crude oil price index to capture fluctuations in intermediate-input prices induced by oil price shocks. The estimated slopes of the Phillips curve are similar to the ones in Table 2 across all time periods, while the coefficient on the relative intermediate-input price term is stable at 0.16 in the pre-COVID, COVID, and post-COVID inflation period and doubles to 0.33 in the post-COVID disinflation

period. In the right panel, we instrument the US manufacturing PPI with the New York Fed GSCP index to isolate changes in intermediate-input prices due to trade and global supply chain disruptions. The coefficients on the unemployment rate term display similar dynamics to the ones in Table 2 with the exception of the pre-COVID period, in which the estimated slope becomes statistically significant. Notably, the estimated coefficients on the relative US manufacturing PPI suggest that the sensitivity of inflation to fluctuations in the relative intermediate-input price induced by global supply chain pressures has progressively increased from the start of the pandemic onward.

6.4 Retail Scanner Data: External Validity, Small Sample Bias, and Clustering

Our main results are based on a relatively small sample of MSAs for which the BLS computes the CPI at a monthly or by-monthly frequency. This could raise some concerns regarding the external validity of our results and their robustness to small sample bias and appropriate clustering of standard errors. To address these concerns, we leverage Nielsen Retail Scanner data and construct a monthly MSA-level retail price index following the methodology of Beraja et al. (2019). This allows us to expand our sample to a total of 851 MSAs.

Figure A.2 plots Nielsen inflation and unemployment across the 21 MSAs included in our main sample. The descriptive relationship between the two variables looks similar to the one in Figure 1, suggesting that our measure of retail inflation captures well MSA-level headline inflation dynamics. Since the Nielsen Retail Scanner dataset disproportionally represents items in the food category of the CPI, we quantitatively assess the accuracy of our measure by regressing 12-month CPI food inflation on 12-month Nielsen inflation in our main sample of 21 MSAs. We estimate a coefficient of 0.999 (0.009), which makes us confident that the Nielsen price index can be used to perform our desired robustness checks.[16]

Table A.5 shows the results from estimating Equation (27) using food inflation (left panel) and Nielsen inflation (right panel) as outcome variables, respectively. The pattern of estimated coefficients on the right panel tracks the pattern obtained in Table 2, with a pronounced steepening in the post-COVID inflation period. The results using Nielsen inflation as an outcome variable in our larger sample of MSAs confirm our main insights. The Nielsen Phillips curve starts steepening during COVID. The increase in the slope is more pronounced during the post-COVID inflation period, but the coefficient remains high also in the post-COVID disinflation phase. We conclude that our findings are externally valid and hardly sensitive to small sample bias.

[16] When regressing 12-month headline inflation on 12-month Nielsen inflation, we obtain an estimated coefficient of 0.736 (0.010).

In our main specification, we leverage variation across 21 MSAs. The small number of potential clusters prevents us to accurately adjust our standard errors for serial correlation in the error terms.[17] Indeed, the standard errors reported in Table 2 are clustered at the MSA-year-quarter level. A larger sample size allows us to verify the robustness of our results to the appropriate clustering of standard errors. Table A.6 shows our estimates of the slope of the Nielsen Phillips curve with standard errors clustered at the MSA-level (left panel) and with two-way clustering at the MSA and year-quarter level (right panel). The estimated coefficients and standard errors indicate that the statistical significance of our results is robust to serial correlation of the error terms.

7 Mechanisms

In this section, we evaluate two potential drivers of the steepening of the Phillips curve documented in our empirical analysis. According to our model, the slope of the Phillips curve is

$$\psi = \frac{\gamma(1-\alpha\beta)(1-\alpha)}{\phi\alpha(1-\beta\rho_u)}. \tag{31}$$

For given values of the discount factor (β), the output elasticity to labor in the final-goods sector (γ), and the autocorrelation coefficient of the unemployment rate (ρ_u), the slope of the Phillips curve depends on the Frisch elasticity of labor supply (ϕ) and the frequency of price change ($1-\alpha$). On the one hand, ψ is decreasing in ϕ capturing the intuition that, if the labor supply becomes more inelastic, demand shocks cause a relatively stronger response of real wages than employment. As a consequence, relatively higher real marginal costs translate into higher prices and inflation. On the other hand, ψ is increasing in $(1-\alpha)$, the frequency of price change. Indeed, if a larger fraction of firms change prices in each period, the response of prices and inflation to a demand shock is more pronounced.

7.1 Estimation of the labor supply elasticity

An important potential driver of the Phillips curve steepening is changes in the Frisch labor supply elasticity, ϕ. Log-linearizing Equation (5), we obtain

$$\left(\hat{\frac{W_{it}}{P_{it}}}\right) = \hat{\chi}_{it} - \frac{1}{\phi}\hat{u}_{it}, \tag{32}$$

[17] See, for instance, Cameron et al. (2008), Angrist and Pischke (2009), and MacKinnon and Webb (2017) for discussions of over-rejection probability and number of clusters.

To estimate ϕ separately for the pre-COVID and post-COVID periods, we leverage quarterly variation in average earnings and unemployment across 385 MSAs. Specifically, we measure MSA-level average nominal monthly earnings, we adjust them by keeping the workforce composition fixed at 2019-Q4,[18] and deflate them by the Nielsen price index to obtain a measure of composition-adjusted average real monthly earnings. We obtain separate measures of composition-adjusted average real monthly earnings for all employees, stable employees (i.e., workers that were working in the past quarter and will work in the next quarter), and new hires (i.e., workers that were hires in the past quarter).

We estimate the following equation:

$$w_{it} = \eta_i + \delta_t - \xi u_{it} + \varepsilon_{it}, \tag{33}$$

where w_{it} denotes log composition-adjusted average real monthly earnings in MSA i and year-quarter t, α_i denotes MSA fixed effects, δ_t denotes year-quarter fixed effects, and u_{it} denotes the unemployment rate,[19] as in Equation (27). The estimated coefficient $\hat{\xi} = 1/\hat{\phi}$ corresponds to the inverse labor supply elasticity (i.e., the percentage change in real wages required to induce a 1% increase in the labor supply) and is allowed to shift across the pre-COVID and the COVID/post-COVID periods.

Properly estimating ξ poses one main challenge, analogous to the one we faced when estimating Equation (27). As Equation (32) shows, shocks to the disutility of labor $\hat{\chi}_{it}$ affect real wages and may correlate with unemployment, thus biasing the OLS estimate of ξ. Therefore, we want to isolate demand-driven fluctuations in MSA-level unemployment rates to identify movements of real wages and unemployment occurring along the labor supply curve, rather than caused by shifts in the labor supply curve. We proceed by instrumenting u_{it} with z^{IV}_{xit}, as in our main specification.

[18] The data allow to adjust for 2-digit NAICS industry and age bins or 2-digit NAICS industry and education bins.
[19] Technically, this always denotes $-\ln(1 - u_{it})$.

Table 6: IV estimates of Equation (33)

	w_{it}: All Workers		w_{it}: Stayers		w_{it}: New Hires	
	(1) Pre-COVID	(2) Post-COVID	(3) Pre-COVID	(4) Post-COVID	(5) Pre-COVID	(6) Post-COVID
u_{it}	-0.26**	-4.85***	-0.48***	-5.17***	-0.93***	-6.89***
	(0.12)	(0.60)	(0.14)	(0.63)	(0.19)	(0.75)
Observations	25025	25025	25025	25025	25025	25025
Units	385	385	385	385	385	385
MSA FE	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times$ $\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓

Notes. This table presents estimates of ξ from Equation (33) for the pre-COVID (i.e., from 2000-Q1 to 2020-Q1) and post-COVID (i.e., 2020-Q2 to 2023-Q4) periods. The dependent variable is log average real monthly earnings of all workers in columns from (1) and (2), log average real monthly earnings of stayers (i.e., workers who worked in the same firm the previous and the following quarter) in columns (3) and (4), and log average real monthly earnings of new hires in columns (5) and (6). All earnings measures have been adjusted for NAICS 2-digit industry and age bins composition, keeping the composition of the workforce fixed at 2019-Q4. All columns feature log of 1 minus the unemployment rate as the main independent variable. All specifications control for MSA fixed effects, year-quarter fixed effects, and intermediate-input prices relative to headline CPI. All columns display IV estimates of ξ obtained by instrumenting u_{it} with the shift-share instrument z_{xit}^{IV}. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6 reports our estimates of ξ, the inverse labor supply elasticity, in the pre-COVID and post-COVID periods for all workers, stayers, and new hires. The estimated coefficients indicate that for a given decrease in the unemployment rate, real wages respond much more substantially in the post-COVID period than in the pre-COVID period. This is true for all categories of workers and, particularly, for new hires. Our benchmark estimate of the inverse labor supply elasticity is 0.26 for the pre-COVID period and 4.85 for the COVID and post-COVID periods. These estimates imply that $\phi = 3.85$ in the pre-COVID period and $\phi = 0.21$ in the COVID and post-COVID periods, documenting a sharp decrease in the estimated labor supply elasticity.

These results indicate that the sensitivity of real wages to the unemployment rate dramatically increased after COVID, causing the New Keynesian Phillips curve to steepen. This likely happened because the labor supply elasticity of the marginal workers changed between the pre-COVID and the COVID/post-COVID periods. Mui and Schoefer (2025) document a non-constant asymmetric aggregate labor supply curve using survey evidence in the US and Germany, noting that existing models fail to match this feature. Specifically, they show that the aggregate labor supply exhibits short-run elasticities above 3 for small labor demand shocks and around 0.5 for large positive labor demand shocks. Our quasi-experimental estimates confirm their findings and suggest that unprecedented labor market tightness in the aftermath of the pandemic made the economy hit the steeper section of the labor supply curve, driving up the slope

of the Phillips curve.

7.2 Drivers of Phillips curve steepening

A second potential driver of the steepening of the Phillips curve is changes in the frequency of price adjustments. In our model, α denotes the probability of keeping the price fixed in a given period. Hence, with probability $1-\alpha$, the firm is able to adjust its price. Through the lenses of our model, we rationalize a change in the frequency of price adjustments as an exogenous change of the Calvo parameter α. A decrease in α corresponds to an increase in the share of firms changing prices in a given period. Following Blanco et al. (2024), we calibrate $(1-\alpha)$ to be equal to 25% in the pre-COVID period and 40% in the post-COVID inflation phase. Blanco et al. (2024) show that, on average, 25% of prices change in a given quarter with respect to the previous year in periods of low inflation, as the one experienced by the US from 1990 to 2019. After March 2021, the fraction of price changes raised, peaking to approximately 50%. We calibrate $1-\alpha$ to be equal to 40%, which is the average of the quarterly frequency of price adjustment during the post-COVID inflation phase.

Table 7 quantifies the contribution of the two potential mechanisms driving the steepening of the Phillips curve. To do so, we start by calibrating or estimating the parameters composing the slope of the Phillips curve, as illustrated in equation 31. We calibrate ϕ to be equal to the values implied by our estimates in Table 6 for the pre-COVID and the post-COVID inflation phase periods. As discussed in the previous paragraph, we calibrate $(1-\alpha)$ following Blanco et al. (2024). We estimate the autocorrelation coefficient of the unemployment rate by regressing the MSA-level unemployment rate on its one-period lag, controlling for MSA and year-quarter fixed effects. We find that ρ_u decreased from 0.91 before COVID to 0.77 in the post-COVID inflation phase. We calibrate ρ_u to be equal to the point estimates. Finally, we calibrate γ and β to be, respectively, 0.3 and 0.99 and we assume they do not change between the two periods of interest. In particular, we set the value of γ to be equal to the average ratio between employee compensations and total sales in the retail sector from 1998 to 2023.

Table 7: Drivers of the Phillips curve steepening

	Pre-COVID		Post-COVID		
	(1) Calibrated	(2) Estimated	(3) Calibrated	(4) Estimated	(5) % of steepening
ϕ	3.85	3.85	0.21	0.21	74.86
$(1-\alpha)$	0.25		0.4		9
ρ_u	0.91	0.91	0.77	0.77	-2.63
γ	0.3		0.3		
β	0.99		0.99		
ψ	-0.07	-0.07 [-0.16, 0.02]	-1.66	-1.41 [-1.69, -1.13]	

Notes. This table reports the contribution to the steepening of the Phillips curve between the pre-COVID and post-COVID periods of changes in the Frisch labor supply elasticity (ϕ), the frequency of price adjustments $(1-\alpha)$, and the autocorrelation coefficient of the unemployment rate (ρ_u). In columns 1 and 2, we report the calibrated and estimated values of the parameters in the pre-COVID period. In columns 3 and 4, we report the calibrated and estimated values of the parameters in the post-COVID inflation phase. The estimated values of ϕ are the ones implied by our estimates in Table 6 and we set the calibrated values equal to the point estimates. We calibrate the values of $(1-\alpha)$ following Blanco et al. (2024). We estimate ρ_u by regressing the MSA-level unemployment rate on its one-period lag, controlling for MSA and year-quarter fixed effects and we set the calibrated values equal to the point estimates. Finally, we assume that γ and β do not change between the pre-COVID and post-COVID periods and we calibrate them to be 0.3 and 0.99, respectively. γ is proxied by the average value of the employees' compensations to total sales in the retail sector from 1998 to 2023. Columns 1 and 3 in the second panel report the calibrated values of ψ according to equation 31, using the calibrated values of the parameters in the first panel. Columns 2 and 4 report the estimated values of ψ from Table 2 and the confidence intervals in the square brackets. Column 5 reports the percentage of the increase in ψ from column 1 to 3 that is attributable to the change in, respectively, ϕ, $(1-\alpha)$, and ρ_u. To compute these values, we proceed in two steps. First, we recalibrate ψ keeping all parameters fixed to their pre-COVID values and allowing one among ϕ, $(1-\alpha)$, and ρ_u to change to its post-COVID value. Second, we compute the percentage of the distance between ψ in columns 1 and 3 that is accounted for by the distance between ψ in column 1 and the recalibrated ψ obtained in the first step.

We use the values of the parameters shown in columns (1) and (3) of the upper panel of Table 7 to calibrate the values of ψ as predicted by the model (equation 31) in the pre-COVID and post-COVID inflation periods. The lower panel of Table 7 shows that the calibrated values of ψ in columns (1) and (3) are close and within the confidence intervals of the estimates of ψ from Table 2, reported in columns (2) and (4). To quantify the contribution of the decrease in the Frisch labor supply elasticity to the steepening of the Phillips curve, we recalibrate ψ allowing ϕ to change from column (1) to (3), while keeping all the other parameters fixed to their pre-COVID values reported in column (1). Then, we compute the relative value of the recalibrated ψ with respect to the value in column (3). We proceed in a similar way to quantify the contribution of the increase in the frequency of price adjustments, recalibrating ψ allowing only the parameter $(1-\alpha)$ to change.

Column (5) of Table 7 reports that the steepening of the Phillips curve between the pre-COVID

period and the post-COVID inflation phase is mostly explained by the decrease of the Frisch elasticity of labor supply over the same two periods. This channel drives almost 75% of the increase in the slope. As the labor supply has become more inelastic in the aftermath of the pandemic, demand shocks transmit into higher prices by inducing a stronger response of real wages with respect to employment. The sizable increase in the frequency of price changes is responsible for 9% of the steepening of the Phillips curve. We also quantify the contribution to the steepening of the decrease in the persistency of the unemployment rate after the pandemic. The decrease in ρ_u has a negative effect on the slope of the Phillips curve, but this counteracting force is small and negligible in magnitude.

8 Conclusion

In this paper, we estimate the slope of the Phillips curve before, during, and after COVID. To do so, we exploit MSA-level variation in inflation and unemployment combined with an instrumental variable approach. We relate our cross-sectional estimates to the aggregate parameter through a two-region New-Keynesian model of a monetary union. The model features labor supply shocks and a production side with a vertical supply chain (i.e., an international intermediate-input market and local final-goods sectors) to capture supply-side shocks relevant to post-COVID inflation dynamics. We derive the regional and aggregate Phillips curves, showing that their slopes coincide.

In our benchmark specification, we estimate that the Phillips curve substantially steepened after COVID. A steep Phillips curve is consistent with models featuring a low degree of nominal and real rigidities that likely provide a good representation of the US 2021-2023 economy. Our results imply that fluctuations in the unemployment gap explained up to 55% of the post-COVID increase in inflation. Not allowing the slope of the Phillips curve to change across the three periods makes the contribution of the unemployment gap small and statistically insignificant. Another important implication of a steeper Phillips curve is that the sacrifice ratio (i.e., the output cost associated with reducing inflation) is lower. Models of optimal monetary policy with a non-linear Phillips curve typically prescribe to "strike while the iron is hot" (Karadi et al., 2024).

Our results indeed point to the presence of non-linearities in the Phillips curve in the aftermath of the pandemic. In our model, the slope of the Phillips curve depends on the frequency of price change and the elasticity of labor supply, among other parameters. We show that an estimated structural change in the labor supply elasticity can account for a large part of the post-COVID New Keynesian

Phillips curve steepening. As pointed our by Mui and Schoefer (2025), current models struggle to capture nonlinearities in the aggregate labor supply, although some recent works make important steps in that direction (Benigno and Eggertsson, 2023; Gitti, 2023; Barnichon and Shapiro, 2024). This is an exciting path for future research.

References

Angrist, Joshua D and Jörn-Steffen Pischke, *Mostly harmless econometrics: An empiricist's companion*, Princeton university press, 2009.

Autor, David, Arindrajit Dube, and Annie McGrew, "The Unexpected Compression: Competition at Work in the Low Wage Labor Market," Technical Report, National Bureau of Economic Research 2023.

Babb, Nathan R. and Alan K. Detmeister, "Nonlinearities in the Phillips Curve for the United States : Evidence Using Metropolitan Data," Finance and Economics Discussion Series 2017-070, Board of Governors of the Federal Reserve System (U.S.) June 2017.

Ball, Laurence, Daniel Leigh, and Prachi Mishra, "Understanding US Inflation during the COVID Era," Working Paper 30613, NBER 2022.

Barnichon, Régis and Adam Hale Shapiro, "Phillips meets beveridge," *Journal of Monetary Economics*, 2024, *148*, 103660.

Bartik, Timothy J., "Who Benefits from State and Local Economic Development Policies?," WE Upjohn Institute for Employment Research 1991.

Beaudry, Paul, Chenyu Hou, and Franck Portier, "On the Fragility of the Nonlinear Phillips Curve View of Recent Inflation," Technical Report, National Bureau of Economic Research 2025.

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," Working Paper 31197, National Bureau of Economic Research April 2023.

Beraja, Martin, Erik Hurst, and Juan Ospina, "The Aggregate Implications of Regional Business Cycles," *Econometrica*, 2019, *87* (6), 1789–1833.

Bernanke, Ben and Olivier Blanchard, "What caused the US pandemic-era inflation?," *Peterson Institute for International Economics Working Paper*, 2023, (23-4).

Binder, Carola Conces, Rupal Kamdar, and Jane M Ryngaert, "Partisan expectations and COVID-era inflation," *Journal of Monetary Economics*, 2024, p. 103649.

Blanchard, Olivier J., Alex Domash, and Lawrence H. Summers, "Bad News for the Fed from the Beveridge Space," Peterson Institute for International Economics 2022.

Blanco, Andrés, Corina Boar, Callum J Jones, and Virgiliu Midrigan, "The inflation accelerator," Technical Report, National Bureau of Economic Research 2024.

Bohr, Clement E, "Capacity buffers: Explaining the retreat and return of the Phillips curve," *Manuscript, UCLA*, 2024.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-experimental Shift-Share Research Designs," *The Review of Economic Studies*, 2022, *89* (1), 181–213.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-based improvements for inference with clustered errors," *The review of economics and statistics*, 2008, *90* (3), 414–427.

Cavallo, Alberto, "Inflation with Covid Consumption Baskets," Working Paper 27352, NBER 2020.

_ , Francesco Lippi, and Ken Miyahara, "Large shocks travel fast," *American Economic Review: Insights*, 2024, *6* (4), 558–574.

Comin, Diego A, Robert C Johnson, and Callum J Jones, "Supply Chain Constraints and Inflation," Working Paper, NBER 2023.

Crump, Richard K., Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin, "The Unemployment-Inflation Trade-off Revisited: The Phillips Curve in COVID Times," Working Paper 29784, NBER 2022.

Crump, Richard K, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin, "The unemployment–inflation trade-off revisited: The Phillips curve in COVID times," *Journal of Monetary Economics*, 2024, *145*, 103580.

Erceg, Christopher, Jesper Lindé, and Mathias Trabandt, "Monetary Policy and Inflation Scares," *Unpublished, International Monetary Fund, Washington, DC*, 2024.

Fitzgerald, Terry J. and Juan Pablo Nicolini, "Is There a Stable Relationship between Unemployment and Future Inflation?," Federal Reserve Bank of Minneapolis Working Paper 2014.

Fitzgerald, Terry J, Callum Jones, Mariano Kulish, Juan Pablo Nicolini et al., "Is There a Stable Relationship between Unemployment and Future Inflation?," Technical Report, Federal Reserve Bank of Minneapolis 2020.

Fornaro, Luca, *Fiscal stimulus with supply constraints*, BSE, Barcelona School of Economics, 2024.

Galí, Jordi, *Monetary policy, Inflation, and the Business Cycle: an Introduction to the New Keynesian Framework and its Applications*, Princeton University Press, 2015.

Ganong, Peter, Fiona Greig, Pascal Noel, Daniel M Sullivan, and Joseph Vavra, "Spending and job-finding impacts of expanded unemployment benefits: Evidence from administrative micro data," *American Economic Review*, 2024, *114* (9), 2898–2939.

Gitti, Giulia, "Nonlinearities in the Regional Phillips Curve with Labor Market Tightness," Job Market Paper 2023.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman, "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 1988, pp. 402–417.

Harding, Martín, Jesper Lindé, and Mathias Trabandt, "Resolving the Missing Deflation Puzzle," *Journal of Monetary Economics*, 2022, *126*, 15–34.

Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jón Steinsson, "The Slope of the Phillips Curve: Evidence from US States," *The Quarterly Journal of Economics*, 2022, *137* (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* (2), 26–52.

Karadi, Peter, Anton Nakov, Galo Nuno, Ernesto Pasten, and Dominik Thaler, "Strike while the iron is hot: optimal monetary policy with a nonlinear Phillips curve," 2024.

Klenow, Peter J. and Oleksiy Kryvtsov, "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent US Inflation?," *The Quarterly Journal of Economics*, 2008, *123* (3), 863–904.

MacKinnon, James G and Matthew D Webb, "Wild bootstrap inference for wildly different cluster sizes," *Journal of Applied Econometrics*, 2017, *32* (2), 233–254.

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock, "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 2014, *52* (1), 124–88.

McLeay, Michael and Silvana Tenreyro, "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual*, 2020, *34* (1), 199–255.

Michaillat, Pascal and Emmanuel Saez, "u*=$\sqrt{uv}$," Working Paper 30211, NBER 2022.

Mondragon, John A. and Johannes Wieland, "Housing Demand and Remote Work," Working Paper 30041, NBER 2022.

Mui, Preston and Benjamin Schoefer, "Reservation Raises: The Aggregate Labour Supply Curve at the Extensive Margin," *Review of Economic Studies*, 2025, *92* (1), 442–475.

Nakamura, Emi and Jón Steinsson, "Five Facts about Prices: a Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.

Pfäuti, Oliver, "The inflation attention threshold and inflation surges," *arXiv preprint arXiv:2308.09480*, 2023, p. 2.

Reinsdorf, Marshall, "COVID-19 and the CPI: Is Inflation Underestimated?," Working Paper, IMF 2020.

Rotemberg, Julio J and Michael Woodford, "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 1997.

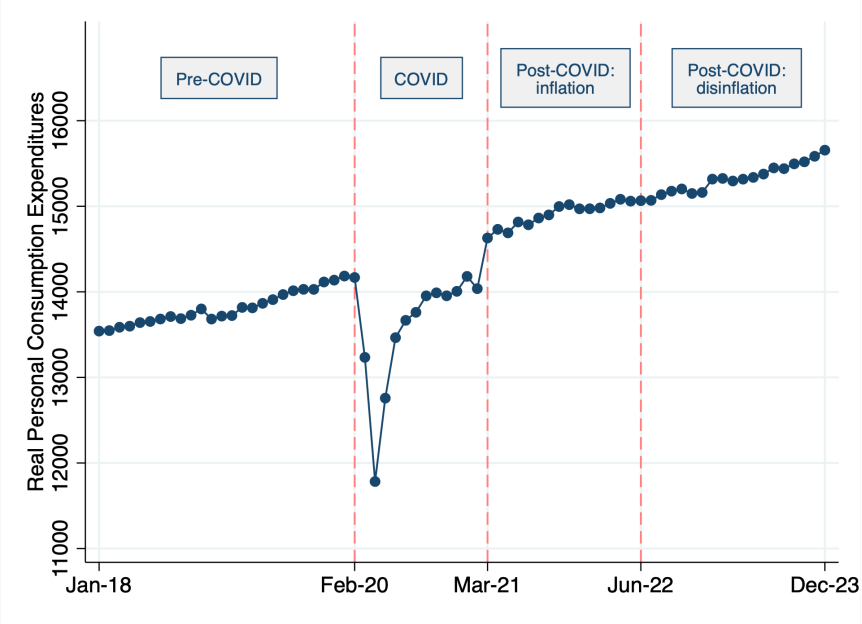
Schmitt-Grohé, Stephanie and Martın Uribe, "What Do Long Data Tell Us About the Inflation Hike Post COVID-19 Pandemic?," 2022.

Smith, Simon, Allan Timmermann, and Jonathan H Wright, "Breaks in the Phillips Curve: Evidence from Panel Data," Technical Report, National Bureau of Economic Research 2023.

Appendix

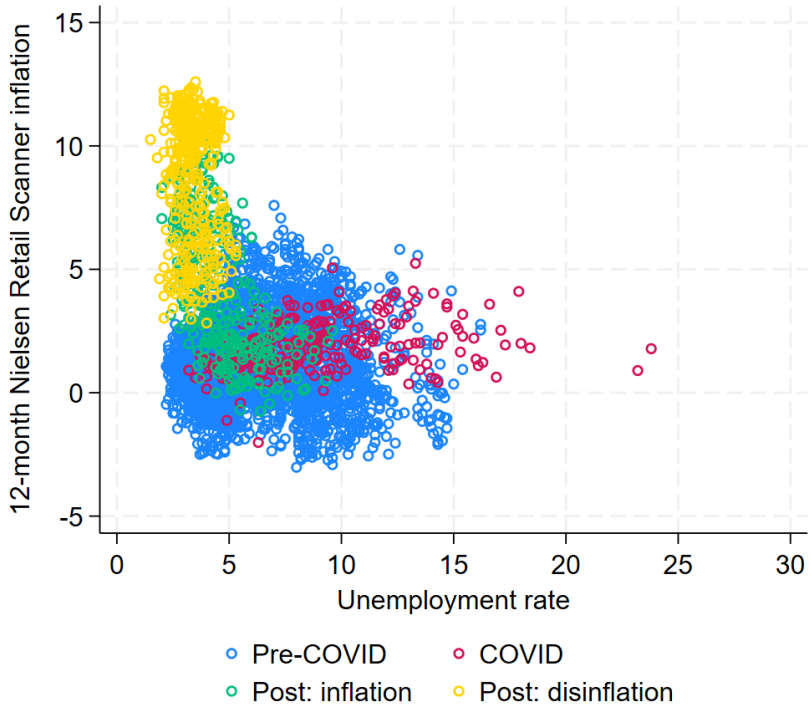
A Figures and Tables

Figure A.1: Real Personal Consumption Expenditures (Jan 2018-Dec 2023)



Notes. The figure shows the time series of US real personal consumption expenditures as made available by the BEA from January 2018 to December 2023. We set the start of the post-COVID period in March 2021, when the time series reverts to its pre-pandemic trend after the pandemic shock. The vertical red lines separate the pre-COVID, COVID, and post-COVID: inflation and post-COVID: disinflation periods according to our definition.

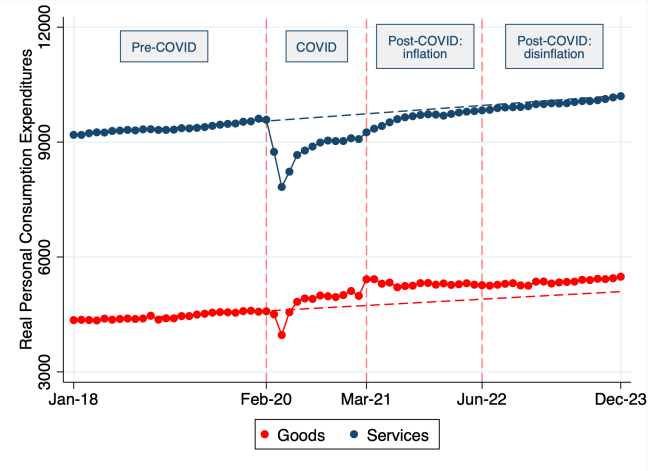
Figure A.2: The Phillips correlation across US Cities



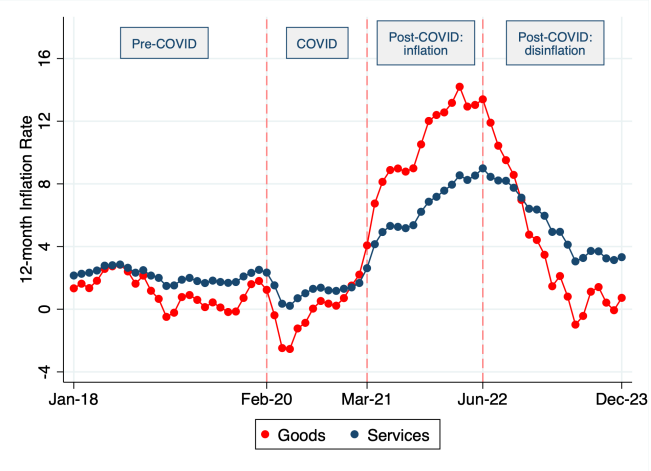
Notes. The scatter plot shows the relationship between the 12-month Nielsen Retail Scanner inflation rate and the unemployment rate for all MSAs included in the BLS sample (i.e., the MSAs for which the BLS computes the CPI every month or every two months). The blue dots denote observations belonging to the pre-COVID period (i.e., Jan 1990-Feb 2020), the green dots denote observations belonging to the COVID period (i.e., Mar 2020-Feb 2021), and the red dots denote observations belonging to the post-COVID period (i.e., Mar 2021-Sep 2022).

Figure A.3: Real Consumption Expenditures and Inflation, Goods vs. Services (Jan 2018-Dec 2023)

(a) Real Consumption Expenditures



(b) Inflation Rates



Notes. The figure shows the time series of US real consumption expenditures for goods vs. services (A.3a) from the BEA and the 12-month inflation rate for goods vs. services (A.3b) made available by the BLS from January 2018 to December 2023. We set the start of the post-COVID period in March 2021, when the time series of real personal consumption expenditures reverts to its pre-pandemic trend after the pandemic shock. The vertical red lines separate the pre-COVID, COVID, and post-COVID: inflation and post-COVID: disinflation periods according to our definition.

Table A.1: Robustness check with inflation expectations

	π_{it}: headline			
	(1) Pre	(2) COVID	(3) Post: inflation	(4) Post: disinflation
u_{it}	-0.09	0.04	-1.18***	-0.71*
	(0.11)	(0.07)	(0.26)	(0.36)
p_{xit}	0.12***	0.44***	0.17***	0.41***
	(0.04)	(0.13)	(0.05)	(0.05)
$E_{it}\pi_{t+5}$	0.91***	-0.52	0.49**	0.06
	(0.18)	(0.39)	(0.22)	(0.25)
Observations	1033	252	336	378
MSA-Period FE	✓	✓	✓	✓
Year-Quarter FE $\times$ $\Sigma_s e_{is}$	✓	✓	✓	✓

Notes. This table presents estimates of ψ from Equation (27) for the pre-COVID (i.e., from January 1990 to February 2020), COVID (i.e., from March 2020 to February 2021), post-COVID inflation (i.e., from March 2021 to June 2022), and post-COVID disinflation (i.e., from July 2022 to December 2023) periods. All columns use 12-month headline inflation as the dependent variable. All columns use the unemployment rate as the main independent variable. All specifications control for MSA fixed effects (allowed to shift across the pre-COVID, COVID, and post-COVID periods), year-quarter fixed effects interacted with the MSA-level sum of intermediate-input employment shares, intermediate-input prices relative to headline CPI, and our proxy for 5-year-ahead inflation expectations (Equation 29). All columns display IV estimates of ψ obtained by instrumenting u_{it} with the shift-share instrument z^x_{it} and $\hat{p}_{xit}$ with $\hat{p}^{IV}_{xit}$. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Robustness check with wealth effects

	π_{it}: headline; τ_i: without UB			π_{it}: headline; τ_i: with UB		
	(1) COVID	(2) Post: inflation	(3) Post: disinflation	(4) COVID	(5) Post: inflation	(6) Post: disinflation
u_{it}	0.00	-1.48***	-0.79**	0.03	-1.35***	-0.76**
	(0.05)	(0.27)	(0.36)	(0.06)	(0.26)	(0.37)
p_{xit}	0.28***	0.18***	0.37***	0.35***	0.19***	0.41***
	(0.07)	(0.05)	(0.05)	(0.08)	(0.05)	(0.05)
Observations	252	336	378	252	336	378
MSA-Period FE	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times$ $\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times$ τ_i	✓	✓	✓	✓	✓	✓

Notes. This table presents estimates of ψ from Equation (27) for the COVID (i.e., from March 2020 to February 2021), post-COVID inflation (i.e., from March 2021 to June 2022), and post-COVID disinflation (i.e., from July 2022 to December 2023) periods. All columns use 12-month headline inflation as the dependent variable. All specifications control for MSA fixed effects (allowed to shift across the COVID, and post-COVID periods), year-quarter fixed effects interacted with the MSA-level sum of intermediate-input employment shares, intermediate-input prices relative to headline CPI, and the 2020-2021 government transfer shock as a share of income (τ_i) interacted with year-quarter fixed effects. The left and right panels of the table include τ_i computed excluding and including unemployment benefits, respectively. All columns display IV estimates of ψ obtained by instrumenting u_{it} with the shift-share instrument z^x_{it} and $\hat{p}_{xit}$ with $\hat{p}^{IV}_{xit}$. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Robustness checks with relative all-commodities PPI and import price index

	π_{it}: headline, p_{xit}: all-commodities PPI				π_{it}: headline, p_{xit}: import price index			
	(1) Pre	(2) COVID	(3) Post: inflation	(4) Post: disinflation	(5) Pre	(6) COVID	(7) Post: inflation	(8) Post: disinflation
u_{it}	-0.07	0.03	-1.33***	-0.69*	-0.03	0.18	-1.59***	-0.43
	(0.09)	(0.06)	(0.27)	(0.37)	(0.09)	(0.11)	(0.29)	(0.39)
p_{xit}	0.09***	0.21***	0.18***	0.52***	0.09***	0.72***	0.14	0.68***
	(0.01)	(0.04)	(0.04)	(0.07)	(0.01)	(0.21)	(0.10)	(0.10)
Observations	4855	252	336	378	4855	252	336	378
MSA-Period FE	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times$ $\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓	✓	✓

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes. This table presents estimates of ψ from Equation (27) for the COVID (i.e., from March 2020 to February 2021), post-COVID inflation (i.e., from March 2021 to June 2022), and post-COVID disinflation (i.e., from July 2022 to December 2023) periods. All columns use 12-month headline inflation as the dependent variable. All specifications control for MSA fixed effects (allowed to shift across the COVID, and post-COVID periods), year-quarter fixed effects interacted with the MSA-level sum of intermediate-input employment shares, and intermediate-input prices relative to headline CPI. The left and right panels of the table measure intermediate-input prices through the all-commodities Producer Price Index (PPI) and the import price index (MPI) provided by the BLS, respectively. All columns display IV estimates of ψ obtained by instrumenting u_{it} with the shift-share instrument z^x_{it} and $\hat{p}_{xit}$ with $\hat{p}^{IV}_{xit}$. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4: Robustness checks with relative oil price and Global Supply Chain Pressure (GSCP) index as IV

	π_{it}: headline, p^{IV}_{xit}: oil price				π_{it}: headline, p^{IV}_{xit}: GSCP index			
	(1) Pre	(2) COVID	(3) Post: inflation	(4) Post: disinflation	(5) Pre	(6) COVID	(7) Post: inflation	(8) Post: disinflation
u_{it}	0.02	-0.03	-1.44***	-0.69*	-0.33***	-0.03	-1.30***	-1.07*
	(0.10)	(0.05)	(0.27)	(0.37)	(0.11)	(0.05)	(0.26)	(0.61)
p_{xit}	0.19***	0.16*	0.16***	0.33***	-0.10	0.17*	0.27***	0.93***
	(0.02)	(0.08)	(0.05)	(0.04)	(0.08)	(0.09)	(0.07)	(0.25)
Observations	4855	252	336	378	3703	252	336	378
MSA-Period FE	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times$ $\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓	✓	✓

Notes. This table presents estimates of ψ from Equation (27) for the COVID (i.e., from March 2020 to February 2021), post-COVID inflation (i.e., from March 2021 to June 2022), and post-COVID disinflation (i.e., from July 2022 to December 2023) periods. All columns use 12-month headline inflation as the dependent variable. All specifications control for MSA fixed effects (allowed to shift across the COVID, and post-COVID periods), year-quarter fixed effects interacted with the MSA-level sum of intermediate-input employment shares, and intermediate-input prices relative to headline CPI. All columns display IV estimates of ψ obtained by instrumenting u_{it} with the shift-share instrument z^x_{it} and $\hat{p}_{xit}$ with $\hat{p}^{IV}_{xit}$. The left and right panels of the table instrument intermediate-input prices through the relative price of oil and the Global Supply Chain Pressure (GSCP) Index provided by the New York Fed, respectively. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Robustness checks with food CPI vs. Nielsen inflation

	π_{it}: food				π_{it}: Nielsen			
	(1) Pre	(2) COVID	(3) Post: inflation	(4) Post: disinflation	(5) Pre	(6) COVID	(7) Post: inflation	(8) Post: disinflation
u_{it}	-0.00	-0.07	-2.01***	-0.58	-0.03	-0.29***	-2.67***	-1.79***
	(0.12)	(0.06)	(0.42)	(0.48)	(0.02)	(0.02)	(0.22)	(0.19)
p_{xit}	0.00	0.03	0.29***	0.27***	0.20***	-0.26***	0.08***	0.05***
	(0.02)	(0.16)	(0.08)	(0.06)	(0.01)	(0.03)	(0.02)	(0.02)
Observations	4399	252	336	378	113147	10212	13616	14670
Unique MSAs	21	21	21	21	851	851	851	851
MSA-Period FE	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓	✓	✓

Notes. This table presents estimates of ψ from Equation (27) for the pre-COVID (i.e., from January 1990 to February 2020), COVID (i.e., from March 2020 to February 2021), post-COVID inflation (i.e., from March 2021 to June 2022), and post-COVID disinflation (i.e., from July 2022 to December 2023) periods. Columns from (1) to (4) use 12-month food inflation measured for the 21 MSAs present in the BLS sample as the dependent variable, while columns from (5) to (8) use 12-month Nielsen inflation measured for the 851 MSAs present in our sample as the dependent variable. All columns feature the unemployment rate as the main independent variable. All specifications control for MSA fixed effects (allowed to shift across the pre-COVID, COVID, and post-COVID periods), year-quarter fixed effects, and intermediate-input prices relative to headline CPI. All columns display IV estimates of ψ obtained by instrumenting u_{it} with the shift-share instrument z^{IV}_{xit} and p_{xit} with p^{IV}_{xit}. Standard errors in parentheses are clustered at the MSA-year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Robustness checks with MSA and two-way clustered standard errors

	π_{it}: Nielsen				π_{it}: Nielsen			
	(1) Pre	(2) COVID	(3) Post: inflation	(4) Post: disinflation	(5) Pre	(6) COVID	(7) Post: inflation	(8) Post: disinflation
u_{it}	-0.03	-0.29***	-2.67***	-1.79***	-0.03	-0.29***	-2.67**	-1.79**
	(0.05)	(0.02)	(0.30)	(0.34)	(0.06)	(0.05)	(0.97)	(0.69)
p_{xit}	0.20***	-0.26***	0.08***	0.05*	0.20***	-0.26*	0.08	0.05
	(0.01)	(0.03)	(0.02)	(0.03)	(0.03)	(0.10)	(0.19)	(0.12)
Observations	113147	10212	13616	14670	113147	10212	13616	14670
Unique MSAs	851	851	851	851	851	851	851	851
MSA-Period FE	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter FE $\times\Sigma_s e_{is}$	✓	✓	✓	✓	✓	✓	✓	✓
Clusters: MSA	✓	✓	✓	✓				
Clusters: MSA + Year-Quarter					✓	✓	✓	✓

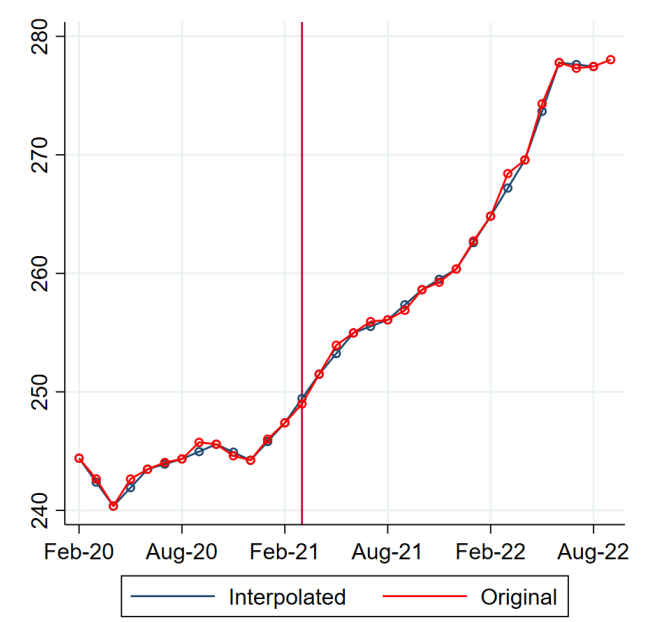
Notes. This table presents estimates of ψ from Equation (27) for the pre-COVID (i.e., from January 1990 to February 2020), COVID (i.e., from March 2020 to February 2021), post-COVID inflation (i.e., from March 2021 to June 2022), and post-COVID disinflation (i.e., from July 2022 to December 2023) periods. All columns feature the 12-month Nielsen price index growth rate as the dependent variable and the unemployment rate as the main independent variable. All specifications control for MSA fixed effects (allowed to shift across the pre-COVID, COVID, and post-COVID periods), year-quarter fixed effects interacted with the MSA-level sum of intermediate-input employment shares, and intermediate-input prices relative to headline CPI. All columns display IV estimates of ψ obtained by instrumenting u_{it} with the shift-share instrument z_{xit}^{IV} and p_{xit} with p_{xit}^{IV}. Standard errors in parentheses are clustered at the MSA level in the left panel and are two-way clustered (at the MSA and year-quarter levels) in the right panel. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Interpolation

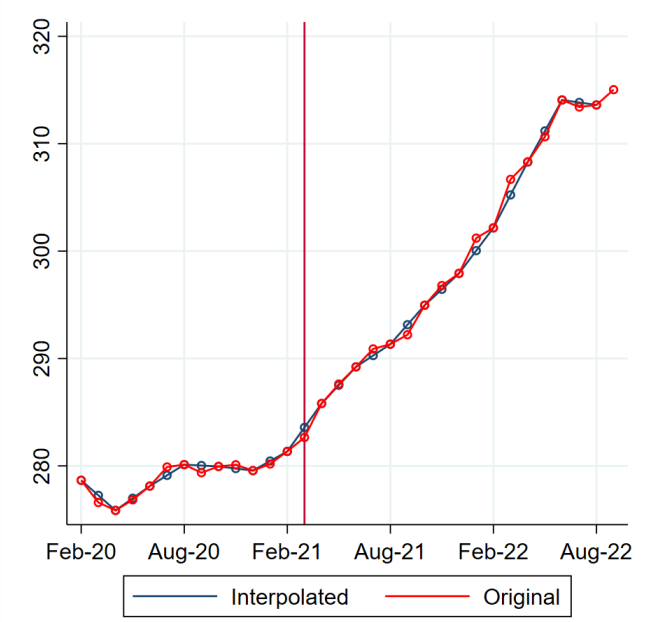
As discussed in Section 3, the CPI data we use from the BLS are made available at a monthly frequency for three metropolitan areas (i.e., Chicago-Naperville-Elgin, Los Angeles-Long Beach-Anaheim, and New York-Newark-Jersey City) and at a bi-monthly frequency for all other MSAs. We linearly interpolate the data for all MSAs for which the CPI is made available at a bi-monthly frequency. On the one hand, by doing so, we obtain a larger sample size, which greatly benefits statistical power for parameter estimation in the post-COVID period, characterized by a narrow time window. On the other hand, our imputed measures of the CPI are affected by measurement error. In this Appendix section, we discuss how large the measurement error we introduce is likely to be and address potential concerns about the correlation between imputed inflation values and our instrument.

Figure B.1: Original vs. Interpolated CPI series, Chicago and Los Angeles (Feb 2020-Sep 2022)

(a) Chicago-Naperville-Elgin (IL-IN-WI)



(b) Los Angeles-Long Beach-Anaheim (CA)



Notes. The figure compares original and linearly interpolated time series of the all-items CPI for Chicago-Naperville-Elgin (B.1a) and Los Angeles-Long Beach-Anaheim (B.1b) from February 2020 to September 2022. The blue line denotes the interpolated time series, while the red line denotes the original time series. The interpolated series are produced by imputing odd months observations. The vertical line indicates the beginning of the post-COVID period, starting in March 2021.

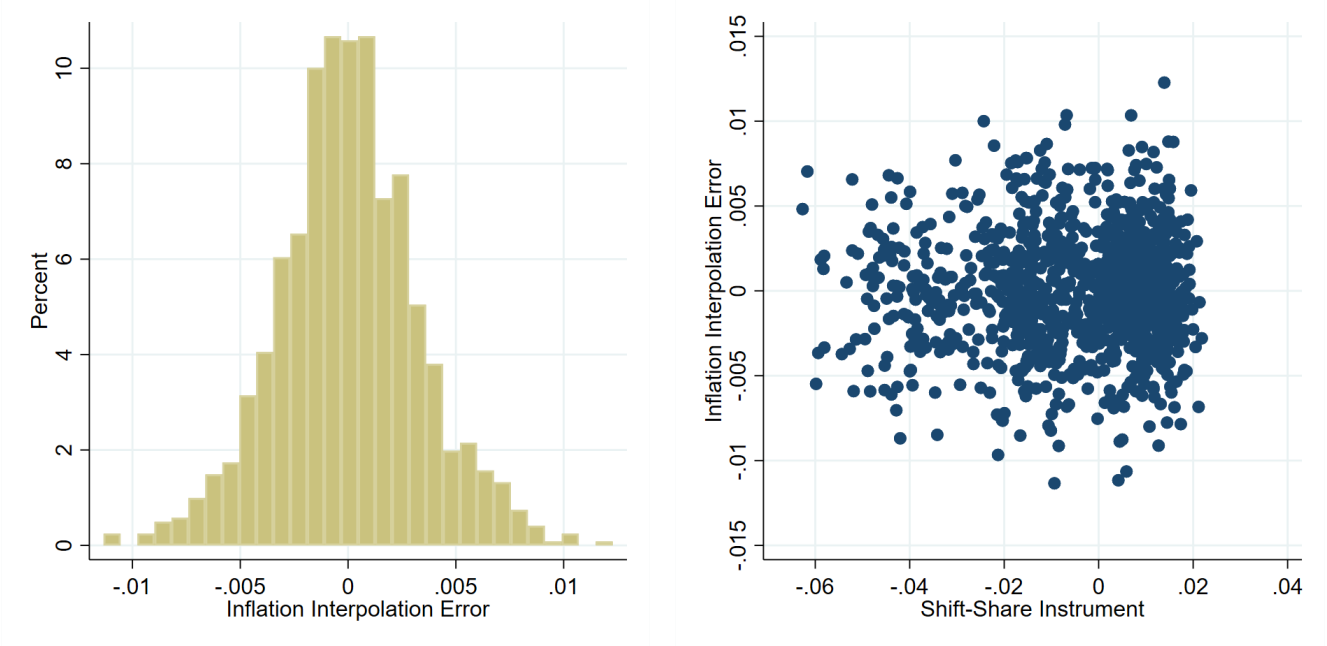
To evaluate how large the interpolation measurement error might be, we conduct the following exercise. For the three metropolitan areas for which CPI data are reported at a monthly frequency, we produce interpolated CPI time series, by declaring CPI observations missing in odd (or even) months

and imputing them through linear interpolation. Next, we compare the original CPI time series with the interpolated one, and compute measurement errors for imputed observations. Figure B.1 compares original and linearly interpolated time series of the CPI obtained by imputing observations in odd months for Chicago-Naperville-Elgin and Los Angeles-Long Beach-Anaheim from February 2020 to September 2022. The interpolated series match closely the original ones for both MSAs. Provided that measurement error is not systematically different for imputed observations in Chicago-Naperville-Elgin, Los Angeles-Long Beach-Anaheim, and New York-Newark-Jersey City, imputed observations for all other MSAs should closely approximate true values.

Figure B.2: Inflation Interpolation Error Distribution and Correlation with z_{it}^x

(a) Inflation Interpolation Error Distribution

(b) Scatter Plot: Inflation Intepolation Errors vs. z_{it}^x



Notes. The figure shows the distribution of inflation interpolation errors in for Chicago-Naperville-Elgin, Los Angeles-Long Beach-Anaheim, and New York-Newark-Jersey City (B.2a) and plots them against the shift-share variable z_{it}^x that we use to instrument u_{it} in Equation 27 (B.2b). The distribution of inflation interpolation errors pools errors obtained by imputing odd- and even-month observations in the three aforementioned MSAs. The coefficient of a linear regression with inflation interpolation errors as an outcome variable and the shift-share instrument as regressor is 0.002 (0.005).

We can construct interpolated inflation series by computing the 12-month percent difference in the interpolated CPI series. The interpolated inflation series for all MSAs are also characterized by measurement error with respect to the original ones. However, to the extent that such error terms are not systematically correlated with the exogenous instrument z_{it}^x, our estimates of ψ should not be biased. Figure B.2 shows the pooled distribution of inflation interpolation errors computed in odd and even

months for Chicago-Naperville-Elgin, Los Angeles-Long Beach-Anaheim, and New York-Newark-Jersey City (left panel) and plots them against the shift-share instrument z_{it}^{x} (right panel). As the right panel shows, there is no correlation between inflation interpolation errors and our instrument. Therefore, linearly interpolating our outcome variable allows us to exploit to the full extent the exogenous variation stemming from our instrumental variable.