

Hazell, Joe , Herreño, Juan, Nakamura, Emi & Steinsson, Jón (2022) The slope of the Phillips Curve: evidence from U.S. States. *Quarterly Journal of Economics, 137*(3), 1299-1344. https://doi.org/10.1093/qje/qjac010

https://researchonline.lse.ac.uk/id/eprint/113326/

Version: Accepted Version

This document is the author's accepted version of the journal article. There may be differences between this version and the published version.

Items deposited in LSE Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws.

LSE Research Online is the repository for research produced by the London School of Economics and Political Science. For more information, please refer to our Policies page or contact lseresearchonline@lse.ac.uk

The Slope of the Phillips Curve: Evidence from U.S. States

Jonathon Hazell	Juan Herreño	Emi Nakamura	Jón Steinsson*
LSE	UC San Diego	UC Berkeley	UC Berkeley

January 4, 2022

Abstract

We estimate the slope of the Phillips curve in the cross section of U.S. states using newly constructed state-level price indexes for non-tradeable goods back to 1978. Our estimates indicate that the slope of the Phillips curve is small and was small even during the early 1980s. We estimate only a modest decline in the slope of the Phillips curve since the 1980s. We use a multi-region model to infer the slope of the aggregate Phillips curve from our regional estimates. Applying our estimates to recent unemployment dynamics yields essentially no missing disinflation or missing reinflation over the past few business cycles. Our results imply that the sharp drop in core inflation in the early 1980s was mostly due to shifting expectations about long-run monetary policy as opposed to a steep Phillips curve, and the greater stability of inflation between 1990 and 2020 is mostly due to long-run inflation expectations becoming more firmly anchored.

JEL Classification: E30

*We would like to thank Tomas Breach, Massimiliano Cologgi, and Victoria de Quadros for excellent research assistance. We thank John Molino, our project coordinator at the BLS, for all his help throughout the project. We thank Hassan Afrouzi, Regis Barnichon, Olivier Blanchard, Jordi Galí, Yuriy Gorodnichenko, Pierre-Olivier Gourinchas, Greg Kaplan, Chen Lian, Silvana Tenreyro, John Williams and seminar participants at various institutions for valuable comments and discussions. We thank the National Science Foundation, the Alfred P. Sloan Foundation, the Smith Richardson Foundation, the Institute for New Economic Thinking, and the Julis-Rabinowitz Center for Public Policy and Finance for financial support.

1 Introduction

The Phillips curve is a formal statement of the common intuition that, if demand is high in a booming economy, this will provoke workers to seek higher wages, and firms to raise prices. A well-known formulation is the New Keynesian Phillips curve:

$$\pi_t = \beta E_t \pi_{t+1} - \kappa(u_t - u_t^n) + \nu_t. \tag{1}$$

According to this formulation, inflation π_t is determined by three factors: expected inflation $E_t\pi_{t+1}$, the output gap — measured here as the difference between unemployment u_t and the natural rate of unemployment u_t^n — and cost-push shocks ν_t. The slope of the Phillips curve κ represents the sensitivity of inflation to the output gap (i.e., to an increase in demand).

The episode in US economic history that has perhaps most strongly influenced the profession's thinking regarding the slope of the Phillips curve is the Volcker disinflation. In the early 1980s, Paul Volcker's Federal Reserve sharply tightened monetary policy. Unemployment rose sharply and inflation fell sharply. The conventional interpretation of this episode is that it provides evidence for a relatively steep Phillips curve.

One way to formalize this conventional interpretation is to assume that inflation expectations are adaptive: $\beta E_t \pi_{t+1} = \pi_{t-1}$ in equation (1). This yields the accelerationist Phillips curve:

$$\Delta\pi_t = -\kappa(u_t - u_t^n) + \nu_t. \tag{2}$$

Stock and Watson (2019) estimate κ in this equation and refer to it as the "Phillips correlation." They measure $\Delta\pi_t$ by the annual change in 12-month core PCE inflation, and $u_t - u_t^n$ by the CBO unemployment gap, both at a quarterly frequency. Figure 1 reproduces this analysis. It suggests that the slope of the Phillips curve was steep prior to and during the Volcker disinflation (0.67 for the period 1960-1983), but has flattened considerably since then (to only 0.03 for the period 2000-2019q1).[1]

The insensitivity of inflation to changes in unemployment between 1990 and 2020 led many economists to suggest that the Phillips curve had disappeared—or was "hibernating." During the Great Recession, unemployment rose to levels comparable to those during the Volcker disinflation, yet inflation fell by much less. The "missing disinflation" during and after the Great Recession

[1] See also Ball and Mazumder (2011), Kiley (2015b), and Blanchard (2016).



Figure 1: Stock and Watson's Changing Phillips Correlation

Note: Black solid line is a regression line for 2000-2019. Dark grey broken line is regression for 1984-1999. Light gray dash-dot line is a regression line for 1960-1983. The year-over-year change in inflation is the 4 quarter change in the (backward-looking) 4 quarter moving average of headline PCE inflation. The unemployment gap is the 4 quarter (backward-looking) moving average of the gap between the unemployment rate and the natural rate of unemployment. Authors' calculations. The figure replicates Figure 1 from Stock and Watson (2019).

then gave way to "missing reinflation" in the late 2010s as unemployment fell to levels not seen in 50 years, but inflation inched up only slightly. A similar debate raged in the late 1990s, when unemployment was also very low without this leading to much of a rise in inflation . Some have argued that the apparent flattening of the Phillips curve signals an important flaw in the Keynesian model.

There is, however, an alternative interpretation of these facts that emphasizes the anchoring of long-term inflation expectations in the United States (Bernanke, 2007; Mishkin, 2007). Figure 2 plots long-term inflation expectations from the Survey of Professional Forecasters. During the 1980s, long-term inflation expectations fluctuated a great deal. In particular, they fell rapidly over the period of the Volcker disinflation. In sharp contrast, since 1998, long-term inflation expectations have been extremely stable.

An alternative to the standard narrative of the Volcker disinflation is that the decline in inflation was driven not by a steep Phillips curve but by shifts in beliefs about the long-run monetary regime in the United States that caused the rapid fall in long-run inflation expectations we observe

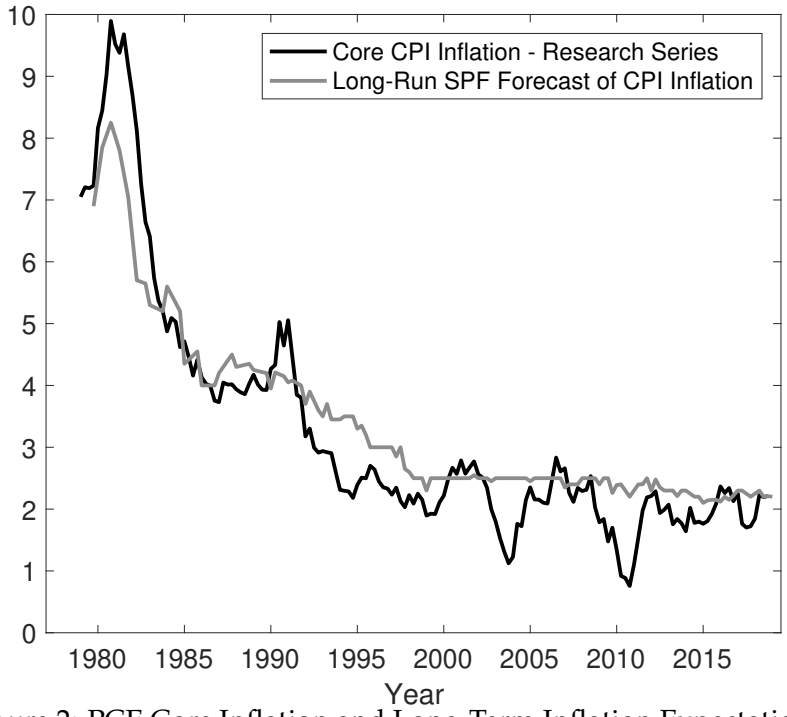


Figure 2: PCE Core Inflation and Long-Term Inflation Expectations

Note: The grey line plots 10 year ahead inflation expectation for the CPI. From 1990 onward, these come from Survey of Professional Forecasters. For the 1980s, these come from Blue Chip and are available on the Research and Data site of the Federal Reserve Bank of Philadelphia. The black line plots 12 month core CPI inflation using the Bureau of Labor Statistics' research series. This research series uses current methods to calculate inflation back in time.

in Figure 2. To see how this can be the case, it is useful to solve equation (1) forward and assume for simplicity that unemployment follows an AR(1) process. This yields

$$\pi_t = -\psi \tilde{u}_t + E_t \pi_{t+\infty} + \omega_t, \tag{3}$$

where $\tilde{u}_t$ denotes the deviation of unemployment from its long-run expected value, $E_t \pi_{t+\infty}$ represents long-term inflation expectations, and the parameter ψ is proportional to κ in equation (1). (Section 2 presents a more detailed derivation.) What this formulation of the Phillips curve makes clear is that changes in beliefs about the long-run monetary regime feed strongly into current inflation: the coefficient on $E_t \pi_{t+\infty}$ in equation (3) is one. Furthermore, in the presence of substantial variation in $E_t \pi_{t+\infty}$, the relationship between π_t and $\tilde{u}_t$ may be essentially uninformative about the slope of the Phillips curve (ψ and κ). In particular, if changes in $E_t \pi_{t+\infty}$ comove negatively with $\tilde{u}_t$ (as they would during an imperfectly credible shift in the long-run inflation target) the Phillips curve would appear to be steeper than it actually was.

Sargent (1982) emphasizes that hyperinflations tend to end quickly, much too quickly to be

explained by even a very large value of κ in the Phillips curve. In these episodes, it is clear that the primary cause of the abrupt fall in inflation is an abrupt fall in $E_t\pi_{t+\infty}$ associated with an abrupt change in the policy regime. Volcker's monetary policy constituted a sharp regime shift that was imperfectly credible at the outset but became gradually more credible as time passed (Erceg and Levin, 2003; Goodfriend and King, 2005; Bianchi and Ilut, 2017). This regime shift led to a large and sustained decline in long-term inflation expectations over the 1980s but also a transitory rise in unemployment. Perhaps it was this large change in inflation expectations that was the primary cause of the rapid fall in inflation over this period rather than high unemployment working through a steep Phillips curve.

This discussion highlights an important identification problem researchers face when they seek to estimate the slope of the Phillips curve: inflation expectations may covary with the output gap. Standard methods for estimating the Phillips curve aim to address this issue by controlling for inflation expectations $E_t\pi_{t+1}$ when estimating equation (1). A challenge with this approach is that estimates are quite sensitive to details of the specification. Mavroeidis, Plagborg-Møller, and Stock (2014) show that reasonable variation in the choice of data series, the specification, and the time period used yield a wide range of estimates for κ roughly centered on a value of zero (i.e., they are equally likely to have the "right" as the "wrong" sign). Mavroeidis, Plagborg-Møller, and Stock (2014) point to a weak instruments problem in driving these results: there simply isn't enough variation available in the aggregate data to separately identify the coefficients on unemployment and expected inflation. They conclude: "the literature has reached a limit on how much can be learned about the New Keynesian Phillips curve from aggregate macroeconomic time series. New identification approaches and new datasets are needed to reach an empirical consensus."

In addition to the identification problem discussed above, researchers seeking to estimate the slope of the Phillips curve also face the classic simultaneity problem of distinguishing demand shocks from supply shocks. Supply shocks (u_t^n and ν_t) yield positive comovement of inflation and unemployment (stagflation). If the variation used to identify the slope of the Phillips curve is contaminated by such shocks, the estimated slope will be biased towards zero and may even have the "wrong" sign. Fitzgerald and Nicolini (2014) and McLeay and Tenreyro (2019) point out that a central bank conducting optimal monetary policy will seek to offset aggregate demand shocks. If the central bank is successful, the remaining variation in inflation will be *only* due to supply shocks, a worst case scenario for the simultaneity problem.

Can cross-sectional data help overcome these problems? Several recent papers have argued

that they can. Fitzgerald and Nicolini (2014) and McLeay and Tenreyro (2019) show that using regional data helps overcome the simultaneity problem of distinguishing demand and supply shocks: central banks cannot offset regional demand shocks using a single national interest rate. These papers as well as Kiley (2015a), Babb and Detmeister (2017), Hooper, Mishkin, and Sufi (2019), and Fitzgerald et al. (2020) make use of city-level inflation data produced by the BLS to estimate regional Phillips curves. Beraja, Hurst, and Ospina (2019) use regional wage data to estimate wage Phillips curves.

We contribute to this regional Phillips curve literature in several ways. First, we show formally how estimating the Phillips curve using regional data provides a solution to the problem of shifting values of $E_t\pi_{t+\infty}$ confounding the estimation of the slope of the Phillips curve. We derive a regional Phillips curve in a "simple benchmark" multi-region model of a monetary union. The model clarifies the interpretation of the slope of regional Phillips curves relative to that of the aggregate Phillips curve. We also use the model to show that changes in the long-run monetary regime are absorbed by time fixed effects when the regional Phillips curve is estimated using a panel data specification. The intuition is that such long-run regime changes are common to all regions and therefore "cancel out" across regions within the monetary union.

Using our cross-section specification, we estimate a modest flattening of the Phillips curve when we split our sample in 1990: the Phillips curve in the post-1990 sample is flatter by a factor of two. This contrasts sharply with empirical specifications that make use of time series variation: a specification without time fixed effects yields a 50-100 times steeper Phillips curve for the pre-1990 sample. We interpret this as evidence that shifting long-run inflation expectations seriously confound estimates of the Phillips curve based on time series variation in the pre-1990 sample.

Our cross-sectional estimates indicate that the slope of the Phillips curve is small and was small even during the 1980s. Combining our estimate of the slope of the Phillips curve with an estimate of the persistence of fluctuations in unemployment, we find that a one percentage point increase in unemployment reduces inflation by about 0.34 percentage points, i.e., $\psi = 0.34$ in equation (3). This implies that only a modest fraction of the large changes in inflation in the early 1980s can be accounted for by the direct effect of increasing unemployment working through the slope of the Phillips curve. In contrast, movements in long-run inflation expectations were large over this period as is evident from Figure 2. In particular, long-run inflation expectations fell by about 4 percentage points from 1981 to 1986, accounting for about 2/3 of the fall in core inflation during this period. We conclude that a majority of the rapid decline in core inflation during the

Volcker disinflation arose from a rapid decline of long-term inflation expectations, associated with a rapidly changing monetary regime.[2]

Our estimates of the slope of the Phillips curve imply essentially no "missing disinflation" during the Great Recession or "missing reinflation" in the late 2010s or late 1990s. In other words, our cross-sectional estimates are consistent with the magnitude of movements in aggregate inflation post 1990. We conclude that the stability of inflation since 1990 is due to long-run inflation expectations becoming more firmly anchored. These conclusions echo those of Jorgensen and Lansing (2019).

Our analysis uses new state-level consumer price indexes for the United States that we have constructed back to the 1970s. Prior to our work, state level price indexes based on BLS micro price data have not existed. The BLS has published city-level inflation series for a group of relatively large cities. But it has refrained from reporting inflation indexes for smaller metropolitan areas (and for states). Our new state-level price indexes use all the available underlying micro-data gathered by the BLS. We also construct state-level price indexes for non-tradeables and tradeables. We focus our analysis on the behavior of the prices of non-tradeable goods. This is important. For prices set at the national level—as is more likely for tradeables—the slope of the regional Phillips curve will be zero no matter how large the slope of the aggregate Phillips curve is.

A notable conclusion of the recent regional Phillips curve literature has been that the estimated slope of the regional Phillips curve has tended to be steeper than the slope estimated for the aggregate Phillips curve. The theoretical framework we develop helps explain why this is the case. We show that panel data estimates of the regional Phillips curve by prior researchers are estimates of ψ in equation (3) as opposed to estimates of κ in equation (1). This means that they are not directly comparable to much of the aggregate literature. We discuss how researchers can convert estimates of ψ to κ and explain what other statistics this conversion depends on (primarily the degree of persistence of the unemployment variation used to estimate ψ). Our analysis highlights the importance of the exact specification used in estimating regional Phillips curves.[3]

The regional setting, along with our new inflation indexes, allow us to leverage new forms of variation in estimating the Phillips curve. We develop a new "tradeable demand spillovers" instrument building on insights from Nguyen (2014). This instrument is based on the idea that

[2]Carvalho et al. (2021) reach a similar conclusion using very different methods. They propose a model for long-run inflation expectations and show how their model generates the result that the Volcker disinflation was driven by shifting long-run inflation expectation and also that long-run inflation expectations become anchored in the 1990s onward.

[3]For example, Nishizaki and Watanabe (2000) find evidence of Phillips curve flattening in their baseline specification with no time fixed effects but this evidence changes dramatically when time fixed effects are added.

supply shocks in tradeable sectors will differentially affect demand in non-tradeable sectors in regions that are differentially exposed to the shocked tradeable sectors: e.g., an oil boom will increase demand for restaurant meals in Texas. In carrying out our regional analysis, we are careful to account for the fact that roughly 42% of the expenditure weight in core inflation is on the shelter component of housing services, which are measured by rents.[4] We estimate the slope of the regional Phillips curve for rents, and show that it is substantially steeper than the regional Phillips curve for non-tradeables excluding housing. We use the combination of these two estimates to predict the behavior of aggregate core inflation, which includes rents, and show that these predictions match the greater aggregate cyclicality of core inflation than core inflation excluding housing, a fact emphasized by Stock and Watson (2019). We conclude from this that the behavior of rent prices play an important role in determining the slope of both the regional and aggregate Phillips curves.

In addition to the papers discussed above, our work builds on the vast empirical and theoretical literature on the Phillips curve. The literature on the Phillips curve originates with Phillips (1958) and Samuelson and Solow (1960). Friedman (1968) and Phelps (1967) emphasized the importance of including an inflation expectations term in the Phillips curve. Gordon (1982) emphasized the importance of supply shocks. Important early papers that estimate the New Keynesian include Roberts (1995), Fuhrer and Moore (1995), Gali and Gertler (1999) and Sbordone (2002), but see also papers cited in Mavroeidis, Plagborg-Møller, and Stock (2014). Important recent papers estimating the Phillips curve include Ball and Mazumder (2011, 2019), Coibion and Gorodnichenko (2015b), Stock and Watson (2019), Barnichon and Mesters (2019), and Del Negro et al. (2020). Our paper is also related to a recent literature that assesses the missing disinflation during the Great Recession (see, e.g., Del Negro et al., 2015; Christiano et al., 2015; Gilchrist et al., 2017; Crump et al., 2019).

The paper proceeds as follows. Section 2 derives equation (3) and explains the problem of regime change in estimating the Phillips curve. Section 3 describes our main framework for interpreting the regional Phillips curve. Section 4 describes our new state-level inflation indexes. Section 5 presents our empirical results. Section 6 concludes.

[4]Much of the expenditure weight for housing derives from owner-occupied housing. However, rents are used to measure inflation for all shelter, due to the difficulty of backing out the user cost of housing from actual house prices in a theoretically appealing way. The expenditure weight of the CPI less food and energy is 77.7%, and 32.3 percentage points out of this expenditure weight are rents.

2 The Power and Problem of Long-Run Inflation Expectations

To appreciate the value of using regional variation to estimate the slope of the Phillips curve, it is useful to understand the central role of long-run inflation expectations in determining aggregate inflation. To this end, we solve equation (1) forward to get

$$\pi_t = -\kappa E_t \sum_{j=0}^{\infty} \beta^j u_{t+j} + \omega_t \tag{4}$$

where $\omega_t \equiv E_t \sum_{j=0}^{\infty} \beta^j (\kappa u^n_{t+j} + \nu_{t+j})$. This equation illustrates how inflation at time t is determined by the path of unemployment out into the infinite future.[5] We can furthermore decompose the variation in future unemployment u_{t+j} into a transitory and permanent component. Define the transitory component of variation in unemployment to be $\tilde{u}_t = u_t - E_t u_{t+\infty}$, where $E_t u_{t+\infty}$ is the permanent component of the variation in unemployment. Using these concepts, we can rewrite equation (4) as

$$\pi_t = -\kappa E_t \sum_{j=0}^{\infty} \beta^j \tilde{u}_{t+j} - \frac{\kappa}{1-\beta} E_t u_{t+\infty} + \omega_t, \tag{5}$$

Assuming that shocks to u^n_t and ν_t are transitory, equation (1) implies that $E_t \pi_{t+\infty} = -\frac{\kappa}{1-\beta} E_t u_{t+\infty}$. We can then rewrite equation (5) as

$$\pi_t = -\kappa E_t \sum_{j=0}^{\infty} \beta^j \tilde{u}_{t+j} + E_t \pi_{t+\infty} + \omega_t. \tag{6}$$

Finally, let's assume for simplicity that $\tilde{u}_t$ follows and AR(1) process with autocorrelation coefficient equal to ρ_u. In this case $E_t \tilde{u}_{t+j} = \rho^j_u \tilde{u}_t$ and we can rewrite equation (6) as

$$\pi_t = -\psi \tilde{u}_t + E_t \pi_{t+\infty} + \omega_t, \tag{7}$$

where $\psi = \kappa/(1 - \beta\rho_u)$.

[5] While the most popular micro-foundation of the New Keynesian Phillips curve—and the one we develop in section 3—is based on the price rigidity assumptions in Calvo (1983), this equation or something very similar arises from several other micro-foundations. Roberts (1995) shows that the same Phillips curve arises from Rotemberg's (1982) quadratic costs of price adjustment model and Taylor's (1979,1980) model of staggered contracts (the timing of the output gap term is slightly different in the Taylor model). Furthermore, Gertler and Leahy (2008) develop the same Phillips curve as a linear approximation of a model with Ss foundations. In the case of the Rotemberg model in continuous time, the derivation does not rely on a linear approximation around a zero inflation steady state. Models based on information frictions yield Phillips curves that are not forward looking. These models, however, typically assume no price rigidity. Incorporating price rigidity into these models would make their Phillips curves forward looking as well. Sbordone (2002), Gali, Gertler, and Lopez-Salido (2005), and Rudd and Whelan (2005) develop approaches to estimating the Phillips Curve on aggregate data using versions of equation (4).

This way of writing the Phillips curve highlights the importance of long-run inflation expectations in determining inflation at the aggregate level. Long-run inflation expectations $E_t\pi_{t+\infty}$ appear with a coefficient of one in equation (7). In other words, current inflation moves one-for-one with changes in long-run inflation expectations. These long-run expectations are determined by the private sector's beliefs about the long-run monetary regime being followed by the central bank (the long-run inflation target). Variation in beliefs about the long-run monetary regime therefore have very large effects on current inflation.[6]

Equation (7) implies that inflation can vary dramatically without *any* variation in $\tilde{u}_t$ if there is substantial variation in long-run inflation expectations. In this case, the relationship between inflation and $\tilde{u}_t$ may be entirely uninformative about the slope of the Phillips curve. Worse still, variation in long-run inflation expectations may be correlated with variation in $\tilde{u}_t$. For example, it seems very plausible that Paul Volcker's willingness to allow unemployment to rise to very high values in the early 1980s—and the fact that Volcker was not forced to resign—signalled to the public that he was serious about bringing down inflation (and had the backing of the president to do this). Such a correlation will impart an upward bias on estimates of the slope of the Phillips curve unless variation in inflation expectations can be controlled for. But in practice, controlling for inflation expectations is hard due to weak instruments (Mavroeidis et al., 2014) and because direct measures of inflation expectations may be imperfect. So, a rapid drop in inflation expectations may masquerade as a steep Phillips curve.

Why has the Phillips curve appeared to flatten over the past few decades? Figure 2 shows that since roughly 1998, long-term inflation expectations have been firmly anchored at close to 2%. This has led to a collapse of the covariance between $E_t\pi_{t+\infty}$ and unemployment and therefore eliminated any bias associated with poorly proxied variation in inflation expectations. A fall in this bias will appear from the perspective of the (misspecified) accelerationist Phillips curve (such as the one we discuss in the introduction) as a flatter curve.

One piece of corroborating evidence for this view is the close relationship between π_t and $E_t\pi_{t+1}$ in the data. Recall that the standard formulation of the New Keynesian Phillips—equation (1)—implies that it is the gap between π_t and $\beta E_t\pi_{t+1}$—let's call this the "inflation gap"—that must be explained by demand pressure (the κu_t term) or supply shocks ($\kappa u_t^n + \nu_t$). Figure 3 plots

[6]Equations (6) and (7) remain valid in the case where the coefficient on $E_t\pi_{t+1}$ in equation (1) is equal to one rather than β. In this case, the long-run Phillips curve is vertical, and long-run unemployment and inflation are unrelated. That is, inflation still satisfies $\pi_t = -\kappa E_t \sum_{j=0}^{\infty} \tilde{u}_{t+j} + \omega_t + E_t\pi_{t+\infty}$, but $E_t u_{t+\infty}$ is now independent of $E_t\pi_{t+\infty}$. The forward sum is bounded even though $\beta = 1$, because $\tilde{u}_t$ has zero unconditional mean.

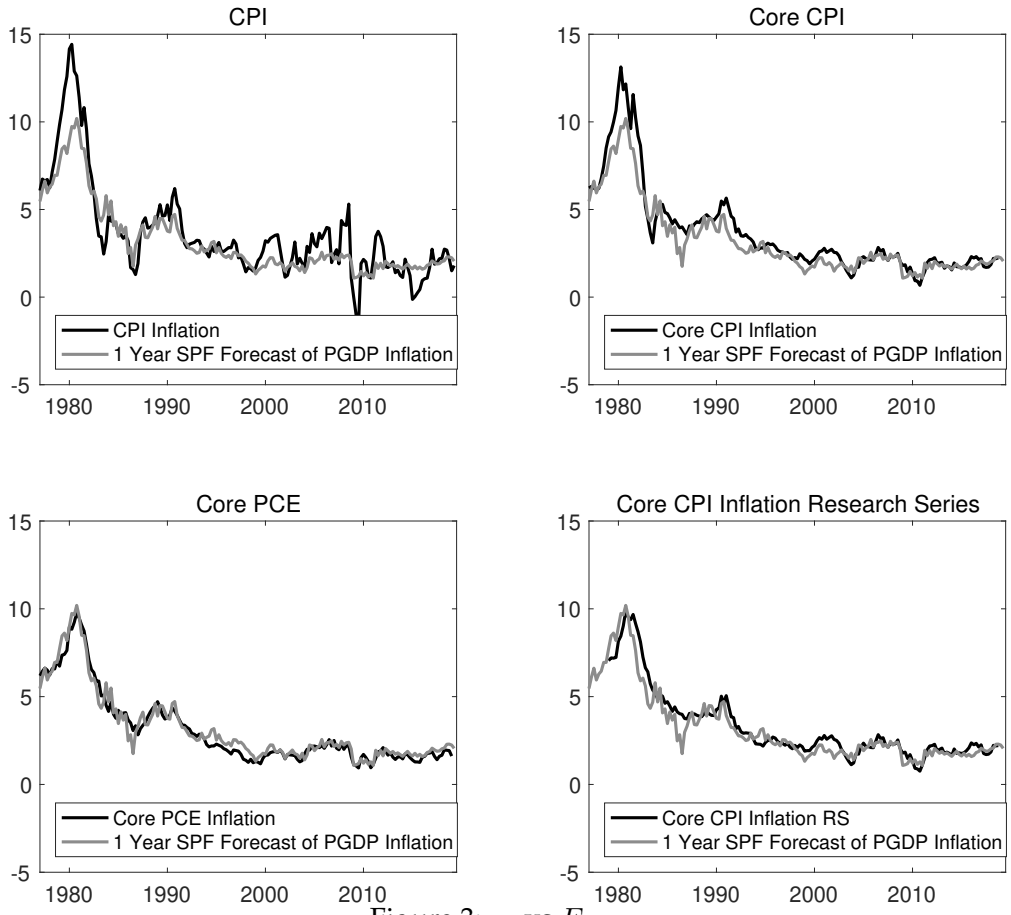


Figure 3: π_t vs $E_t\pi_{t+1}$

Note: Each panel shows the comparison of the one-year ahead forecast of the GDP deflator coming from the Survey of Professional of Forecasters and a measure of inflation. The top left panel uses the published headline CPI. The top right panel excludes food and energy by plotting the published measure of the Core CPI. The bottom panels correct for changes in the methodology of inflation measurement. The bottom left panel uses PCE inflation which has maintained a stable methodology, while the bottom right panel uses the Constant Methodology Research Series for Core CPI published by the Bureau of Labor Statistics. We use forecasts of the GDP deflator because forecasts for the CPI are not available before 1980.

SPF forecasts of inflation over the next year along with four different measures of current inflation. The difference between the two series is approximately equal to the inflation gap $\pi_t - \beta E_t\pi_{t+1}$.

The measure of current inflation plotted in the top-left panel of Figure 3 is the 12-month change in the overall CPI. This conventional way of comparing current inflation and inflation expectations over the next year suggests that these series are closely related, but that there is nevertheless substantial variation in the gap between them (the inflation gap). Moving to the top-right panel, we measure current inflation by the 12-month change in core CPI inflation, excluding food and energy. The inflation gap measured this way is quite a bit smaller. Evidently, commodities account for a large part of the inflation gap for the overall CPI. However, a substantial inflation gap remains in the early 1980s.

The measure of current inflation plotted in the bottom-left panel of Figure 3 is the 12-month change in the core PCE. The advantage of this series is that it makes use of current measurement methods, retroactively applied back in time. In this case, the inflation gap is very small. A similar message emerges in the bottom-right panel using the 12-month change in the core CPI research series published by the BLS. This series also uses consistent, modern methods to calculate inflation back in time. A particularly important measurement change for our purposes occurred in 1983, when the BLS switched to using rent inflation as a proxy for overall housing inflation, including for owner-occupied housing ("rental equivalence"). Before that time, housing services inflation in the CPI was constructed from a weighted average of changes in house prices and mortgage costs (i.e., interest rates). This earlier approach essentially "baked in" a strong relationship between Volcker's actions to curb the Great Inflation and measured CPI inflation, since interest rates (and house prices) fed directly into the CPI.[7]

The overall message that emerges from Figure 3 is that the inflation gap for core inflation measured using modern methods is tiny throughout our sample period. Importantly, this includes the period of the Volcker disinflation. This is suggestive evidence that the slope of the Phillips curve was small throughout our sample period: unemployment varied a great deal both in the early 1980s and again in the Great Recession without much variation in the inflation gap. However, the four panels in Figure 3 illustrate well that this conclusion is sensitive to the details of how inflation is measured.[8] It is also sensitive to whether the expectations data used come from the SPF or from the Michigan Survey of Consumers as Coibion and Gorodnichenko (2015b) emphasize, and also sensitive to the exact timing of the variables.

3 A Model of the Regional Phillips Curve

We now develop a two-region, New Keynesian, open economy model featuring tradeable and non-tradeable sectors. We derive a regional Phillips curve in this model and show how it relates to the aggregate Phillips curve. The model demonstrates a chief benefit of regional data: time and state fixed effects "difference out" changes in long run inflation expectations. The model also illustrates the importance of using non-tradeable inflation when estimating the slope of the

[7] These choices are consequential since the housing component of the CPI has a weight of roughly one-third in the overall CPI. Appendix B.2 presents our attempt to replicate the pre-1983 BLS housing methodology on more modern data. The main conclusion from this is that this methodology would have led to much more variable (and cyclical) inflation over the past few decades.

[8] We discuss this in more detail in appendix B.1.

Phillips curve using regional data.

3.1 Model Setup

Our model consists of two regions that belong to a monetary and fiscal union. We refer to the regions as Home (H) and Foreign (F). The population of the entire economy is normalized to one. The population of the home region is denoted by ζ. Labor is immobile across regions. Within each region, there is a single labor market. Household preferences, market structure, and firm behavior take the same form in both regions. Below, we describe the economy of the home region. All prices in the economy are denominated in "dollars," a digital currency issued by the federal government.[9] Throughout, we adopt the following conventions unless otherwise stated. Lower case variables are the logs of upper case variables. Hatted variables denote the percentage deviation of a variable from its steady state value. Steady state values are recorded without time subscripts.

3.1.1 Households

The representative household in the home region seeks to maximize the utility function

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

where C_{Ht} is per capita consumption of a composite consumption good, N_{Ht} is per capita employment, and β is the household's subjective discount factor. We follow Greenwood, Hercowitz, and Huffman (1988) in assuming that the function $u(C_{Ht}, N_{Ht})$ takes the form

$$u(C_{Ht}, N_{Ht}) = \frac{\left(C_{Ht} - \chi \frac{N_{Ht}^{1+\varphi^{-1}}}{1+\varphi^{-1}}\right)^{1-\sigma^{-1}}}{1-\sigma^{-1}} \tag{8}$$

where φ is the household's Frisch elasticity of labor supply, σ determines the household's elasticity of intertemporal substitution, and χ governs the intensity of the household's disutility of labor. We refer to this preference specification as GHH preferences.

The composite consumption good C_{Ht} is a constant elasticity of substitution (CES) index over

[9] In other words, we are considering an economy in the cashless limit (Woodford, 1998, 2003).

tradeables C_{Ht}^T and non-tradeables C_{Ht}^N given by

$$C_{Ht} = \left[\phi_N^{\frac{1}{\eta}} {C_{Ht}^N}^{\frac{\eta-1}{\eta}} + \phi_T^{\frac{1}{\eta}} {C_{Ht}^T}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where η is the elasticity of substitution between tradeables and non-tradeables and ϕ_T and ϕ_N are the household's steady state expenditure shares on tradeable and non-tradeable goods, respectively. C_{Ht}^N and C_{Ht}^T are themselves composite goods described further below. Non-tradeable goods are only consumed in the region in which they are produced. In contrast, the market for tradeable goods is completely integrated across regions. Hence, home and foreign households may face different prices for non-tradeables, but face the same prices for tradeable goods. The expenditure share on tradeable and non-tradeable goods must sum to one, i.e., $\phi_N + \phi_T = 1$.

The composite non-tradeable good C_{Ht}^N is given by

$$C_{Ht}^N = \left[\int_0^1 C_{Ht}^N(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}}$$

where $C_{Ht}^N(z)$ denotes consumption of variety z of non-tradeable goods in the home region. The home price of this non-tradeable variety is $P_{Ht}^N(z)$. The parameter $\theta > 1$ denotes the elasticity of substitution between different non-tradeable varieties.

Home tradeable consumption C_{Ht}^T is a CES aggregate over tradeable goods produced in the home and foreign regions given by

$$C_{Ht}^T = \left[{\tau_{Ht}^H}^{\frac{1}{\eta}} {C_{Ht}^{TH}}^{\frac{\eta-1}{\eta}} + {\tau_{Ht}^F}^{\frac{1}{\eta}} {C_{Ht}^{TF}}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \tag{9}$$

where C_{Ht}^{TH} and C_{Ht}^{TF} are home consumption of composite tradeable goods produced in the home and foreign regions, respectively. We assume (for simplicity) that the elasticity of substitution between home-produced and foreign-produces tradables is η (the same as the elasticity of substitution between tradeables and non-tradeables). Demand for home-produced and foreign-produced tradeables is subject to shocks denoted by τ_{Ht}^H and τ_{Ht}^F, respectively. We normalize $\tau_{Ht}^H + \tau_{Ht}^F = 1$. For simplicity, we do not allow for home bias in tradeable consumption. Thus, we set $\tau_H^H = \tau_F^H = \zeta$, i.e., the share of spending on goods from the home region in each region is equal to the size of the home region.

The home and foreign composite tradeable goods are CES indexes given by

$$C_{Ht}^{TH} = \left[\int_0^1 C_{Ht}^{TH}(z)^{\frac{\theta-1}{\theta}} dz\right]^{\frac{\theta}{\theta-1}} \quad \text{and} \quad C_{Ht}^{TF} = \left[\int_0^1 C_{Ht}^{TF}(z)^{\frac{\theta-1}{\theta}} dz\right]^{\frac{\theta}{\theta-1}}.$$

where $C_{Ht}^{TH}(z)$ and $C_{Ht}^{TF}(z)$ are home consumption of varieties of tradeable goods produced in the home and foreign region, respectively. The prices of these home-produced and foreign-produced tradeable good varieties are $P_{Ht}^{T}(z)$ and $P_{Ft}^{T}(z)$, respectively.

Households maximize utility subject to a sequence of budget constraints

$$C_{Ht}^{N} P_{Ht}^{N} + C_{Ht}^{TH} P_{Ht}^{T} + C_{Ht}^{TF} P_{Ft}^{T} + E_t\left[M_{Ht,t+1} B_{H,t+1}\right] \leq B_{Ht} + W_{Ht} N_{Ht} + \Xi_{Ht}^{N} + \Xi_{Ht}^{T}$$

where B_{Ht} is a random variable denoting payoffs of the state contingent portfolio held by households in period t; $M_{Ht,t+1}$ is the one-period-ahead stochastic discount factor of the home representative household; P_{Ht}^{N}, P_{Ht}^{T}, and P_{Ft}^{T} are price indexes that give the minimum cost of purchasing a unit of C_{Ht}^{N}, C_{Ht}^{TH}, and C_{Ht}^{TF}, respectively; W_{Ht} is the nominal wage received by workers in region H; and Ξ_{Ht} and Ξ_{Ht} are the profits of non-tradeable and tradeable firms in the home region. There is a complete set of financial markets across the two regions. To rule out Ponzi schemes, we assume that household debt cannot exceed the present value of future income in any state.

We present the first order necessary conditions for household optimization in Appendix A.1. As we noted above, the problem of the foreign household is analogous. We therefore refrain from describing it in detail here. For simplicity, we do not allow for tradeable demand shocks to foreign tradeable consumption as we do for home tradeable consumption.

3.1.2 Firms

There is a continuum of firms in each of the tradeable and non-tradeable sectors. Firms are indexed by z and firm z specializes in the production of differentiated good z. Labor is the only variable factor of production used by firms.

We begin by discussing the non-tradeable sector. The output of good z in the non-tradeable sector is denoted $Y_{Ht}^{N}(z)$. The production function of firm z in this sector is

$$Y_{Ht}^{N}(z) = Z_{Ht}^{N} N_{Ht}^{N}(z), \tag{10}$$

where $N_{Ht}^{N}(z)$ is the amount of labor demanded by firm z and Z_{Ht}^{N} is a productivity shock.

Firm z in the non-tradable sector maximizes its value:

$$E_t \sum_{j=0}^{\infty} M_{Ht,t+j} \left[P^N_{H,t+j}(z) Y^N_{H,t+j}(z) - W_{H,t+j} N^N_{H,t+j}(z) \right]$$

given demand for its good, which is

$$Y^N_{Ht}(z) = \zeta C^N_{Ht} \left(\frac{P^N_{Ht}(z)}{P^N_{Ht}} \right)^{-\theta}.$$

Firm z can set its price freely with probability $1-\alpha$ as in Calvo (1983). With probability α the firm must keep its price unchanged.

Analogously to the non-tradeable sector, the output of firm z in the tradeable sector is denoted $Y^T_{Ht}(z)$. Its production function is

$$Y^T_{Ht}(z) = Z^T_{Ht} N^T_{Ht}(z)$$

where $N^T_{Ht}(z)$ is the amount of labor demanded by the firm producing good z and Z^T_{Ht} is a productivity shock.

Firm z in the tradeable sector maximizes its value:

$$E_t \sum_{j=0}^{\infty} M_{Ht,t+j} \left[P^T_{H,t+j}(z) Y^T_{H,t+j}(z) - W_{H,t+j} N^T_{H,t+j}(z) \right]$$

given demand for its good. Demand in the tradeable sector comes from both the home and foreign regions. Firm z's demand is thus given by

$$Y^T_{Ht}(z) = \left(\zeta C^{TH}_{Ht} + (1-\zeta) C^{TH}_{Ft} \right) \left(\frac{P^T_{Ht}(z)}{P^T_{Ht}} \right)^{-\theta}.$$

The tradeable goods firms also have an opportunity to change their price with probability $1-\alpha$ each period and must otherwise keep their prices fixed.

We present the first order necessary conditions for firm optimization in Appendix A.2. The problems of foreign firms are analogous to those of home firms.

3.1.3 Government Policy and Equilibrium

The federal government operates a common monetary policy for the two regions. This policy takes the form of the following interest rate rule

$$\hat{r}_t^n = \varphi_\pi \left(\pi_t - \bar{\pi}_t\right) - \varphi_u \left(\hat{u}_t - \bar{u}_t\right) + \varepsilon_{rt},$$

where, as elsewhere in the paper, hatted variables denote deviations from a zero inflation steady state and lower case variables are the logs of upper case variables. Economy-wide inflation π_t is a population weighted average of inflation in the two regions: $\pi_t \equiv \zeta \pi_{Ht} + (1-\zeta)\pi_{Ft}$, where $\pi_{Ht} = p_{Ht} - p_{H,t-1}$ is consumer price inflation in the home region and π_{Ft} is defined analogously for the foreign region. In our model, we define unemployment in the home region simply as $u_{Ht} = 1 - N_{Ht}$. We define foreign unemployment analogously. This implies that to a first order $\hat{u}_{Ht} = -\hat{n}_{Ht}$ and $\hat{u}_{Ft} = -\hat{n}_{Ft}$. Economy-wide unemployment is a population weighted average of unemployment in the two regions, so $\hat{u}_t = \zeta \hat{u}_{Ht} + (1-\zeta)\hat{u}_{Ft}$.

Importantly, we allow the monetary authority to have a time-varying inflation target $\bar{\pi}_t$.[10] Since the long-run Phillips curve in our model is not vertical, variation in long-run inflation yields variation in long-run unemployment. We assume that the monetary authority targets an unemployment rate that is consistent with its long-run inflation target, i.e., $\bar{u}_t = (1-\beta)\bar{\pi}_t/\kappa$. We assume that φ_π and φ_u obey the Taylor principle, ensuring that the economy has a unique locally bounded equilibrium. ε_{rt} is a transitory monetary shock, which we assume follows an exogenous AR(1) process.

For simplicity, the government levies no taxes, engages in no spending, and issues no debt. In other words, there is no fiscal policy. The digital currency issued by the government is in zero net supply. The government's monetary policy, therefore, has no fiscal implications. An equilibrium in this economy is an allocation that satisfies household optimization, firm optimization, the government's interest rate rule, and market clearing. We focus on the unique locally bounded equilibrium of the model. Implicitly we rule out equilibria in which the inflation rate rises without bound using the trigger strategy argument presented in Obstfeld and Rogoff (1983).

[10] Prior work that allows for a time-varying inflation target includes Stock and Watson (2007), Ireland (2007) and Cogley and Sbordone (2008).

3.2 Regional and Aggregate Phillips Curves

Taking a log-linear approximation of the model presented in section 3.1 around a zero-inflation steady state with balanced trade yields the following regional Phillips Curve for the inflation of non-tradeable goods:

$$\pi_{Ht}^{N} = \beta E_t \pi_{H,t+1}^{N} - \kappa \hat{u}_{Ht} - \lambda \hat{p}_{Ht}^{N} + \nu_{Ht}^{N}, \tag{11}$$

and aggregate Phillips Curve for overall inflation:

$$\pi_t = \beta E_t \pi_{t+1} - \kappa \hat{u}_t + \nu_t, \tag{12}$$

where $\pi_{Ht}^{N} = p_{Ht}^{N} - p_{H,t-1}^{N}$ is home non-tradeable inflation, $\hat{p}_{Ht}^{N} = P_{Ht}^{N}/P_{Ht} - 1$ is the percentage deviation of the home relative price of non-tradables from its steady state value of one, ν_{Ht}^{N} is a non-tradeable home supply shock, ν_t is a corresponding aggregate supply shock, and the parameter $\kappa = \lambda\varphi^{-1}$, where $\lambda = (1-\alpha)(1-\alpha\beta)/\alpha$. We provide a detailed derivation of these equations in Appendix A.

Equations (11) and (12) yield an important result: The slopes of the regional Phillips curve for non-tradeables and the aggregate Phillips curve are the same in our model. These slopes are both equal to κ. This result holds for the *non-tradeable* regional Phillips curve, but does not carry over to the regional Phillips curve for overall consumer price inflation—which includes both tradeable and non-tradeable inflation in the region. As we show in Appendix A.8, the slope of the regional Phillips curve for overall consumer price inflation is smaller by a factor equal to the expenditure share on non-tradeable goods.

Intuitively, the difference in the slope between the non-tradeable and overall regional Phillips curves arises because all regions share the tradeable goods and these goods are priced nationally. The tradeable goods therefore don't contribute to difference in inflation across regions, which means that the regional CPI is made up partly of goods whose regional prices are insensitive to regional variation in unemployment. This makes the regional CPI less sensitive to regional unemployment than the aggregate CPI is to aggregate unemployment.

Our result that the slope of the non-tradeable regional Phillips curve is equal to the slope of the aggregate Phillips curve leads us to focus our cross-sectional empirical work on inflation for non-tradeable goods. Earlier research that has estimated regional Phillips curves has done so for overall consumer price inflation at the regional level (e.g., Fitzgerald and Nicolini, 2014;

McLeay and Tenreyro, 2019). Our model suggests that results from such analysis are less directly informative about the slope of the aggregate Phillips curve.

Our assumption that households have GHH preferences helps simplify the derivation of the regional and aggregate Phillips curves in our model—equations (11) and (12). GHH preferences imply that wealth effects on labor supply are zero, which eliminates the dependence of marginal costs on consumption. The absence of a consumption term in the Phillips curve plays a role in the derivation of our result that the non-tradeable regional Phillips curve and the aggregate Phillips curve have the same slope. We discuss this point at greater length in Appendix A.9. The form of the Phillips curve in our model does not, however, depend on the structure of financial markets. We have assumed complete financial markets across regions, but the Phillips curve is the same in a model with incomplete markets across regions.

An important difference between equations (11) and (12) is the presence of the relative price of non-tradeables term $\lambda \hat{p}^N_{Ht}$ in equation (11). This term implies that inflation in the non-tradeables sector will be lower the higher is the relative price of non-tradeables. Conceptually, this term is very important. It pushes relative prices towards parity in the long run. Also, it implies that even if prices in the economy are very flexible—κ is very large—a local boom will not result in unbounded inflation of home non-tradeable prices since demand for these goods is affected by their prices relative to other prices in the economy. The mechanical reason this term appears is that the inflation rate for non-tradeable goods is driven by variation in the real wage deflated by non-tradeable prices. Labor supply in the home region, however, is a function of the real wage deflated by the home consumer price index. The real marginal cost variable in the home non-tradeable Phillips curve therefore gives rise to an unemployment term and a relative price of non-tradeables term.

3.3 Estimating the Slope of the Phillips Curve with Regional Data

Next we solve the regional Phillips curve — equation (11) — forward to obtain

$$\pi^N_{Ht} = -E_t \sum_{j=0}^{\infty} \beta^j \left(\kappa \tilde{u}_{H,t+j} + \lambda \hat{p}^N_{H,t+j} \right) + E_t \pi^N_{t+\infty} + \omega^N_{Ht}, \tag{13}$$

where $\tilde{u}_{Ht} = u_{Ht} - E_t u_{H,t+\infty}$ and $\omega^N_{Ht} = E_t \sum_{j=0}^{\infty} \beta^j \nu^N_{H,t+j}$.

A major benefit of estimating the slope of the Phillips curve using regional data from a monetary union is that variation in long-run inflation expectations — the $E_t \pi^N_{t+\infty}$ term in equation (13)

— is constant across regions. This implies that variation in long-run inflation expectations will be absorbed by time fixed effects in a panel specification. Intuitively, while short-run inflation expectations ($E_t\pi^N_{t+1}$) will differ across regions due to differences in their economic circumstances, long-run inflation expectations ($E_t\pi^N_{t+\infty}$) are independent of the current business cycle. They are determined by beliefs about the long-run monetary regimes. In a monetary union like the US, these beliefs will vary uniformly across regions. This means that these expectations are "differenced out" in a panel regression with time fixed effects.

The result that long-run inflation expectations are constant across regions (and sectors) in our model relies on productivity and other drivers of real costs having a common trend in the long run. If productivity growth (say) differs across regions even in the long run, this will lead to persistent differences in non-tradeable inflation (a Balassa-Samuelson effect). However, if this difference is constant over time, it will be absorbed by region fixed effects in a panel specification.

These observations imply that that we can adopt an empirical specification that replaces the $E_t\pi^N_{t+\infty}$ term in equation (13) with time and region fixed effects:

$$\pi^N_{it} = -E_t \sum_{j=0}^{\infty} \beta^j \left(\kappa u_{i,t+j} + \lambda \hat{p}^N_{i,t+j}\right) + \alpha_i + \gamma_t + \tilde{\omega}^N_{it}, \tag{14}$$

where i denotes region, α_i denotes a set of region fixed effects, and γ_t denotes a set of time fixed effects. Variation in $E_t\pi^N_{t+\infty}$ in equation (13) that is common across regions will be absorbed by the time fixed effect.[11] Constant differences across regions in $E_t\pi^N_{t+\infty}$ will be absorbed by the state fixed effects. To the extent that there is remaining variation in $E_t\pi^N_{t+\infty}$ across regions (e.g., due to changing trends), it will be a part of the error term $\tilde{\omega}^N_{it}$.

It is useful to relate equation (14) to the empirical specifications used in the recent regional Phillips curve literature. If we assume that both u_{Ht} and $\hat{p}^N_{Ht}$ follow AR(1) processes with autocorrelation coefficients equal to ρ_u and ρ_{pN}, respectively, equation (14) simplifies to

$$\pi^N_{it} = -\psi u_{it} - \delta \hat{p}^N_{it} + \alpha_i + \gamma_t + \tilde{\omega}^N_{it}, \tag{15}$$

where $\psi = \kappa/(1-\beta\rho_u)$ and $\delta = \lambda/(1-\beta\rho_{pN})$. This equation is similar to the empirical specification used by much of the recent regional Phillips curve literature. Comparing equations (14) and (15),

[11] The time fixed effects also absorb time variation in the long-run expected unemployment $E_t u_{t+\infty}$. We have therefore replaced $\tilde{u}_{i,t+j}$ in equation (13) with $u_{i,t+j}$ in equation (14). This equation remains valid if $\beta = 1$. The forward sum in the equation is still bounded, because u_{it} has zero mean conditional on time fixed effects.

we see that an important difference between these two specifications is that the slope coefficient is not the same. The slope coefficient in equation (14) is κ (which is the same as the slope coefficient in equation (11) and (12)), while the slope coefficient in equation (15) is $\psi = \kappa/(1 - \beta\rho_u)$. Since unemployment is quite persistent, ψ is likely to be substantially larger than κ. Note that the AR(1) assumption we use to derive equation (15) is not used in our estimation of κ.

A curious feature of the recent regional Phillips curve literature is that it has tended to yield larger estimates of the slope of the Phillips curve than more traditional estimation strategies based on aggregate data (Fitzgerald and Nicolini, 2014; Babb and Detmeister, 2017; McLeay and Tenreyro, 2019; Hooper et al., 2019). Comparing equations (14) and (15) provides a simple explanation for this discrepancy. The regional Phillips curve literature has been estimating ψ in equation (15), while the more traditional literature using aggregate variation has typically been estimating κ. Since $\psi >> \kappa$, it is not surprising that the slope of the Phillips curve estimated in the regional literature has seemed large relative to traditional estimates.[12]

The difference between κ and ψ arises due to the different ways equations (11) and (15) capture the effects of expected future unemployment on current inflation. In equation (11), the effects of expected future unemployment on current inflation are captured by the inflation expectations term $E_t\pi_{t+1}$ and the coefficient on current unemployment κ only reflects the effect of current unemployment on current inflation. In contrast, the slope coefficient in equation (15) captures both the effect of current unemployment and the effect of expected future unemployment into the indefinite future on current inflation—i.e., the fact that high unemployment today forecasts high unemployment in future periods.[13]

An advantage of estimating specifications such as equations (14) and (15) rather than equation (11) is that the identification of the slope coefficient is less sensitive to the exact timing of changes in inflation relative to inflation expectations. In Figure 3, we show that the difference between inflation and inflation expectations is quite sensitive to the exact measure of inflation.

[12]This same type of lack of comparability arises in some cases for different estimates based on aggregate data. Some researchers use longer-term inflation expectations, rather than one-period ahead inflation expectations, to proxy for $E_t\pi^N_{t+1}$ when estimating the Phillips curve using aggregate data. Our analysis shows, however, that when researchers choose to use data on long-term inflation expectations, they (perhaps inadvertently) end up estimating ψ, not κ. To compare such estimates with those based on a specification that controls for one-period ahead expectations, one must translate between the two, e.g., by using the formula $\psi = \kappa/(1 - \beta\rho_u)$ or a version of this formula appropriate for (say) 10-year ahead inflation expectations.

[13]McLeay and Tenreyro (2019) control for inflation expectations at the Census Region level when they estimate the regional Phillips curve. The variation across regions in these inflation expectations data is quite minimal. It may therefore be that the variation in this variable is quite attenuated relative to actual variation in inflation expectations across the MSA areas that form the regional units in their analysis.

We have so far manipulated the Phillips curve under the standard assumption of full-information rational expectations. However, the arguments we make above — i.e., solving the Phillips curve forward — rely only on the weaker assumption, that the law of iterated expectations holds. We elaborate on this point in Appendix A.10, drawing on results from Adam and Padula (2011) and Coibion, Gorodnichenko, and Kamdar (2018).

To derive a tractable empirical specification for the regional Phillips curve in which the coefficient on unemployment is the same as in the aggregate Phillips curve, we have made a number of strong assumptions (perfect labor mobility within region, no labor mobility across regions, GHH preferences, production linear in labor, etc.). In the world, these assumptions are unlikely to hold exactly and the empirical specification we estimate is thus unlikely to yield exactly the slope of the aggregate Phillips curve. Deriving an exact analytical mapping is nontheless useful since it highlights in a transparent way the importance of certain forces (e.g., inflation expectations). In more general models in which no exact analytical mapping between the slope of the regional and aggregate Phillips curves exists, our regional slope estimates can be used as empirical targets in a moment matching exercise. Even away from the simple case we analyze where the aggregate and regional slopes are equal, these moments are likely to provide valuable information about the slope of the aggregate Phillips curve (Nakamura and Steinsson, 2018; Andrews, Gentzkow, and Shapiro, 2020).

4 Data and Construction of State-Level Price Indexes

The BLS does not publish state-level price indexes. Prior work has used metropolitan level BLS price indexes and cost of living estimates from the American Chamber of Commerce Realtors Association (ACCRA) to construct state-level price indexes (see, e.g., Del Negro, 1998; Nakamura and Steinsson, 2014). An important drawback of this approach is that the BLS imputes missing data using data from other regions. Recent work has used scanner price data to construct state-level price indexes (Beraja, Hurst, and Ospina, 2019). An important drawback of scanner data is the short sample period available.

We construct new state-level price indexes for the US based on the micro-price data the BLS collects for the purpose of constructing the CPI. Our sample period is 1978 to 2018 (with a 26 month gap in 1986-1988 due to missing micro-data). The micro-data that we base our price indexes on are available in the CPI Research Database at the BLS. The data for the period 1978-1987

were constructed by Nakamura et al. (2018). The micro-price data in the CPI Research Database cover thousands of individual goods and services, constituting about 70% of consumer expenditures. They are collected by BLS employees who visit outlets to record prices. The database does not include the rent prices used to construct the shelter component of the CPI. For this reason, we analyze the behavior of rents separately. Prices are sampled in 87 geographical areas across the United States. In New York, Chicago, and Los Angeles, all prices are collected at a monthly frequency. In other locations, food and energy prices are collected monthly and the prices of other items are collected bimonthly. The CPI Research Database is described in more detail in Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008).

4.1 State-Level Price Index Construction

Our methodology for constructing price indexes is a simplified version of the procedure used by the BLS to construct the CPI. One key difference versus the BLS procedure, and a key reason why we do not simply employ the BLS's own price index software, is that we do not impute missing price observations using inflation rates calculated for other sectors or regions. We describe our procedure below.

We start by calculating price relatives for individual products. These are the fundamental building blocks of a matched-model price index. For product i at time t, the formula we use to calculate the price relative is

$$r_{i,t} = \left(\frac{P_{i,t}}{P_{i,t-\tau}} \right)^{1/\tau}. \tag{16}$$

where $r_{i,t}$ denotes the price relative, $P_{i,t}$ denote the effective price, and τ denotes the number of months since the last time a price was collected for this product. Several details are important. First, it is important to use the effective price rather than the raw "collected price." The difference between the collected and effective prices is that the latter adjusts for changes in the number and size of the items being priced (e.g. a 2L bottle of Diet Coke vs. a two-pack of 2L bottles of Diet Coke).

Second, we define a product not only by its characteristics (e.g., 2L bottle of Diet Coke), but also by the location in which it is sold. To be precise, in the CPI Research Database, each product is indexed by outlet, quote, and version. The quote is a very narrowly described product, and the version is the exact specification of the item that the price collector identifies in the store. We hold all three of these parameters—outlet, quote, and version—fixed in constructing a product's price

relative.

Third, we must decide what to do when prices are missing. Missing prices occur when the product is unavailable due to a temporary stockout, or as a consequence of the bimonthly pricing schedule used by the BLS for most products in most cities. Our procedure is to divide the price change evenly among the periods between successive price observations by taking the τ-th root of the price change and applying this price relative to all τ periods. This implies that $r_{i,t} = ... = r_{i,t-\tau+1}$ where again τ is the number of periods between successive price changes. There are several other important details of our index construction procedure that we describe in Appendix B.3.

We aggregate the price relatives in several steps. First, we compute an unweighted geometric average of the price relatives within each Entry Level Item ("ELI") product category and state. ELIs are relatively narrow product categories such as "Full Service Meals and Snacks" (restaurants) and "Motorcycles" defined by the BLS for the purpose of calculating the CPI.[14] We then calculate sectoral state-level price indexes by computing a weighted geometric average of the ELI-state indexes across the ELIs within that state and sector. We use national weights from the Consumer Expenditure Survey (CEX) for 1998 to perform this aggregation.[15]

Our empirical analysis focuses on non-tradeables but we also construct state-level price indexes for tradeables—which we simply define as the complement of non-tradeables—and overall state-level price indexes. We construct a price index for non-tradeables based on our own categorization of BLS's ELI product categories. In doing this, we attempt to be conservative in our definition of what constitutes a non-tradeable good, since including tradable goods could lead to attenuation of the slope of the Phillips curve if tradeable goods are priced nationally. In contrast, the main downside of excluding some non-tradeable goods is less precise estimates. The goods we classify as non-tradeables account for roughly 44% of non-housing consumer expenditures. Importantly, our index of non-tradeables does not include housing services or transportation goods (mainly airline tickets).[16] We estimate regional Phillips curves for housing services separately in section 5 using different data. Appendix B.4 provides a detailed list of which ELI categories we

[14] See the appendix to Nakamura and Steinsson (2008) for a list of the ELIs used in the construction of the CPI.

[15] Here we follow the BLS in using consumption weights. Rubbo (2020) argues that production networks imply that product-level inflation should be weighted by sales shares.

[16] We find that there is much more variablility across states in non-tradeable inflation than tradeable inflation. For non-tradeables, the first principal component of state-level inflation captures only about 37% of the variance in the underlying state-level series. In contract, for tradeables, the first principal component captures about 71% of the variance in the underlying state-level series. This pattern is consistent with our argument in section 3.2 that many tradeable goods are priced nationally, and do not respond to regional marginal costs.

classify as non-tradeable.

Our method for calculating state-level price indexes aims to approximate the non-shelter price index published by the BLS. Appendix Figure C.1 illustrates our ability to match the official BLS data by comparing the evolution of 12-month inflation at the aggregate level using our methodology with official CPI inflation excluding housing. The figure shows that we are able to approximate the official BLS data very closely. This is true even for the pre-1988 period when we rely on the micro-data recovered by Nakamura et al. (2018) which likely have greater measurement error.[17]

4.2 Employment data

The measure of unemployment that we use as our measure of labor market slack in the Phillips curve is the quarterly, seasonally adjusted, state unemployment rate from the Local Area Unemployment Statistics (LAUS) published by the BLS. We also make use of employment data in constructing our tradeable demand spillovers instrument discussed in section 5. This instrument is a shift-share instrument, similar to the one used in Bartik (1991). It is constructed using employment shares of individual industries at the state level.[18] We seasonally adjust the resulting series by regressing it on an exponentially weighted moving average of its lags as well as state by quarter-of-year fixed effects. We use the variation not explained by the quarter-of-year dummies as our instrument.[19] We define the tradeable employment share in the same way as Mian and Sufi (2014). Appendix B.5 discusses this in more detail.

5 Empirical Results

We now turn to our empirical results. We present estimates both of the structural parameter κ from equation (14) and ψ from equation (15). Recall that κ is structural slope coefficient in the regional Phillips curve for non-tradeables from our model, while ψ is the reduced form slope coefficient in the type of regional empirical specification often run in prior work. To be able to estimate equation (14), we replace expected future unemployment and relative prices with their realized values and

[17] In the present draft, we drop Arizona due to anomalous trends that we have not been able to investigate due Covid-19 related access restrictions at the BLS.

[18] Industry-state employment data are available from the QCEW, at quarterly frequency for 2-digit SIC codes (1975-2000) and 3-digit NAICS codes (1990-2017). Before 1990 we use 2-digit SIC codes to define industry, whereas after 2000 we use 3-digit NAICS Code. For the period 1990-2000, when both the NAICS and SIC code classifications are available, we construct both versions of the instrument and use a simple average of the two.

[19] Using the X-11 algorithm for seasonal adjustment yields virtually identical results.

an expectation error. We also truncate the infinite sum in equation (14) at $j = T$. Doing this yields

$$\pi_{it}^N = \alpha_i + \gamma_t - \kappa \sum_{j=0}^{T} \beta^j u_{i,t+j} - \lambda \sum_{j=0}^{T} \beta^j \hat{p}_{i,t+j}^N + \tilde{\omega}_{it}^N + \eta_{it}^N, \tag{17}$$

where η_{it}^N denotes an expectations error that is orthogonal to information known at time t (and a truncation error). Equation (17) can now be estimated with standard GMM methods, i.e., by instrumenting for the two forward sums. We do not attempt to estimate β. Rather, we set it to a standard quarterly value of $\beta = 0.99$.

We present results for two approaches to identifying the coefficients κ and λ in equation (17). Our first approach is to instrument for the two forward sums with 4-quarter lagged unemployment $u_{i,t-4}$ and the 4-quarter lagged relative price of non-tradeables $\hat{p}_{i,t-4}^N$. Assuming rational expectations, these lagged variables will be uncorrelated with the expectations error η_{it}^N. The identifying assumption regarding supply shocks is that when one state experiences a boom or bust relative to another state, it does not systematically experience non-tradeable supply shocks relative to this other state. For example, when Texas experiences a recession relative to Illinois, this is not systematically correlated with changes in restaurant technology in Texas relative to Illinois. Notice, that national supply shocks are absorbed by the time fixed effects. So, only regional non-tradeable supply shocks are potential confounders.

Our second approach to identification is to construct an instrumental variable that captures variation in demand. The idea behind our instrumental variable is the notion that national variation in demand for specific tradeable goods will differentially affect labor demand for non-tradeable goods in states that produce those tradeable goods. For example, an increase in oil prices will differentially affect labor demand in Texas (and other oil producing states). As a result, wages in Texas will rise differentially affecting costs of non-tradeables in Texas. Building on this idea, we construct a "tradeable demand spillovers" instrument as

$$\text{Tradable Demand}_{i,t} = \sum_{x} \bar{S}_{x,i} \times \Delta_{3Y} \log S_{-i,x,t}, \tag{18}$$

where $\bar{S}_{x,i}$ is the average employment share of industry x in state i over time, and $\Delta_{3Y} \log S_{-i,x,t}$ is the three-year growth in national employment of industry x at time t excluding state i. This shift-share instrument builds on Bartik (1991) and more closely on Nguyen (2014). The identifying assumption in this case is that there are no supply factors that are both correlated with the shifts

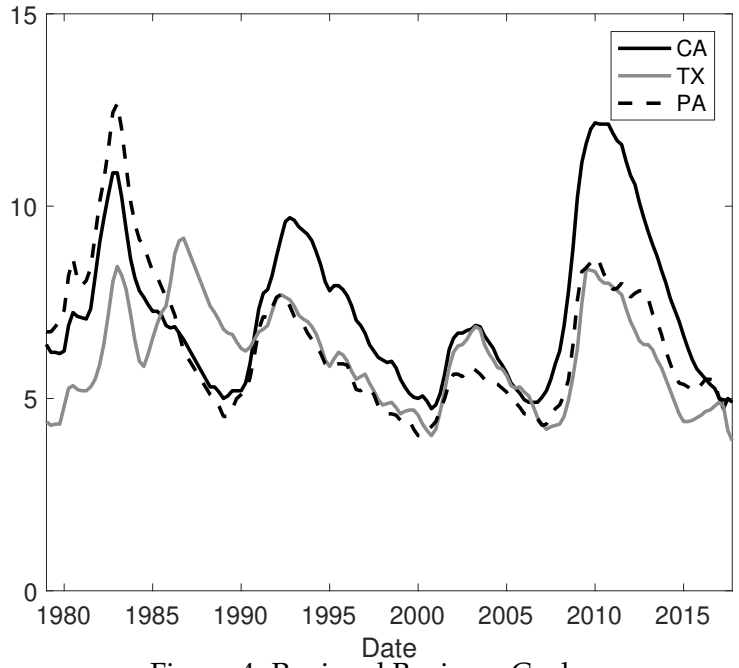


Figure 4: Regional Business Cycles

Note: This figure plots the unemployment rate for California, Pennsylvania and Texas.

$\Delta_{3Y} \log S_{-i,x,t}$ in the time series and correlated with the shares $\bar{S}_{x,i}$ in the cross section. For example, costs will increase as a result of an increase in oil prices. But if such cost increases are no larger on average for restaurants in Texas than Illinois they will be uncorrelated with our instrument.[20]

Our panel data approach implies that we are relying on cross-state variation in unemployment to identify the slope of the Phillips curve. Figure 4 depicts the evolution of the unemployment rate for three states, California, Texas and Pennsylvania, over our sample period. While there is certainly a great deal of comovement, this figure illustrates well that there is also substantial cross-state variation. One example is that both the 1991 and 2007-2009 recessions affected California much more than Texas and Pennsylvania. Another is that Texas experienced a recession in the mid-1980s (widely thought to stem from a dramatic fall in oil prices) while most other states experienced a continued fall in unemployment. Estimates of unemployment at the state level may be plagued by measurement error. Our IV estimation will address this insofar as the measurement error is classical.

The dependent variable in our regressions is $\pi_{it}^N = p_{it}^N - p_{i,t-4}^N$, i.e., state-level non-tradeable inflation over the previous 12 months. Studying inflation over four quarters allows us to reduce measurement error and eliminate seasonality. In Appendix A.11, we show that using twelve-

[20] In a related approach, McLeay and Tenreyro (2019) use identified demand shocks from government spending to estimate the slope of the regional Phillips Curve.

month inflation as our dependent variable implies that we need to divide our estimates of κ and λ from equation (17) by four to account for the time aggregation. Recall that the inflation rate in our model in section 3 is a quarterly inflation rate.

We truncate the discounted sums on the right-hand-side of equation (17) at $T = 20$ quarters. Table C.4 presents robustness regarding this choice for our main specification. Our results are similar for values between $T = 20$ and $T = 40$. In Appendix A.12, we estimate κ using equation (17) with $T = 20$ on data simulated from our model from section 3. We find that our empirical procedure is able to accurately estimate the true value κ in this setting for a very wide range of true values of κ.

The forward sums in equation (17) imply that we lose five years of observations at the end of our sample when we set $T = 20$. To minimize the impact of this, we use a two-sample two stage least squares regression. We estimate the first stage on a reduced sample without the last five years, and the second stage on the full sample. We cluster standard error at the state level and apply a correction to our standard errors appropriate for two-sample 2SLS developed by Chodorow-Reich and Wieland (2019).[21]

Our empirical specification for estimating ψ is

$$\pi_{it}^N = \alpha_i + \gamma_t - \psi u_{i,t-4} - \delta p_{i,t-4}^N + \varepsilon_{it}. \tag{19}$$

We use beginning-of-period unemployment and relative price of non-tradeables as regressors for consistency with previous studies such as Ball and Mazumder (2019). We present results for two identification approaches analogous to those we use for κ. The first approach is to estimate equation (19) by OLS (i.e., instrumenting for lagged unemployment and relative price of non-tradeables with themselves). The second approach replaces lagged unemployment among the instruments with our tradeable-demand instrument.

5.1 Full-Sample Results

Table 1 presents estimates of κ and ψ for our full sample period of 1978-2018. Let's start by considering the estimates of κ in Panel A. When we estimate equation (17) without fixed effects, our estimate has the "wrong" sign, i.e., higher unemployment is associated with higher rather than

[21] Our tradeable demand instrument uses all the information in national industry employment growth rates. So, our standard errors are not subject to the concerns about inference with shift share instruments raised by Adao, Kolesar, and Morales (2019).

Table 1: Slope of the Regional Phillips Curve

	No Fixed Effects	No Time Effects	Lagged Unempl.	Tradeable Demand IV
	(1)	(2)	(3)	(4)
	Panel A: Estimates of κ from equation (17)			
κ	-0.0037	0.0003	0.0062	0.0062
	(0.0013)	(0.0019)	(0.0028)	(0.0025)
	Panel B: Estimates of ψ from equation (19)			
ψ	-0.103	0.017	0.112	0.339
	(0.036)	(0.027)	(0.057)	(0.126)
State Effects		✓	✓	✓
Time Effects			✓	✓

Note: This table presents estimates of κ and ψ from regression specifications (17) and (19), respectively. The outcome variable is cumulative non-tradeable inflation over four quarters, measured in percentage points. We include time and state fixed effects as noted at the bottom of each column. In Panel A, the regressors are the discounted future sum of quarterly state unemployment, in percentage points, and the discounted future sum of the relative price of non-tradeables, in 100 x log points. For both variables, we truncate the discounted furture sum at 20 quarters. In Panel B, the regressors are the fourth lags of quarterly state unemployment, measured in percentage points, and the relative price of non-tradeables. In the first three columns we instrument using the fourth lags of quarterly state unemployment and the relative price of non-tradeables (this is OLS for ψ). In the fourth column, we replace lagged unemployment with our tradeable demand instrument among the instruments. In all columns, we estimate κ by two-sample two stage least squares, and apply the correction to our standard errors from Chodorow-Reich and Wieland (2019). The sample period is 1978-2018. Standard errors are reported in parentheses, clustered by state. All regressions are unweighted. The number of observations is 3323 in the first three columns of panel A, with slightly fewer in the last column due to differencing. Likewise, the number of observations is 4490 in the first three columns of panel B.

lower inflation ($\kappa = -0.0037$). Adding state fixed effects raises the estimate of κ to 0.0003. Adding time fixed effects further raises the estimate of κ to 0.0062. As we stress throughout the paper, time fixed effects eliminate changes in long-run inflation expectations. Finally, using our tradeable demand instrument as opposed to instrumenting with lagged unemployment yields virtually the same estimate for κ of 0.0062. The fact that our estimate of κ does not change between columns (3) and (4) suggests that the fixed effects we include are sufficient to absorb supply shocks.

Our estimated slope of the Phillips curve is statistically significantly different from zero. In absolute size, however, the slope is small in the sense that it is consistent with the modest response of inflation to changes in unemployment seen in the aggregate time series since 1990. We develop this implication in section 5.4. Table C.1 presents estimates of the "first stage" regressions for our IV estimates of equation (17). These first stage regressions show that our instruments are strong instruments. We separately regress the present value of unemployment, and the present value of relative prices, on the reduced form regressors. Lagged unemployment and tradeable demand

both strongly predict the present value of unemployment and weakly predict the present value of relative prices. Lagged relative prices strongly predict the present value of relative prices and weakly predict the present value of unemployment.

Table C.2 reports our estimates of λ for regression specification (17)—the coefficient on the relative price of non-tradeables. In our preferred specifications with time and state fixed effects, we estimate values of λ between 0.002 and 0.003. In the model we present in section 3, λ provides an estimate of the degree of nominal rigidities. In a world with flexible prices, our estimate of λ would be large. The fact that our estimate of λ is very small provides further support—over and above our estimate of κ—for the notion that prices are quite rigid in the U.S. economy.

In our baseline results, we calibrate $\beta = 0.99$. It may, however, be that firms are considerably less forward looking when they set prices than this calibration implies. Recent work has shows that plausible deviations from full rationality or common knowledge yield a Phillips curve that is less forward looking (Angeletos and Lian, 2018; Gabaix, 2020). Also, a model with a combination of sticky information and sticky prices yields a Phillips curve that is less forward looking. Table C.3 presents estimates of κ where we calibrate β to lower values. As we vary our quarterly calibration of β from 0.99 to 0.9, κ doubles in size. The absolute size of the increase is small because our initial estimate of κ is small.

Our estimates of ψ in Panel B of Table 1 have a similar pattern to our estimates of κ discussed above. The estimate without time or state fixed effects is negative and the estimate increases as we include state and then time fixed effects. An important difference is that the absolute size of our estimates of ψ are much larger than our estimates of κ. This reflects the fact that in equation (19) the lagged unemployment rate is standing in for the entire future sum in equation (17). Since unemployment is quite persistent, time variation in the future sum is much larger than time variation in the unemployment rate, which results in a much larger coefficient in equation (19) than in equation (17).

Another difference is that ψ is much larger in column (4) than in column (3), while κ is virtually identical. This reflects the fact that the tradeable demand instrument we use in column (4) is more persistent than the unemployment rate itself. The coefficients in column (4) are therefore identified using more persistent variation which results in a larger value of ψ, but not a larger value of κ. This highlights an important advantage of estimating κ as opposed to ψ: estimates of ψ are hard to interpret because they are sensitive to the persistence of the variation that is used to identify them. More generally, κ is a structural parameter, while ψ is not. This implies that ψ may differ

Table 2: Has the Phillips Curve Flattened?

	Lagged Unempl. IV Without Time Fixed Effect		Lagged Unempl. IV With Time Fixed Effect		Tradeable Demand IV With Time Fixed Effect	
	Pre-1990 (1)	Post-1990 (2)	Pre-1990 (3)	Post-1990 (4)	Pre-1990 (5)	Post-1990 (6)
	Panel A: Estimates of κ from equation (17)					
κ	0.0278	0.0002	0.0107	0.0050	0.0109	0.0055
	(0.0025)	(0.0017)	(0.0080)	(0.0040)	(0.0062)	(0.0028)
	Panel B: Estimates of ψ from equation (19)					
ψ	0.449	0.009	0.198	0.090	0.422	0.332
	(0.063)	(0.025)	(0.113)	(0.057)	(0.232)	(0.157)

Note: The table presents estimates of κ and ψ, before and after 1990. Columns (1), (3) and (5) present results for the sample period 1978-1990; and columns (2), (4) and (6) for the sample period 1991-2018. All specifications include state fixed effects. Specifications in columns (3)-(6) include time fixed effects. The instruments in columns (1)-(4) are the fourth lag of the tradeable demand instrument and the relative price of non-tradeables (i.e., OLS in Panel B). In columns (5) and (6), the instrument are the fourth lag of the tradeable demand instrument and the relative price of non-tradeables. In all columns, we estimate κ by two-sample two stage least squares and apply the correction to our standard errors from Chodorow-Reich and Wieland (2019). Standard errors are reported in parentheses, clustered by state. All regressions are unweighted.

depending on the setting being considered (e.g., may be low in response to a policy change that may be reversed due to a future change in government), while κ is policy invariant.

5.2 Subsample Results

We next analyze to what extent the Phillips curve was steeper during the period of the Volcker disinflation than in subsequent years. Table 2 presents estimates of κ and ψ for the periods 1978-1990 and 1991-2018. We present these estimates for specifications with and without time fixed effects. All specifications include state fixed effects and control for the relative price of non-tradeables.

Consider first the specification without time fixed effects reported in columns (1) - (2). For the pre-1990 sample, κ is estimated to be 0.0278, while ψ is estimated to be 0.449. In sharp contrast, for the post-1990 sample, κ is estimated to be 0.0002 and ϕ is estimated to be 0.009. The difference across samples is roughly a factor of 100 for κ and 50 for ψ. In other words, aggregate inflation became much less sensitive to unemployment after 1990 than it was during the Volcker disinflation.

Contrast this with the results in columns (3) - (4) where time fixed effects are included in the regressions. In this case, the estimated values of κ and ψ fall only modestly between the early

part of the sample and the later part of the sample. For the pre-1990 sample, κ is estimated to be 0.0107 and ψ is estimated to be 0.198. For the post-1990 sample, κ is estimated to be 0.0050 and ψ is estimated to be 0.090. The difference across samples is roughly a factor of two and is not statistically significant. The estimate for κ in columns (5) and (6) are very similar to the estimates in columns (3) and (4), while the estimate of ψ in columns (5) and (6) show an even smaller difference across sample periods.

As we emphasize in section 2, estimates of the Phillips curve based on time-series variation — such as the estimates without time fixed effects in Table 2 — are likely to be heavily influenced by time-series variation in long-run inflation expectations $E_t\pi_{t+\infty}$. In contract, the specifications in Table 2 that include time fixed effects difference out the influence of long-run inflation expectations. The results in Table 2 therefore suggest that the apparent flattening of the Phillips curve in the time series is largely due to inflation expectations becoming more firmly anchored over time. In the early part of the sample, inflation expectations shifted a great deal and these shifts were negatively correlated with the unemployment rate, which meant that shifts in inflation expectations masqueraded as a steep Phillips curve. The cross-sectional results in columns (3) - (6) of Table 2 reveal that in fact the Phillips curve has always been quite flat (at least since 1978).

Figure 5 provides a visual representation of the results in Table 2. In the left panel, we plot a binned scatterplot of state-level non-tradeable inflation against state-level unemployment after removing state fixed effects and the effects of the relative price of non-tradeables. We plot the data separately for the period 1978-1990 and 1991-2018. The plot also includes regression lines for each subsample. The data in this panel does not account for time fixed effects and therefore includes aggregate time-series variation. As a consequence, we see a huge flattening of the Phillips curve in this case.

Contrast this with the right panel in Figure 5. This is an analogous figure to the left panel except that we also demean by time fixed effects. These data therefore only reflect regional variation in inflation. In this case, the difference in the slope of the Phillips curve between the early sample and the late sample is modest. The modest flattening of the Phillips curve that we find over our sample (once we account for time fixed effects) seems consistent with the fact that the frequency of price change in the U.S. has declined by about 40% as inflation has fallen since the early 1980's (Nakamura et al., 2018).

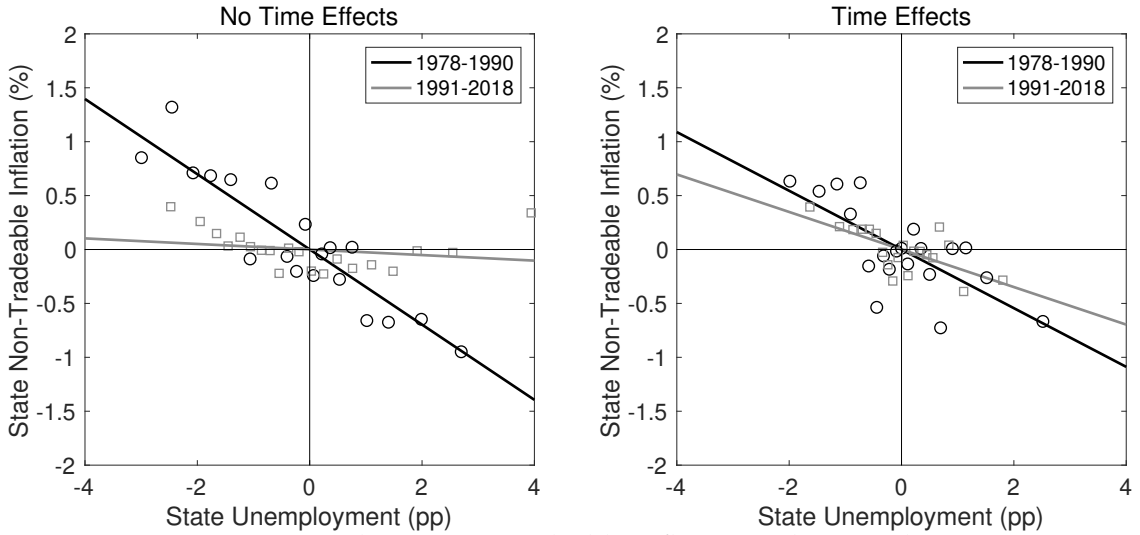


Figure 5: Scatterplots—Non-Tradeable Inflation and Unemployment

Note: in the the left hand size panel we residualize state non-tradeable inflation and unemployment against state fixed effects and the relative price of non-tradeables, before and after 1990. We then plot residualized inflation and unemployment, before and after 1990, grouped by 20 bins of state unemployment. The right hand side panel carries out the same exercise after further residualizing against time fixed effects. In both panels inflation is cumulated over the previous four quarters and unemployment is lagged by four quarters.

5.3 How Do Our Estimates Compare to Prior Work?

It is instructive to compare our estimate of κ to values of κ arrived at by means of structural estimation or calibration of New Keynesian models. Table 3 reports three such estimates from Rotemberg and Woodford (1997), Gali (2008), and Nakamura and Steinsson (2014). In all cases, we have adjusted the reported value of κ in these papers by the elasticity of output with respect to employment in the models used in these papers. As is well known, the value of κ in a New Keynesian model is highly dependent on both the degree of nominal and real rigidities assumed. The values for κ used in these papers ranges from about an order of magnitude larger than our estimated value to a value roughly equal to our estimated value. The main difference between Gali's relatively high value and the much lower values in Rotemberg and Woodford (1997) and Nakamura and Steinsson (2014) lies in the degree of real rigidity that the models used in these papers imply. Gali's model is a relatively simple (textbook) version of the New Keynesian model, which does not incorporate strong sources of real rigidity. Rotemberg and Woodford (1997) and Nakamura and Steinsson (2014) use models with heterogeneous labor markets, which yields a much larger amount of real rigidity. In both cases, the large amount of real rigidity helps these authors match moments that they target in their analysis. Similarly, our estimates imply that the data we have analyzed is also more consistent with New Keynesian models that incorporate a

Table 3: Our Estimates Compared to Prior Work

	κ
Rotemberg and Woodford (1997)	0.019
Gali (2008)	0.085
Nakamura and Steinsson (2014)	0.0077
Our Estimate	
Full Sample IV Estimate	0.0062

Note: We adjust the estimates from Rotemberg and Woodford (1997), Gali (2008), and Nakamura and Steinsson (2014) by the elasticity of output with respect to employment in the model in these papers. For Nakamura and Steinsson (2014), we use the calibration with GHH preferences.

large amount of real rigidity.

5.4 Aggregate Implications

A question that naturally arises regarding our cross-sectional estimates of κ is whether they can explain the aggregate time-series variation in inflation over our sample. A number of researchers and commentators have suggested that the stability of inflation at the aggregate level in the U.S. has been surprising over the past 25 years ("missing disinflation" during the Great Recession and "missing reinflation" during the late 1990s and late 2010s). Some researcher have recently argued that cross-sectional variation suggests a steeper Phillips curve than time-series variation for the past few decades. Here, we assess whether this is the case for our estimates.

We start with the solved-forward aggregate Phillips curve—equation (6). In section 2, we made the simplifying assumption that the unemployment rate follows an AR(1). This assumption allowed us to derive a simple aggregate relationship between the discounted future sum of unemployment rates in equation (6) and the current unemployment rate—see equation (7). In reality, however, the dynamics of the US unemployment rate differ substantially from an AR(1) (see, e.g., Neftçi, 1984; Sichel, 1993; Dupraz, Nakamura, and Steinsson, 2020). For this reason, we adopt an approach of estimating a scaling factor ζ that relates the current unemployment rate to the discounted future sum in equation (6) using the following regression

$$\sum_{j=0}^{T} \beta^j \tilde{u}_{t+j} = \zeta \tilde{u}_t + \alpha + \epsilon_t. \tag{20}$$

The series we use for $\tilde{u}_t$ in this regression is the difference between the aggregate unemployment

rate in the U.S. and the CBO's estimate of the natural rate of unemployment at each point in time.[22] We run this regression for the sample period 1979Q4-2017Q4. This yields an estimate of ζ for aggregate variation in the unemployment rate of 6.16 with a Newey-West standard error of 1.80. Using equation (20), we can rewrite equation (6) as

$$\pi_t - E_t\pi_{t+\infty} = -\kappa\zeta\tilde{u}_t + \omega_t. \tag{21}$$

Our cross-sectional estimates of κ are for non-tradeables excluding housing services. As we emphasize in section 2, the treatment of housing services has important implications for the behavior of inflation. Appendix Table C.5 presents estimates of κ and ψ using state-level annual rent inflation data from the American Community Survey for the years 2001 to 2017. For our baseline specification with state and time fixed effects, we estimate κ to be 0.0243. This estimate of κ is roughly four times larger than our estimate of κ for non-housing non-tradeable goods reported in Table 1. We account for this difference below by taking a weighted average of our full-sample κ estimate for non-tradeables and this κ estimate for housing services.[23]

Figure 6 plots the left-hand side of equation (21) (black line) against the first term on the right-hand side of equation (21) (gray line) using our estimates of κ and ζ from above, which yields $\kappa\zeta = 0.34$.[24] We use the 10 year ahead SPF inflation expectations for the CPI as our measure of long-term inflation expectations. The gray line is the demand-induced variation in inflation predicted by our estimates. The figure indicates that the amplitude of inflation fluctuations over the last few business cycles has been roughly in line with what our cross-sectional estimates of κ suggest. In particular, the disinflation during the Great Recession and reinflation during the 2010s lines up well with what our estimate of κ implies. If the gray line had a larger amplitude than the black line over the business cycle, this would indicate missing disinflation and missing reinflation. In fact, the amplitude of the gray line is very similar to that of the black line for the Great Recession, the post-Great Recession recover, and the long 1990s expansion. By this metric, there is thus no missing disinflation or missing reinflation over this period. These findings echo the results of Ball and Mazumder (2019).[25]

[22] We are thus treating the CBO's estimate of the natural rate of unemployment as a forecast of long run unemployment $E_t u_{t+\infty}$.

[23] We use the shelter and non-shelter expenditure weights in the core CPI. These are 0.42 and 0.58, respectively.

[24] The coefficient $\kappa\zeta$ is calculated as $4 \times (0.58 \times 0.0062 + 0.42 \times 0.0243) \times 6.16$, where the factor of 4 accounts for time aggregation to annual inflation.

[25] Figure C.2 shows that modest flattening of the Phillips curve we estimate over our sample period has a minimal effect on the fluctuations in the gray line in Figure 6. Figure C.3 shows that a disproportionate share of the systematic

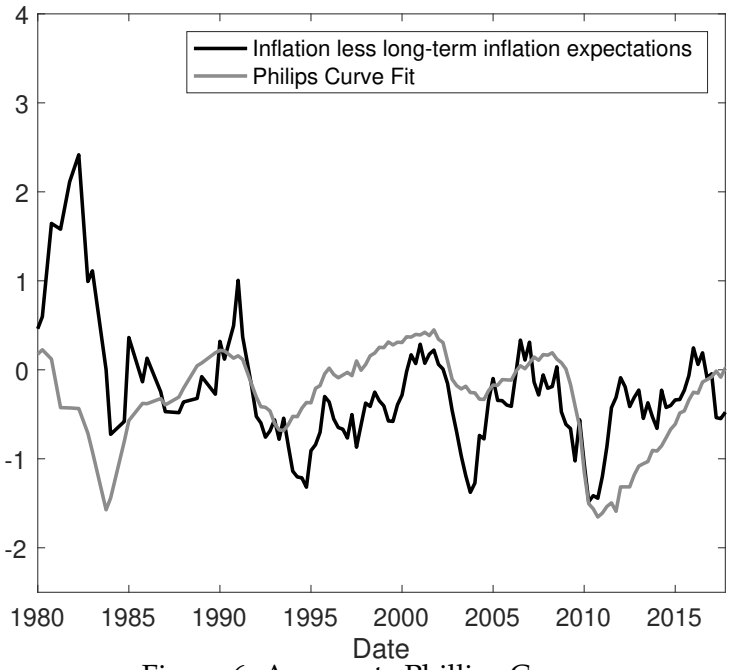


Figure 6: Aggregate Phillips Curve

Note: The figure shows the fit of the aggregate Phillips curve for core inflation. The black line is the difference between core inflation and the 10 year ahead SPF inflation expectation for the CPI. The gray line plots the first term on the right-hand side of equation (21) which is the demand-induced variation in inflation predicted by our estimates.

The most substantial deviation between the actual and fitted values arises during the Volcker period when actual inflation relative to long-run expectations lies far above the fitted value. While the conventional view is that the Phillips curve has broken down after 1990, we are finding the opposite: a poor fit of our cross-sectional estimate of the Phillips curve when applied to aggregate inflation dynamics over the Volcker period. A natural interpretation of this discrepancy is the presence of adverse supply shocks in the early 1980s, for example, associated with the oil price shocks.

How much of the fall in inflation during the Volcker disinflation can be attributed to the causal effect of higher unemployment working through the slope of the Phillips curve according to our estimates? Unemployment rose by about 5 percentage points between 1979 and 1982. Using a weighted average of our slightly higher pre-1990 non-shelter estimate for κ and our estimate of κ for shelter, we find that this increase in unemployment caused inflation to fall by only about 2 percentage points (see gray line in Figure C.4). Core CPI inflation first rose from 7% to 10%

variation in inflation and the fitted value predicted by our model comes from the housing services (rent) component of the CPI. The figure is analogous to Figure 6 except that the black line excludes housing and the fitted value uses only the κ estimate for the non-shelter component of inflation. We see that core inflation excluding housing services varies much less systematically than core inflation including housing services.

from 1979 to 1981 and then fell to 4% by 1986. Clearly, the direct causal effect of unemployment working through the slope of the Phillips curve explains only a modest amount of this variation in inflation. Over this same period, long-run inflation expectations first rose from 7% to 8% and then fell to 4%. Our estimates, therefore, suggest that the bulk of the variation in inflation over the early 1980s is due to changes in long-run inflation expectations, with supply shocks also playing an important role.

6 Conclusion

This paper provides new estimates of the slope of the Phillips curve. We estimate that the slope of the Phillips curve is small, and was small even during the Volcker disinflation of the early 1980s. Our results indicate that shifts in expectations about the conduct of monetary policy explain much of the drop of inflation in the early 1980s and more firmly anchored inflation expectations explain the stability of inflation since the mid-1990s. Our estimates are consistent with the insensitivity of inflation to unemployment during the both the Great Recession and during the low unemployment periods of the late 1990s and late 2010s.

To reach these conclusions, we estimate the Phillips curve in the cross-section of U.S. states. We use newly constructed state-level price indexes for non-tradeable goods starting in 1978. We map from our regional estimates to the slope of the aggregate Phillips curve using a multi-region New Keynesian model. The model clarifies that the slope of the aggregate Phillips curve is equal to the slope of the regional Phillips curve for non-tradeable goods. We also use the model to show that regional data "difference out" the effects of the long-run monetary regime, which otherwise confound estimates of the slope of the Phillips curve. Guided by the model, we show that the conventional empirical specification used to estimate regional Phillips curves must be scaled by a factor relating to the persistence of unemployment fluctuations to yield an estimate of the slope of the Phillips curve. Finally, we develop a new "tradeable demand spillover" instrument that allows for flexible patterns of supply shocks at the local level.

An important lesson from our analysis is that when it comes to managing inflation, the elephant in the room is long-run inflation expectations. This view contrasts sharply with the conventional view that managing inflation is about moving up and down a steep Phillips curve. A crucial question for inflation dynamics is why long-run inflation expectations are sometimes so firmly anchored but at other times move sharply? Beliefs about inflation in the long run are gov-

erned by beliefs about the long-run behavior of the monetary authority and ultimately the political process that shapes the long-run behavior of the monetary authority. Since this is fundamentally a very low-frequency phenomenon, it is not easily pinned down by half a century or so of data from a single country. While much interesting research has sought to understand the behavior of long-run inflation expectations, we believe it is still not sufficiently well understood and its crucial importance for the conduct of monetary policy implies that even more research should focus on this question.

A Theoretical Appendix

A.1 Household Optimality Conditions

Households optimally trade off current consumption and current labor supply. This implies the following labor supply curve must hold in our model:

$$-\frac{u_n(C_{Ht}, N_{Ht})}{u_c(C_{Ht}, N_{Ht})} = \frac{W_{Ht}}{P_{Ht}},$$

where P_{Ht} denotes the lowest cost of purchasing a unit of the composite consumption good C_{Ht} and subscripts on the utility function denote partial derivatives. Using expressions for $u_n(C_{Ht}, N_{Ht})$ and $u_c(C_{Ht}, N_{Ht})$, we can rewrite the home labor supply curve as

$$\chi N_{Ht}^{\varphi^{-1}} = \frac{W_{Ht}}{P_{Ht}}. \tag{22}$$

Households optimally trade off current consumption and consumption in the next period. This implies the following consumption Euler equation must hold in our model:

$$\beta R_t^n E_t \left[\frac{u_c(C_{H,t+1}, N_{H,t+1})}{u_c(C_{Ht}, N_{Ht})} \frac{P_{Ht}}{P_{H,t+1}} \right] = 1. \tag{23}$$

where R_t^n is the gross nominal interest rate, which is common to both regions in the monetary union. Household optimization also implies a standard transversality condition must hold in model and it implies that the stochastic discount factor takes a standard form.

Households choose how much to purchase of the various goods in the economy to minimize the cost of attaining the level of consumption C_{Ht} they choose. This implies the following demand curves for home and foreign tradeable and non-tradeable goods:

$$C_{Ht}^N = \phi_N C_{Ht} \left(\frac{P_{Ht}^N}{P_{Ht}} \right)^{-\eta}, \tag{24}$$

$$C_{Ht}^{TH} = \phi_T \tau_{Ht}^H C_{Ht} \left(\frac{P_{Ht}^T}{P_{Ht}} \right)^{-\eta}, \quad \text{and} \quad C_{Ht}^{TF} = \phi_T \tau_{Ht}^F C_{Ht} \left(\frac{P_{Ft}^T}{P_{Ht}} \right)^{-\eta}. \tag{25}$$

Utility maximization, furthermore, implies the following demand curves for each of the varieties

of goods produced in the economy:

$$C_{Ht}^N(z) = C_{Ht}^N \left(\frac{P_{Ht}^N(z)}{P_{Ht}^N} \right)^{-\theta} \quad C_{Ht}^{TH}(z) = C_{Ht}^{TH} \left(\frac{P_{Ht}^T(z)}{P_{Ht}^T} \right)^{-\theta} \quad C_{Ht}^{TF}(z) = C_{Ht}^{TF} \left(\frac{P_{Ft}^T(z)}{P_{Ft}^T} \right)^{-\theta} \tag{26}$$

The cost minimizing price indexes are given by

$$P_{Ht}^N = \left[\int_0^1 P_{Ht}^N(z)^{1-\theta} dz \right]^{\frac{1}{1-\theta}}, \quad P_{Ht}^T = \left[\int_0^1 P_{Ht}^T(z)^{1-\theta} dz \right]^{\frac{1}{1-\theta}}, \quad P_{Ft}^T = \left[\int_0^1 P_{Ft}^T(z)^{1-\theta} dz \right]^{\frac{1}{1-\theta}},$$

$$P_{Ht} = \left[\phi_N {P_{Ht}^N}^{1-\eta} + \phi_T \tau_{Ht}^H {P_{Ht}^T}^{1-\eta} + \phi_T \tau_{Ht}^F {P_{Ft}^T}^{1-\eta} \right]^{\frac{1}{1-\eta}}. \tag{27}$$

A.2 Firm Optimality Conditions

Firms in our model must satisfy demand. For firms in the non-tradeable sector, this implies that

$$\zeta C_{Ht}^N \left(\frac{P_{Ht}^N(z)}{P_{Ht}^N} \right)^{-\theta} \leq Z_{Ht}^N N_{Ht}^N(z).$$

Optimal choice of labor by the firm z implies that

$$W_{Ht} = S_{Ht}^N(z) Z_{Ht}^N, \tag{28}$$

where $S_{Ht}^N(z)$ is the firm's nominal marginal cost, i.e. the Lagrange multiplier on the constraint above.

Non-tradeable firms z that are able to reoptimize their price in period t set it to satisfy

$$\sum_{k=0}^{\infty} \alpha^k E_t \left[M_{Ht,t+k} Y_{H,t+k}^N(z) \left(P_{Ht}^{N*}(z) - \frac{\theta}{\theta - 1} S_{H,t+k}^N(z) \right) \right] = 0, \tag{29}$$

where $P_{Ht}^{N*}(z)$ is the price the firm chooses. Intuitively, the firm sets its price equal to a constant markup over a weighted average of current and expected future marginal cost taking into account the probability that their price will remain unchanged in future periods. We can divide both sides of this equation by $P_{H,t-1}^N$ and rewrite it as

$$\sum_{k=0}^{\infty} \alpha^k E_t \left[M_{Ht,t+k} Y_{H,t+k}^N(z) \left(\frac{P_{Ht}^{N*}(z)}{P_{H,t-1}^N} - \frac{\theta}{\theta - 1} MC_{H,t+k}^N(z) \frac{P_{H,t+k}^N}{P_{H,t-1}^N} \right) \right] = 0, \tag{30}$$

where $MC^N_{H,t+k}(z) = S^N_{H,t+k}(z)/P^N_{H,t+k}$ is real marginal cost in the non-tradeable sector.

Analogously to the non-tradeable sector, optimal firm labor demand in the tradeable sector is

$$W_{Ht} = S^T_{Ht}(z) Z^T_{Ht}(z). \tag{31}$$

where $S^T_{Ht}(z)$ is the tradeable goods firm's nominal marginal cost. Optimal choice of a new reset price by tradeable goods firms implies

$$\sum_{k=0}^{\infty} \alpha^k E_t \left[M_{Ht,t+k} Y^T_{H,t+k}(z) \left(P^T_{Ht}(z) - \frac{\theta}{\theta - 1} S^T_{H,t+k}(z) \right) \right] = 0. \tag{32}$$

We can divide both sides of this equation by $P^T_{H,t-1}$ and rewrite it as

$$\sum_{k=0}^{\infty} \alpha^k E_t \left[M_{Ht,t+k} Y^T_{H,t+k}(z) \left(\frac{P^T_{Ht}(z)}{P^T_{H,t-1}} - \frac{\theta}{\theta - 1} MC^T_{H,t+k}(z) \frac{P^T_{H,t+k}}{P^T_{H,t-1}} \right) \right] = 0 \tag{33}$$

where $MC^T_{H,t+k} = S^T_{H,t+k}(z)/P^T_{H,t+k}$ is the real marginal cost in the tradeable sector.

A.3 Zero Inflation Steady State

In the next section, we take a log-linear approximation of the equilibrium conditions of our model around a steady state with zero inflation and balanced trade. In this section, we solve for this steady state. The steady state of equations (28) and (31) are $W_H = S^N_H = S^T_H$. The steady state of equations (29) and (32) imply

$$\frac{W_H}{P_H} = \frac{\theta - 1}{\theta}. \tag{34}$$

The steady state of equation (22) implies

$$\chi N_H^{\varphi^{-1}} = \mu^{-1}, \tag{35}$$

where N_H is the steady state per capita employment of households in the home region and $\mu = \theta/(\theta - 1)$. The steady state of equation (23) implies $\beta R^n = 1$.

Since all firms face the same marginal cost in the steady state we consider, all prices will be equal and all relative goods prices will be one. This implies that in steady state the demand

curves for home tradeable and non-tradeable goods—equations (24) and (25)—imply

$$C_H^N = \phi_N C_H, \tag{36}$$

$$C_H^{TH} = \phi_T \tau_H^H C_H, \tag{37}$$

$$C_H^{TF} = \phi_T \tau_H^F C_H. \tag{38}$$

We define total labor in the home non-tradeable and tradeable sectors as $N_{Ht}^N \equiv \int_0^1 N_{Ht}^N(z)dz$ and $N_{Ht}^T \equiv \int_0^1 N_{Ht}^T(z)dz$, respectively. In steady state, total non-tradeable employment, N_H^N, must equal total non-tradeable consumption, $\zeta\phi_N C$. This implies that

$$N_H^N = \zeta C_H^N = \zeta\phi_N C$$

where the second equality substitutes in equation (36). Similarly, steady state home tradeable employment, N_H^T, equals steady state total consumption for home tradeables. This implies that

$$N_H^T = \zeta C_H^{TH} + (1-\zeta)\, C_F^{TH}.$$

Using equations (37) and (38), we get that

$$N_H^T = \zeta\phi_T\tau_H^H C_H + (1-\zeta)\phi_T\tau_F^H C_F$$

Using that fact that $C_F = C_H = C$ in a symmetric steady state, we get that

$$N_H^T = \zeta\phi_T C\left(\tau_H^H + \frac{1-\zeta}{\zeta}\tau_F^H\right)$$

Finally, using the fact that $\tau_H^H = \tau_F^H = \zeta$ (no home bias in steady state), we get that

$$N_H^T = \zeta\phi_T C.$$

A.4 Derivation of Regional Phillips Curves

A first order Taylor-series expansion of equation (30) around the zero inflation and balanced trade steady state yields

$$p_{Ht}^{N*}(z) - p_{H,t-1}^{N} = (1-\alpha\beta)\sum_{k=0}^{\infty}(\alpha\beta)^{k} E_t\left[\widehat{mc}_{H,t+k}^{N} - \left(p_{H,t+k}^{N} - p_{H,t-1}^{N}\right)\right].$$

Rearranging this equation yields

$$p_{Ht}^{N*}(z) - p_{H,t-1}^{N} = \alpha\beta E_t\left[p_{H,t+1}^{N*}(z) - p_{Ht}^{N}\right] + (1-\alpha\beta)\,\widehat{mc}_{Ht}^{N} + \pi_{Ht}^{N}. \tag{39}$$

The expression for P_{Ht}^{N} given in the line above equation (27) implies

$$P_{Ht}^{N1-\theta} = \alpha P_{H,t-1}^{N1-\theta} + (1-\alpha)\,P_{Ht}^{N*1-\theta}$$

where P_{Ht}^{N*} denotes the reset price of firms that are able to change their price in period t. Here we exploit the fact that a random set of firms change their prices at time t and all of these firms set the same price. A first order Taylor series approximation of this last expression is

$$p_{Ht}^{N} = \alpha p_{H,t-1}^{N} + (1-\alpha)\,p_{Ht}^{N*}$$

which implies that

$$\pi_{Ht}^{N} = (1-\alpha)\left(p_{Ht}^{N*} - p_{H,t-1}^{N}\right). \tag{40}$$

Manipulation of equations (39) and (40) yields that

$$\pi_{Ht}^{N} = \beta E_t \pi_{H,t+1}^{N} + \lambda\widehat{mc}_{Ht}^{N} \tag{41}$$

where

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

We can derive an analogous equation to equation (41) for the tradeable sector. In the tradeable sector we have that

$$\pi_{Ht}^{T} = \beta E_t \pi_{H,t+1}^{T} + \lambda\widehat{mc}_{Ht}^{T}, \tag{42}$$

where $\pi_{Ht}^{T} = p_{Ht}^{T} - p_{H,t-1}^{T}$ is producer price inflation in the home tradeable sector.

Taking logs of equation (28) implies that

$$\hat{mc}_{Ht}^N = \hat{w}_{Ht} - p_{Ht}^N - z_{Ht}^N.$$

Taking logs of labor supply—equation (22)—implies that

$$\hat{w}_{Ht} - p_{Ht}^N = \varphi^{-1}\hat{n}_{Ht}.$$

Combining these two equation yields

$$\hat{mc}_{Ht}^N = \varphi^{-1}\hat{n}_{Ht} + \left(p_{Ht} - p_{Ht}^N\right) - z_{Ht}^N. \tag{43}$$

We can substitute this equation into equation (42) to get that

$$\pi_{Ht}^N = \beta E_t \pi_{H,t+1}^N + \kappa \hat{n}_{Ht} - \lambda \hat{p}_{Ht}^N + \nu_{Ht}^N \tag{44}$$

where $\nu_{Ht}^N = -\lambda z_{Ht}^N$ and $\kappa = \lambda\varphi^{-1}$. This is the regional non-tradeable Phillips Curve in our model.

An analogous sequence of steps yields

$$\pi_{Ht}^T = \beta E_t \pi_{H,t+1}^T + \kappa \hat{n}_{Ht} - \lambda \hat{p}_{Ht}^T + \nu_{Ht}^T. \tag{45}$$

This is the regional tradeable Phillips Curve in our model.

A.5 Aggregate Phillips Curve Derivation

Aggregate non-tradeable inflation can be written as $\pi_t^N = \zeta \pi_{Ht}^N + (1-\zeta)\,\pi_{Ft}^N$. Using the Phillips curve for home non-tradeable inflation—equation (44)—and its foreign counterpart, we get that satisfies

$$\pi_t^N = \beta E_t \pi_{t+1}^N + \kappa \hat{n}_t + \nu_t^N - \lambda\left[\zeta \hat{p}_{Ht}^N + (1-\zeta)\,\hat{p}_{Ft}^N\right], \tag{46}$$

Similarly, a weighted average of the Phillips curve for home tradeable inflation—equation (45)—and its foreign counterpart yields

$$\pi_t^T = \beta E_t \pi_{t+1}^T + \kappa \hat{n}_t + \nu_t^T - \lambda\left[\zeta \hat{p}_{Ht}^T + (1-\zeta)\,\hat{p}_{Ft}^T\right]. \tag{47}$$

First order expansions of equation 27 and its foreign counterpart around the zero inflation steady state yield that

$$p_{Ht} = \phi_N p_{Ht}^N + \phi_T \tau_H^H p_{Ht}^T + \phi_T \tau_H^F p_{Ft}^T \tag{48}$$

and

$$p_{Ft} = \phi_N p_{Ft}^N + \phi_T \tau_F^H p_{Ht}^T + \phi_T \tau_F^F p_{Ft}^T. \tag{49}$$

Then the aggregate price level, p_t satisfies $p_t = \zeta p_{Ht} + (1-\zeta) p_{Ft}$. Combining this equation with the previous two equations and using the fact that $\tau_H^H = \tau_F^H = \zeta$ yields

$$p_t = \phi_N \left(\zeta p_{Ht}^N + (1-\zeta) p_{Ft}^N\right) + \phi_T \left(\zeta\zeta p_{Ht}^T + \zeta(1-\zeta) p_{Ft}^T + (1-\zeta)\zeta p_{Ht}^T + (1-\zeta)(1-\zeta) p_{Ft}^T\right)$$

This equation simplifies to

$$p_t = \phi_N p_t^N + \phi_T p_t^T \tag{50}$$

where we use the notation $p_t^N = \zeta p_{Ht}^N + (1-\zeta) p_{Ft}^N$ and $p_t^T = \zeta p_{Ht}^T + (1-\zeta) p_{Ft}^T$.

Equation (50) implies that

$$\pi_t = \phi_N \pi_t^N + \phi_T \pi_t^T.$$

Combining this equation with equations (44) and (45) yields the aggregate Phillips curve

$$\pi_t = \beta E_t \pi_{t+1} + \kappa \hat{n}_t + \nu_t,$$

where $\nu_t \equiv \phi_N \nu_t^N + \phi_T \nu_t^T$ and we make use of the fact that

$$\left[\zeta \hat{p}_{Ht}^N + (1-\zeta)\hat{p}_{Ft}^N\right] + \left[\zeta \hat{p}_{Ht}^T + (1-\zeta)\hat{p}_{Ft}^T\right] = 0.$$

A.6 Deriving the Other Log-Linearized Equations

In this section, we will assume that supply shocks are zero for expositional simplicity. A log-linear approximation of the home consumption Euler equation—equation (23)—yields

$$\hat{c}_{Ht} + \frac{u_{cn}}{u_{cc}} \hat{n}_{Ht} = E_t \left[\hat{c}_{H,t+1} + \frac{u_{cn}}{u_{cc}} \hat{n}_{H,t+1}\right] + \frac{u_c}{u_{cc} C} \left(\hat{r}_t^n - E_t \pi_{H,t+1}\right), \tag{51}$$

where we use the fact that home per capita consumption and labor equal aggregate per capita consumption and labor at the steady state.

Next, we solve for the partial derivatives in the previous equation using the functional form for preferences, equation (8). We have that

$$\begin{aligned}
\frac{u_{cc}C}{u_c} &= \frac{-\sigma^{-1}C\left(C-\chi\frac{N^{1+\varphi^{-1}}}{1+\varphi^{-1}}\right)^{-\sigma^{-1}-1}}{\left(C-\chi\frac{N^{1+\varphi^{-1}}}{1+\varphi^{-1}}\right)^{-\sigma^{-1}}} \\
&= -\sigma^{-1}C\left(C-\chi\frac{N^{1+\varphi^{-1}}}{1+\varphi^{-1}}\right)^{-1} \\
&= -\sigma^{-1}\left(C^{-1}\left(C-\chi\frac{N^{1+\varphi^{-1}}}{1+\varphi^{-1}}\right)\right)^{-1} \\
&= -\sigma^{-1}\left(1-C^{-1}\chi N^{1+\varphi^{-1}}\left(1+\varphi^{-1}\right)^{-1}\right)^{-1} \\
&= -\sigma^{-1}\left(1-\left(\frac{C}{N}\right)^{-1}\chi N^{\varphi^{-1}}\left(1+\varphi^{-1}\right)^{-1}\right)^{-1} \\
&= -\sigma^{-1}\left(1-\left(\frac{C}{N}\right)^{-1}\mu^{-1}\left(1+\varphi^{-1}\right)^{-1}\right)^{-1} \\
&= -\sigma^{-1}\left(1-\mu^{-1}\left(1+\varphi^{-1}\right)^{-1}\right)^{-1},
\end{aligned}$$

where we use the steady state labor supply curve—equation (35)—and the fact that in the steady state $C = N$. Furthermore, we have that

$$\begin{aligned}
u_{cn} &= -\sigma^{-1}\left(C-\chi\frac{N^{1+\varphi^{-1}}}{1+\varphi^{-1}}\right)^{-\sigma^{-1}-1}\times-\frac{\chi}{1+\varphi^{-1}}\left(1+\varphi^{-1}\right)N^{\varphi^{-1}} \\
&= -u_{cc}\chi N^{\varphi^{-1}} \\
&= -u_{cc}\mu^{-1},
\end{aligned}$$

where we again make use of equation (35).

Combining this last to equations with equation (51) yields

$$\hat{c}_{Ht}-\mu^{-1}\hat{n}_{Ht} = E_t\left[\hat{c}_{H,t+1}-\mu^{-1}\hat{n}_{H,t+1}\right]-\sigma_c\left(\hat{r}_t^n-E_t\pi_{H,t+1}\right)$$

where $\sigma_c = \sigma\left(1-\mu^{-1}\left(1+\varphi^{-1}\right)^{-1}\right)$.

Solving this last equation forward yields

$$\begin{aligned}\hat{c}_{Ht} - \mu^{-1}\hat{n}_{Ht} &= -\sigma_c E_t \sum_{j=0}^{\infty} \left(\hat{r}^n_{t+j} - E_t \pi_{H,t+1+j}\right) \\ &= -\sigma_c E_t \sum_{j=0}^{\infty} \hat{r}^n_{t+j} + \sigma_c E_t \sum_{j=0}^{\infty} \pi_{H,t+1+j} \\ &= -\sigma_c E_t \sum_{j=0}^{\infty} \hat{r}^n_{t+j} + \sigma_c E_t \sum_{j=0}^{\infty} \left(p_{H,t+1+j} - p_{H,t+j}\right) \\ &= -\sigma_c E_t \sum_{j=0}^{\infty} \hat{r}^n_{t+j} - \sigma_c p_{Ht}. \end{aligned} \tag{52}$$

Similarly, for foreign households we have

$$\hat{c}_{Ft} - \mu^{-1}\hat{n}_{Ft} = -\sigma_c E_t \sum_{j=0}^{\infty} \hat{r}^n_{t+j} - \sigma_c p_{Ft}. \tag{53}$$

Combining equations (52) and (53) yields

$$\hat{c}_{Ht} - \mu^{-1}\hat{n}_{Ht} = \hat{c}_{Ft} - \mu^{-1}\hat{n}_{Ft} + \sigma_c \left(p_{Ft} - p_{Ht}\right),$$

which is the Backus-Smith condition for our model.

Define $\hat{\xi}_{Ht} = \log \tau_{Ht} - \log \tau_H$ and $\hat{\xi}_{Ft} = \log \tau_{Ft} - \log \tau_F$. With this notation, log-linear approximations of (24) and (25) as well as their foreign counterparts yields

$$\hat{c}^N_{Ht} = \hat{c}_{Ht} - \eta \left(p^N_{Ht} - p_{Ht}\right) \tag{54}$$

$$\hat{c}^T_{Ht} = \hat{\xi}_{Ht} + \hat{c}_{Ht} - \eta \left(p^T_{Ht} - p^T_{Ht}\right) \tag{55}$$

$$\hat{c}^T_{Ft} = \hat{\xi}_{Ft} + \hat{c}_{Ht} - \eta \left(p^T_{Ft} - p^T_{Ht}\right) \tag{56}$$

$$\hat{c}^N_{Ft} = \hat{c}_{Ft} - \eta \left(p^N_{Ft} - p_{Ft}\right) \tag{57}$$

$$\hat{c}^{TH}_{Ft} = \hat{c}_{Ft} - \eta \left(p^T_{Ht} - p_{Ft}\right) \tag{58}$$

$$\hat{c}^{TF}_{Ft} = \hat{c}_{Ft} - \eta \left(p^T_{Ft} - p_{Ft}\right). \tag{59}$$

Note that the expenditure share on tradeable and non-tradeable goods always sums to 1: $\tau_{Ht} +$

$\tau_{Ft} = 1$. This implies that

$$\hat{\xi}_{Ht}\tau_H^H + \hat{\xi}_{Ft}\tau_H^F = 0$$

which in turn implies that

$$\hat{\xi}_{Ft} = -\frac{\zeta}{1-\zeta}\hat{\xi}_{Ht}.$$

First differencing equations (48) and (49) implies

$$\pi_{Ht} = \phi_N \pi_{Ht}^N + \phi_T \tau_H^H \pi_{Ht}^T + \phi_T \tau_H^F \pi_{Ft}^T$$

and

$$\pi_{Ft} = \phi_N \pi_{Ft}^N + \phi_T \tau_F^H \pi_{Ht}^T + \phi_T \tau_H^F \pi_{Ft}^T.$$

Note that without supply shocks, output and employment are equal. This implies that

$$Y_{Ht}^N = N_{Ht}^N, \quad Y_{Ht}^T = N_{Ht}^T, \quad Y_{Ft}^N = N_{Ft}^N, \quad Y_{Ft}^T = N_{Ft}^T.$$

We furthermore have that

$$\zeta N_{Ht} = N_{Ht}^N + N_{Ht}^T.$$

This equation says that total labor supplied by households in the home region equals total labor demanded by firms. The ζ on the left-hand-side reflects the fact that N_{Ht} is per capita labor supply. A log-linear approximation of this last expression around the symmetric steady state yields

$$\begin{aligned}\hat{n}_{Ht} &= \frac{N^N}{N^N + N^T}\hat{n}_{Ht}^N + \frac{N^T}{N^N + N^T}\hat{n}_{Ht}^T \\ &= \phi_N \hat{n}_{Ht}^N + \phi_T \hat{n}_{Ht}^T\end{aligned}$$

Similarly, in the foreign region we have that

$$\hat{n}_{Ft} = \phi_N \hat{n}_{Ft}^N + \phi_T \hat{n}_{Ft}^T.$$

Aggregate employment is

$$N_t = \zeta N_{Ht} + (1-\zeta) N_{Ft}.$$

Log-linearizing this equation around the symmetric steady state yields

$$\begin{aligned}\hat{n}_t &= \frac{\zeta N}{\zeta N + (1-\zeta) N}\hat{n}_{Ht} + \frac{(1-\zeta) N}{\zeta N + (1-\zeta) N}\hat{n}_{Ft} \\ &= \zeta \hat{n}_{Ht} + (1-\zeta)\hat{n}_{Ft},\end{aligned}$$

where N is steady state household labor supply, equal across the two regions at the symmetric steady state.

Market clearing conditions in the non-tradeable sector implies that $N_{Ht}^N = \zeta C_{Ht}^N$. A log-linear approximation of this expression yields that

$$\hat{n}_{Ht}^N = \hat{c}_{Ht}^N.$$

Using equation (54), we get that

$$\hat{n}_{Ht}^N = \hat{c}_{Ht} - \eta\left(\hat{p}_{Ht}^N - \hat{p}_{Ht}\right)$$

A similar set of steps for the foreign region yields

$$\hat{n}_{Ft}^N = \hat{c}_{Ft} - \eta\left(p_{Ft}^N - p_{Ft}\right).$$

In the tradeable sector market clearing implies

$$\begin{aligned}N_{Ht}^T &= \zeta C_{Ht}^{TH} + (1-\zeta) C_{Ft}^{TH} \\ &= \zeta \phi_T \tau_{Ht}^H C_{Ht}\left(\frac{P_{Ht}^T}{P_{Ht}}\right)^{-\eta} + (1-\zeta)\phi_T \tau_{Ft}^H C_{Ft}\left(\frac{P_{Ft}^T}{P_{Ft}}\right)^{-\eta}\end{aligned}$$

where the second line follows from equations (24) and (25). Log-linearizing around the symmetric steady state implies

$$\hat{n}_{Ht}^T = \zeta\left[\hat{c}_{Ht} - \eta\left(p_{Ht}^T - p_{Ht}\right) + \hat{\xi}_{Ht}\right] + (1-\zeta)\left[\hat{c}_{Ft} - \eta\left(p_{Ht}^T - p_{Ft}\right)\right].$$

Similarly, in the foreign region we have

$$\hat{n}_{Ft}^T = \zeta\left[\hat{c}_{Ht} - \eta\left(p_{Ft}^T - p_{Ht}\right) + \hat{\xi}_{Ft}\right] + (1-\zeta)\left[\hat{c}_{Ft} - \eta\left(p_{Ft}^T - p_{Ft}\right)\right].$$

Finally, we define deviations of unemployment from the steady state as

$$\hat{n}_{Ht} = \log N_{Ht} - \log N_H \approx (N_{Ht} - 1) - (N_H - 1) = -(u_{Ht} - u_H) = -\hat{u}_{Ht}.$$

A.7 Log-Linearized Equations of the Model

For convenience we repeat the full set of log-linearized equilibrium conditions of the model.

- Parameters:
  - $\sigma_c = \sigma \left(1 - \mu^{-1} \left(1 + \varphi^{-1}\right)^{-1}\right)$
  - $\kappa = \lambda \varphi^{-1}$
  - $\lambda = (1 - \alpha)(1 - \alpha\beta)/\alpha$
  - $\mu = \theta/(\theta - 1)$
  - $\tau_H^H = \tau_F^H = \zeta$
- The law of motion for tradeable demand is

$$\hat{\xi}_{Ht} = \rho_\xi \hat{\xi}_{Ht} + \varepsilon_t$$

  and

$$\hat{\xi}_{Ft} = -\frac{\zeta}{1 - \zeta} \hat{\xi}_{Ht}.$$

- The home non-tradeable Phillips Curve is:

$$\pi_{Ht}^N = \beta E_t \pi_{H,t+1}^N - \kappa \hat{u}_{Ht} - \lambda \hat{p}_{Ht}^N + \nu_{Ht}^N$$

- The home tradeable Phillips Curve is:

$$\pi_{Ht}^T = \beta E_t \pi_{H,t+1}^T - \kappa \hat{u}_{Ht} - \lambda \hat{p}_{Ht}^T + \nu_{Ht}^T.$$

- The home Euler equation is:

$$\hat{c}_{Ht} - \mu^{-1} \hat{n}_{Ht} = E_t \left[\hat{c}_{H,t+1} - \mu^{-1} \hat{n}_{H,t+1}\right] - \sigma_c \left(\hat{r}_t^n - E_t \pi_{H,t+1}\right)$$

- The Backus-Smith condition is:

$$\hat{c}_{Ht} - \mu^{-1}\hat{n}_{Ht} = \hat{c}_{Ft} - \mu^{-1}\hat{n}_{Ft} + \sigma_c \left(p_{Ft} - p_{Ht}\right)$$

- The foreign non-tradeable Phillips Curve is:

$$\pi^N_{Ft} = \beta E_t \pi^N_{F,t+1} + \kappa \hat{n}_{Ft} - \lambda \hat{p}^N_{Ft} + \nu^N_{Ft}$$

- The foreign tradeable Phillips Curve is:

$$\pi^T_{Ft} = \beta E_t \pi^T_{F,t+1} + \kappa \hat{n}_{Ft} - \lambda \hat{p}^T_{Ft} + \nu^T_{Ft}$$

- Definitions of inflation:

$$\pi_{Ht} = p_{Ht} - p_{H,t-1}$$

$$\pi_{Ft} = p_{Ft} - p_{F,t-1}$$

$$\pi^N_{Ht} = p^N_{Ht} - p^N_{H,t-1}$$

$$\pi^T_{Ht} = p^T_{Ht} - p^T_{H,t-1}$$

$$\pi^N_{Ft} = p^N_{Ft} - p^N_{F,t-1}$$

$$\pi^T_{Ft} = p^T_{Ft} - p^T_{F,t-1}$$

$$\pi_{Ht} = \phi_N \pi^N_{Ht} + \phi_T \tau^H_H \pi^T_{Ht} + \phi_T \tau^F_H \pi^T_{Ft}$$

$$\pi_{Ft} = \phi_N \pi^N_{Ft} + \phi_T \tau^H_F \pi^T_{Ht} + \phi_F \tau^F_F \pi^T_{Ft}$$

- The home resource constraint in the non-tradeable sector is:

$$\hat{n}^N_{Ht} = \hat{c}_{Ht} - \eta \left(p^N_{Ht} - p_{Ht}\right)$$

- The foreign resource constraint in the non-tradeable sector is:

$$\hat{n}^N_{Ft} = \hat{c}_{Ft} - \eta \left(p^N_{Ft} - p_{Ft}\right)$$

- The home resource constraint in the tradeable sector is:

$$\hat{n}_{Ht}^{T} = \zeta \left[\hat{c}_{Ht} - \eta \left(p_{Ht}^{T} - p_{Ht} \right) + \hat{\xi}_{Ht} \right] + (1 - \zeta) \left[\hat{c}_{Ft} - \eta \left(p_{Ht}^{T} - p_{Ft} \right) \right]$$

- The foreign resource constraint in the tradeable sector is:

$$\hat{n}_{Ft}^{T} = \zeta \left[\hat{c}_{Ht} - \eta \left(p_{Ft}^{T} - p_{Ht} \right) + \hat{\xi}_{Ft} \right] + (1 - \zeta) \left[\hat{c}_{Ft} - \eta \left(p_{Ft}^{T} - p_{Ft} \right) \right]$$

- Aggregate labor in the home region then satisfies the log-linear equations

$$\hat{n}_{Ht} = \phi_N \hat{n}_{Ht}^{N} + \phi_T \hat{n}_{Ht}^{T}$$

- Aggregate labor in the foreign region satisfies

$$\hat{n}_{Ft} = \phi_N \hat{n}_{Ft}^{N} + \phi_T \hat{n}_{Ft}^{T}$$

- Monetary policy is

$$\hat{r}_t^n = \varphi_\pi \left(\pi_t - \bar{\pi}_t \right) + \varphi_n \left(\hat{n}_t - \bar{n}_t \right) + \varepsilon_{rt}$$

- Aggregate employment satisfies

$$\hat{n}_t = \zeta \hat{n}_{Ht} + (1 - \zeta) \hat{n}_{Ft}$$

- Aggregate inflation satisfies

$$\hat{\pi}_t = \zeta \hat{\pi}_{Ht} + (1 - \zeta) \hat{\pi}_{Ft}$$

- The deviation of unemployment from its steady state value is

$$\hat{u}_t = -\hat{n}_t.$$

A.8 The Importance of Non-Tradeable Inflation

Here, we show that the slope of the regional Phillips Curve for overall regional consumer price inflation is smaller than the slope of the aggregate Phillips Curve, by a factor equal to the expen-

diture share on non-tradeable goods. For simplicity, we present this derivation with all supply shocks ν_t set to zero.

Consider the Phillips curves for home non-tradeables, home tradeables, and foreign tradeables:

$$\pi^N_{Ht} = \beta E_t \pi^N_{H,t+1} - \kappa \hat{u}_{Ht} - \lambda \hat{p}^N_{Ht}$$

$$\pi^T_{Ht} = \beta E_t \pi^T_{H,t+1} - \kappa \hat{u}_{Ht} - \lambda \hat{p}^T_{Ht}$$

$$\pi^T_{Ft} = \beta E_t \pi^T_{F,t+1} - \kappa \hat{u}_{Ft} - \lambda \hat{p}^T_{Ft}.$$

Substituting these three equations into the definition for home consumer price inflation

$$\pi_{Ht} = \phi_N \pi^N_{Ht} + \phi_T \tau^H_H \pi^T_{Ht} + \phi_T \tau^F_H \pi^T_{Ft}$$

yields

$$\pi_{Ht} = \beta E_t \pi_{H,t+1} - \left(\phi_N + \phi_T \tau^H_H\right) \kappa \hat{u}_{Ht} - \lambda \left(\phi_N \hat{p}^N_{Ht} + \phi_T \tau^H_H \hat{p}^T_{Ht}\right) - \phi_T \tau^F_H \kappa \hat{u}_{Ft} - \lambda \phi_T \tau^F_H \hat{p}^T_{Ft}.$$

An analogous derivation yields the following Phillips curve for foreign consumer prices

$$\pi_{Ft} = \beta E_t \pi_{F,t+1} - \left(\phi_N + \phi_T \tau^F_F\right) \kappa \hat{u}_{Ft} - \lambda \left(\phi_N \hat{p}^N_{Ft} + \phi_T \tau^F_F \hat{p}^T_{Ft}\right) - \phi_T \tau^H_F \kappa \hat{u}_{Ht} - \lambda \phi_T \tau^H_F \hat{p}^T_{Ht}.$$

Subtracting the second of these last two equations from the first (and using the fact that $\tau^H_H = \tau^H_F = \zeta$) yields

$$\pi_{Ht} - \pi_{Ft} = \beta \left(E_t \pi_{H,t+1} - E_t \pi_{F,t+1}\right) - \phi_N \kappa \left(\hat{u}_{Ht} - \hat{u}_{Ft}\right) - \phi_N \lambda \left(\hat{p}^N_{Ht} - \hat{p}^N_{Ft}\right). \tag{60}$$

The coefficient in a regional panel regression corresponds to the coefficient in a differenced equation like this one. Notice that the coefficient on unemployment is $\phi_N \kappa$ rather than κ. In other words, the coefficient differs from the coefficient in the aggregate Phillips curve by the factor ϕ_N.

A.9 The Role of GHH Preferences

The key feature of GHH preferences that we exploit is that, with GHH preferences, there are no wealth effects on labor supply either at the aggregate or the regional level. In contrast, with separable preferences, wealth effects on labor supply are an important determinant of marginal

cost and therefore influence the Phillips curve.

To see this more clearly, consider the non-tradeable regional Phillips curve under separable preferences:

$$\pi^N_{Ht} = \beta E_t \pi^N_{H,t+1} - \kappa \hat{u}_{Ht} + \lambda \sigma^{-1} \hat{c}_{Ht} - \lambda \hat{p}^N_{Ht} + \nu^N_{Ht}, \tag{61}$$

and the aggregate Phillips Curve under separable preferences:

$$\pi_t = \beta E_t \pi_{t+1} - \kappa \hat{u}_t + \lambda \sigma^{-1} \hat{c}_t + \nu_t. \tag{62}$$

Relative to the GHH case, both the non-tradeable regional Phillips curve and aggregate Phillips curve include a consumption term. These terms appear because of wealth effects on labor supply affect marginal cost in this model. These wealth effects complicate the comparison between the regional and aggregate Phillips curve because the relationship between employment and consumption is different at the aggregate level than at the regional level. At the aggregate level, $\hat{c}_t = \hat{n}_t + z_t$. This implies that we can replace the $\hat{c}_t$ term with $\hat{n}_t + z_t$ in equation (62) and get a consolidated coefficient of $\kappa + \lambda\sigma^{-1}$ on unemployment. At the regional level, however, this is not possible because risk-sharing across regions implies that $\hat{c}_{Ht} \neq \hat{n}_{Ht} + z_{Ht}$. This difference implies that the slope of the non-tradeable regional Phillips curve will differ from the slope of the aggregate Phillips curve when preferences are separable.

A.10 Relaxing Full Information Rational Expectations

To derive the solved forward Phillips Curve, equation (13) in the main text, we manipulated the Phillips curve under the standard assumption of full-information rational expectations. However, this type of derivation actually only relies on the weaker assumption that the law of iterated expectations holds. Let's consider the aggregate Phillips curve—equation (12)—for simplicity. Under the assumption that the law of iterated expectations holds and the additional simplifying assumption that the unemployment rate follows an AR(1) process, we can solve this equation forward to get that

$$\pi_t = -\frac{\kappa}{1 - \rho^F_{\tilde{u}} \beta} \tilde{u}_t + F_t \pi_{t+\infty} + \tilde{\omega}_t, \tag{63}$$

where F_t denotes agents' expectations conditional on information at time t, $F_t \pi_{t+\infty}$ is the agent's subjective forecast about the inflation target, $\tilde{\omega}_t \equiv F_t \sum_{j=0}^{\infty} \beta^j \nu_{t+j}$, and $\rho^F_{\tilde{u}}$ is agents' subjective belief about the autoregressive coefficient governing the persistence of fluctuations in unem-

ployment. Notice that if $\rho_{\tilde{u}}^F < \rho_{\tilde{u}}$, the Phillips Curve is less forward looking than the rational expectations Phillips curve. Rational expectations is the special case where $\rho_{\tilde{u}}^F = \rho_{\tilde{u}}$ and $F_t\pi_{t+\infty} = E_t\pi_{t+\infty}$. Coibion and Gorodnichenko (2012, 2015a) provide evidence consitent with the law of iterated expectations holding but full information rational expectations not holding. See Adam and Padula (2011), Coibion and Gorodnichenko (2015a), Coibion, Gorodnichenko, and Kamdar (2018) for further discussion of these issues.

A.11 Time Aggregation

Here, we show how time aggregation associated with using four-quarter inflation as our dependent variable implies that we should divide our estimate of κ by 4 since our model in section 3 is written in terms of quarterly inflation. Consider the non-tradeable regional Phillips Curve—equation (13) from the main text:

$$\pi_{Ht}^N = -E_t \sum_{j=0}^{\infty} \beta^j \left(\kappa \tilde{u}_{H,t+j} + \lambda \hat{p}_{H,t+j}^N\right) + E_t\pi_{t+\infty}, \tag{64}$$

where for simplicity we have set the supply shock ω_{Ht}^N equal to zero. We can rewrite this last equation as

$$p_{Ht}^N - p_{H,t-1}^N = -\kappa PV_{Ht}^u - \lambda PV_{Ht}^p + E_t\pi_{t+\infty}$$

where

$$PV_{Ht}^u = E_t \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j}$$

$$PV_{Ht}^p = E_t \sum_{j=0}^{\infty} \beta^j \hat{p}_{H,t+j}^N.$$

This same equation hold for periods t, $t-1$, $t-2$, and $t-3$:

$$\begin{aligned}
p_{Ht}^N - p_{H,t-1}^N &= -\kappa PV_{Ht}^u - \lambda PV_{Ht}^p + E_t\pi_{t+\infty} \\
p_{H,t-1}^N - p_{H,t-2}^N &= -\kappa PV_{H,t-1}^u - \lambda PV_{H,t-1}^p + E_{t-1}\pi_{t+\infty} \\
p_{H,t-2}^N - p_{H,t-3}^N &= -\kappa PV_{H,t-2}^u - \lambda PV_{H,t-2}^p + E_{t-2}\pi_{t+\infty} \\
p_{H,t-3}^N - p_{H,t-4}^N &= -\kappa PV_{H,t-3}^u - \lambda PV_{H,t-3}^p + E_{t-3}\pi_{t+\infty}
\end{aligned}$$

Summing the preceding four equations together yields

$$\begin{aligned} p_{Ht}^N - p_{H,t-4}^N = &-\kappa \left(PV_{Ht}^u + PV_{H,t-1}^u + PV_{H,t-2}^u + PV_{H,t-3}^u \right) \\ &- \lambda \left(PV_{Ht}^p + PV_{H,t-1}^p + PV_{H,t-2}^p + PV_{H,t-3}^p \right) \\ &+ E_t \pi_{t+\infty} + E_{t-1} \pi_{t+\infty} + E_{t-2} \pi_{t+\infty} + E_{t-3} \pi_{t+\infty}. \end{aligned}$$

Taking expectations at time $t-4$ then yields

$$\begin{aligned} E_{t-4} p_{Ht}^N - p_{H,t-4}^N = &-\kappa \left(E_{t-4} PV_{Ht}^u + E_{t-4} PV_{H,t-1}^u + E_{t-4} PV_{H,t-2}^u + E_{t-4} PV_{H,t-3}^u \right) \\ &- \lambda \left(E_{t-4} PV_{Ht}^p + E_{t-4} PV_{H,t-1}^p + E_{t-4} PV_{H,t-2}^p + E_{t-4} PV_{H,t-3}^p \right) \\ &+ 4 E_{t-4} \pi_{t+\infty}. \end{aligned}$$

Adding and subtracting p_{Ht}^N yields

$$\begin{aligned} p_{Ht}^N - p_{H,t-4}^N = &-\kappa \left(E_{t-4} PV_{Ht}^u + E_{t-4} PV_{H,t-1}^u + E_{t-4} PV_{H,t-2}^u + E_{t-4} PV_{H,t-3}^u \right) \\ &- \lambda \left(E_{t-4} PV_{Ht}^p + E_{t-4} PV_{H,t-1}^p + E_{t-4} PV_{H,t-2}^p + E_{t-4} PV_{H,t-3}^p \right) \\ &+ 4 E_{t-4} \pi_{t+\infty} - \left(E_{t-4} p_{Ht}^N - p_{Ht}^N \right). \end{aligned}$$

We now assume that PV_{Ht}^u and PV_{Ht}^p are well approximated by univariate driftless random walks. We present empirical evidence supporting this assumption in section A.11.1 below. Given this

assumption, the preceding equation simplifies to

$$\begin{aligned}
p^N_{Ht} - p^N_{H,t-4} &= -4\kappa E_{t-4} PV^u_{Ht} \\
&\quad - 4\lambda E_{t-4} PV^p_{Ht} \\
&\quad + 4E_{t-4}\pi_{t+\infty} - \left(E_{t-4} p^N_{Ht} - p^N_{Ht}\right) \\
&= -4\kappa \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} - 4\lambda \sum_{j=0}^{\infty} \beta^j \hat{p}^N_{H,t+j} + 4E_{t-4}\pi_{t+\infty} \\
&\quad - 4\kappa \left(E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} - \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} \right) \\
&\quad - 4\lambda \left(E_{t-4} \sum_{j=0}^{\infty} \beta^j \hat{p}^N_{H,t+j} - \sum_{j=0}^{\infty} \beta^j \hat{p}^N_{H,t+j} \right) \\
&\quad - \left(E_{t-4} p^N_{Ht} - p^N_{Ht}\right) \\
&= -4\kappa \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} - 4\lambda \sum_{j=0}^{\infty} \beta^j \hat{p}^N_{H,t+j} + 4E_{t-4}\pi_{t+\infty} + \upsilon_{Ht}
\end{aligned}$$

where υ_{Ht} is a rational expectations error uncorrelated with variables at time $t-4$. This last equation shows that estimating equation (17) with the dependent variable defined as $\pi^N_{it} = p^N_{it} - p^N_{i,t-4}$ yields estimates of 4κ and 4λ provided that the forward sums in equation (17) are well approximated by a random walk.

A.11.1 The Dynamics of the Present Value of Unemployment and Relative Prices

The derivation above relied on the simplifying assumption that $\sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j}$ and $\sum_{j=0}^{\infty} \beta^j \hat{p}^N_{H,t+j}$ follow univariate random walks. We can assess the accuracy of this assumption by running the regressions

$$\sum_{j=0}^{T} \beta^j u_{i,t+j} = \alpha_i + \gamma_t + \rho_{uu} \sum_{j=0}^{T} \beta^j u_{i,t+j-1} + \rho_{up} \sum_{j=0}^{T} \beta^j p^N_{i,t+j-1}$$

$$\sum_{j=0}^{T} \beta^j p^N_{i,t+j} = \alpha_i + \gamma_t + \rho_{pu} \sum_{j=0}^{T} \beta^j u_{i,t+j-1} + \rho_{pp} \sum_{j=0}^{T} \beta^j p^N_{i,t+j-1}.$$

As in the main text, we truncate the infinite sums at $T = 20$. Table A.1 presents results for these regression. In both regressions, the coefficient on the lag of the dependent variable is very close to 1, whereas the coefficient on the lag of the other future sum is near zero. This shows that both the present value of both unemployment and the present value of relative prices are well

Table A.1: Random Walks in the Present Values of Unemployment and Relative Prices

	PV of Unemployment		PV of Relative Prices	
	(1)	(2)	(3)	(4)
Lag PV of Unemployment	0.999	0.997	0.020	0.007
	(0.001)	(0.002)	(0.005)	(0.008)
Lag PV of Rel. Prices	-0.002	-0.001	0.994	0.992
	(0.001)	(0.002)	(0.001)	(0.005)
State Effects	✓	✓	✓	✓
Time Effects		✓		✓

Note: In the first two columns, we regress $\sum_{j=0}^{T} \beta^j \tilde{u}_{i,t+j}$ on its quarterly lag and the quarterly lag of $\sum_{j=0}^{T} \beta^j p^N_{i,t+j}$, where $\tilde{u}_{it}$ is unemployment in state i in quarter t and p^N_{it} is relative non-tradeable prices. In the last two columns we repeat the exercise with $\sum_{j=0}^{T} \beta^j p^N_{i,t+j}$ as the outcome. The sample period is 1978-2018. We set T = 20. Unemployment is in percentage points and relative prices are in 100 x log points. The regression is unweighted. Standard errors are in parentheses. These are two-way clustered by date and state. The number of observations is 3495.

approximated by univariate random walk processes at quarterly frequency.

A.12 Applying Our Estimation Procedure to Model Generated Data

Our empirical specification—equation (17)—approximates the structural regional Phillips curve in the model we present in section 3—equation (14). The approximation is that we truncate the present sums of unemployment and relative prices. We can assess the error associated with this approximation by estimating κ with equation (17) using data generated by our model. This will also more generally verify that our empirical method is able to identify κ when applied to data generated from our model.

To simulate data from the model, we adopt a quarterly calibration. We simulate the model for a wide range of values for the slope of the Phillips curve κ. We vary κ across these simulations by varying the frequency of price change α. There are 12 remaining parameters in the model for which we must select values. The values that we choose for these parameters are listed Table A.2. We first calibrate 9 parameters to standard values, using external sources. We set the quarterly discount factor to $\beta = 0.99$ consistent with an annual riskless real interest rate of 4 percent. We set the elasticity of intertemporal substitution equal to $\sigma = 1$. We set the Frisch elasticity of labor supply to $\varphi = 1$ roughly in line with the mix of micro- and macro-economic evidence in Chetty et al. (2011). We follow Itskhoki and Mukhin (2017) and Feenstra et al. (2018) in setting the elasticity of substitution between tradeables and non-tradeables to $\eta = 1.5$. We set the elasticity of substitution across varieties to $\theta = 4$. We set the size of the home region to $\zeta = 0.05$. This results in a home

Table A.2: Calibrated Parameters

Externally Calibrated Parameters		
Intertemporal Elasticity of Substitution	σ	1
Discount factor	β	0.99
Elasticity of substitution between tradeables and non-tradeables	η	1.5
Elasticity of substitution between varieties	θ	4
Size of home region	p	0.05
Steady state consumption share of non-tradeables	ϕ_H^N	0.66
Taylor Rule coefficient on inflation	φ_π	1.5
Taylor Rule coefficient on unemployment	φ_u	1.5
Frisch elasticity of labor supply	φ	1
Parameters Calibrated From the Data		
Persistence of tradeable demand	ρ_ξ	0.9
Parameters Targeting Moments from Data		
Standard deviation of tradeable demand innovation	σ_ξ	2.1
Standard deviation of supply shock	σ_ν	2.1

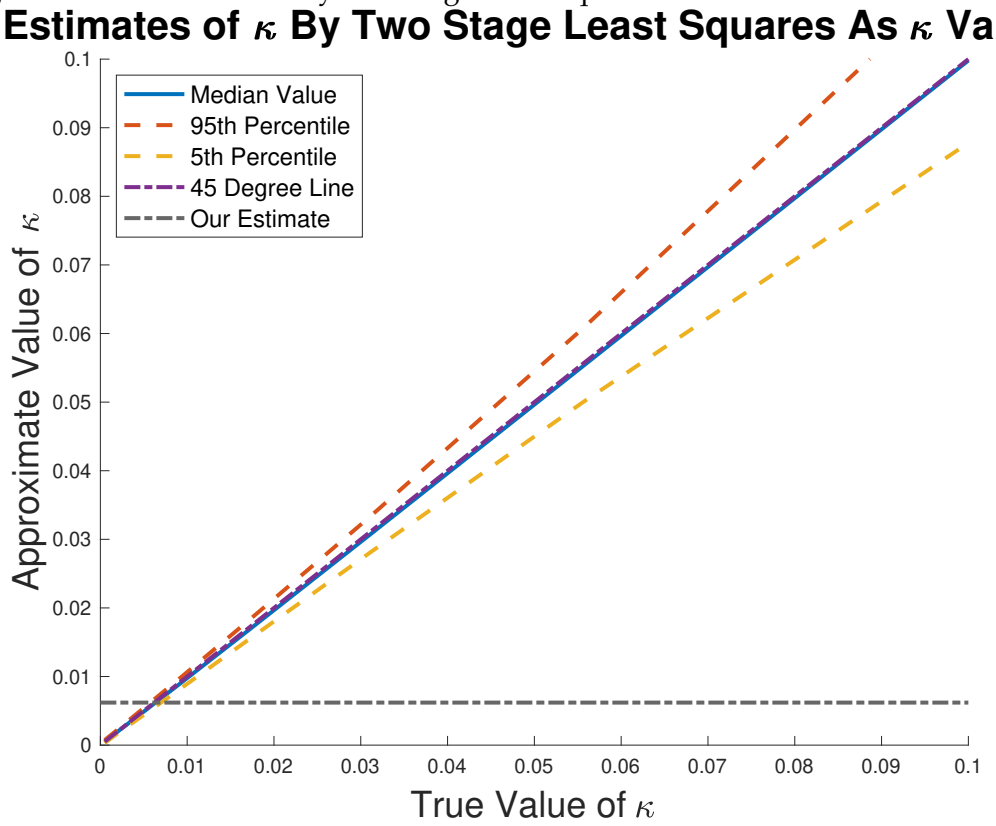
Targeted Moments	Model	Data
Standard deviation of annual unemployment	2.1	2.1

regions that is in between the size of New York and Pennsylvania. We set the steady state share of consumption of non-tradeables to $\phi_N = 0.66$ following Nakamura and Steinsson (2014). We set the parameters of the central bank's interest rate rule to $\varphi_\pi = 1.5$ and $\varphi_u = 1.5$.

We then calibrate the first order autocorrelation of tradeable demand to 0.9. We obtain this value by regressing tradeable demand on its 16th lag, and taking the 16th root of the regression coefficient. We assume that the supply shock is i.i.d. We choose the value of the remaining two parameters, the standard deviations of the innovations to the supply and tradeable demand shocks, to jointly match two criteria. First, we require that the standard deviation of annual unemployment generated from the model equals its value in the data conditional on time and state fixed effects. Second, we require that the contribution of supply and demand shocks to the variance of unemployment is equal. To calculate the standard deviation of unemployment, we simulate data from the model at a quarterly frequency, setting $\kappa = 0.0062$ (the estimate of κ in Column (4) of Table 1). We time-aggregate the simulated data from the model, in order calculate the standard deviation of unemployment in simulated data in the same way as in real world data.

We now use the model to show that our GMM procedure consistently estimates κ. We simulate the model for a range of values for κ between 0.0005 and 0.1. For each value of κ, we simulate the model 1,000 times. For each of these simulation, we generate data for 8,000 periods, roughly

Figure A.1: Estimates of κ By Two Stage Least Squares on Simulated Data from our Model



the size of our quarterly dataset. We then estimate κ for each simulation using our two stage least squares procedure.

Figure A.1 plots the median estimated κ as a function of the true value of κ along with the 5th and 95th quantile of the distribution of the estimated κ's as function of the true value of κ. For visual aid, we also plot the 45 degree line and a horizontal line at our estimated value of κ.

Figure A.1 shows that our two stage least squares procedure consistently estimates κ. The median estimate of κ always lies very close to the 45 degree line. This implies that inconsistency due to truncating the present values of unemployment and relative prices is not important. Also, our procedure is quite precise. In the region of the value of κ that we have estimated from the data, the 5th and 95th percentiles are close to one another and to the median estimate of κ.

Table B.1: Slope of the Aggregate Phillips Curve

	Pre-1990	Post-1990
	(1)	(2)
Core CPI	0.796	0.120
	(0.120)	(0.026)
Median CPI	0.386	0.247
	(0.136)	(0.032)
Shelter CPI	1.624	0.397
	(0.350)	(0.048)
PCE	0.416	0.040
	(0.078)	(0.019)
Core less Shelter CPI	0.221	-0.069
	(0.103)	(0.026)
Core CPI RS	0.182	0.149
	(0.108)	(0.027)

B Data Appendix

B.1 Sensitivity of the Phillips Curve Slope using Aggregate Data

Table B.1 presents estimates of the slope of the Phillips curve using aggregate data for several different measures of inflation. We present estimates separately for the period 1978-1990 and 1991-2018. In each case, we run the regression

$$\pi_t - E_t\pi_{t+\infty} = \alpha + \psi\tilde{u}_{t-4} + \epsilon_t, \tag{65}$$

with 10-year ahead inflation expectations from the Survey of Professional Forecasters serving as a proxy for $E_t\pi_{t+\infty}$, and the 4 quarter moving average of the CBO unemployment gap serving as a proxy for $\tilde{u}_t$. We present results for six measures of inflation: the Core CPI, the Median CPI produced by the Federal Reserve Bank of Cleveland, the CPI for shelter, the PCE, the Core CPI less shelter, and the Core CPI research series. These regressions are run on quarterly data.

The results in B.1 show that the slope of the Phillips curve estimated using aggregate data is highly sensitive to seemingly minor changes in the inflation measure used. This is particularly the case in the pre-1990 sample where the slope estimates vary by roughly a factor of 10 from 0.182 to 1.624. The estimates for the post-1990 sample also vary a great deal, but somewhat less than the

pre-1990 estimates.

Table B.1 also illustrates that inference about the degree to which the Phillips curve flattens based on aggregate data is highly sensitive to the inflation measure used. For some measures, the Phillips curve flattens a great deal (e.g., Core CPI and CPI for shelter). But for others it does not flatten much at all (e.g., median CPI and Core CPI research series).

There has been extensive discussion in the literature behind this. Stock and Watson (2019) discuss how certain sub-indices — such as shelter — are more cyclical than others. Ball and Mazumder (2019) argue that the Median CPI has advantages arising from the elimination of large fluctuations in certain components of the CPI. The difference between CPI inflation and PCE inflation arises to a significant degree from differences in the treatment of housing services in the early 1980s and the fact that the BLS does not revise the CPI, while the PCE is revised.

B.2 CPI Inflation Using Pre- and Post-1983 Housing Methodology

The BLS made a significant change to the methods used to calculate inflation for owner-occupied housing in 1983. This was important given the sizable weight of owner-occupied housing in the CPI (22.8%). Before 1983, the component of the CPI having to do with owner-occupied housing was constructed from a weighted average of changes in house prices and mortgage costs (i.e., interest rates). More specifically, it was made up of home purchases (9.9 percentage points); mortgage interest cost (6.5 percentage points), other financing, taxes and insurance (2.7 percentage points); and maintenance and repairs (3.7 percentage points). For further discussion, see Bureau of Labor Statistics (1982) and Poole, Ptacek, and Verbrugge (2005).

In 1983, the BLS shifted to using changes in rents as a proxy for inflation of owner occupied housing. Figure B.1 plots CPI inflation from 1972 to 2018 (gray line). It also plots our attempt at estimating what CPI inflation would have been had the BLS not changed the methodology for calculating the shelter component in 1983 (black line). Evidently, the pre-1983 methodology yields a much more variable (and cyclical) measure of inflation over the last few decades. The difference between the gray line and the black line in Figure B.1 prior to 1983 gives a sense for how accurately we can replicate the BLS's pre-1983 methodology.

B.3 Price Index Construction

Here we discuss several details of our procedure for constructing state-level price indexes.

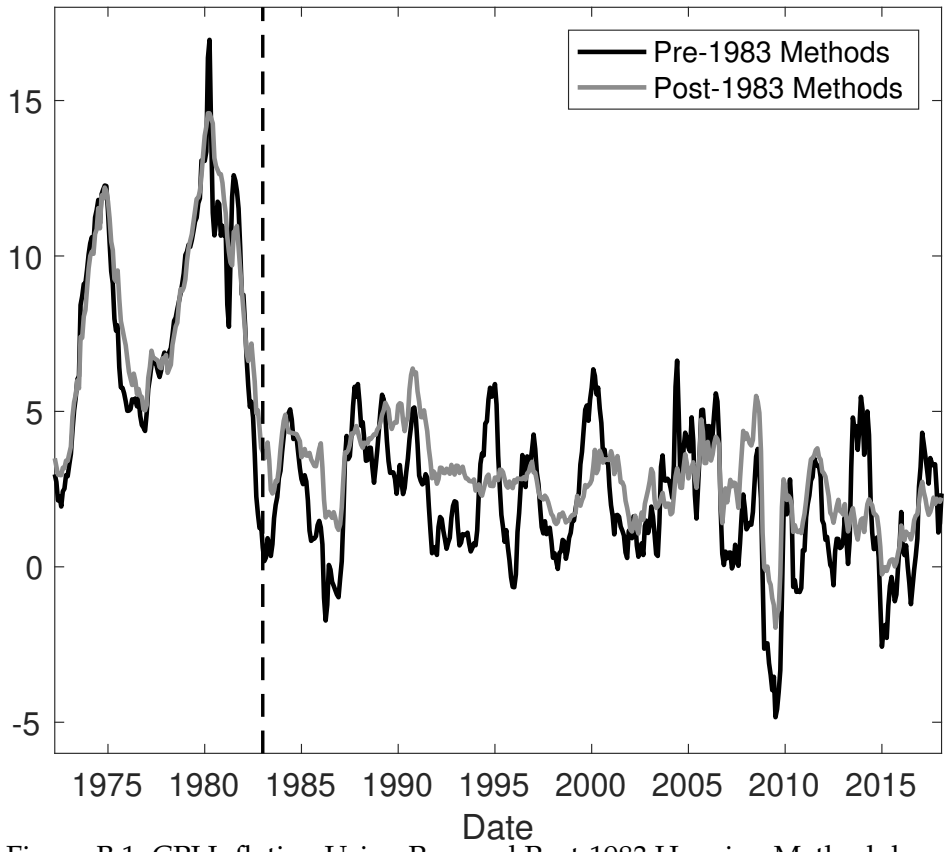


Figure B.1: CPI Inflation Using Pre- and Post-1983 Housing Methodology

Note: This figure plots overall CPI inflation in the US (gray line) and our attempt at estimating what CPI inflation would have been had the BLS not changed the methodology for calculating the shelter component in 1983 (black line). We present these results for the sample period 1972 to 2018. The difference between the gray and the black line before 1983 gives a sense for how accurately we can replicate the BLS's pre-1983 methodology.

B.3.1 Sample Restrictions

We restrict the sample we use in several ways. First, we exclude from our sample price relatives involving a product replacement when the size of the new product is unobserved. This reduces sampling error in our price indexes. Second, we Winsorize price relatives that are larger than 10 or smaller than 0.1. Third, we drop quote lines that include collected prices that are smaller than a tenth of a cent. A quote line includes all versions of a particular "quote-outlet" pair. Recall that a "quote-outlet" pair represents a specific product in a specific location, such as a 2L bottle of Diet Coke from the Westside Market at 110th Street in New York City.

Fourth, we drop observations associated with clearance sales at the end of a quote line. Intuitively, if products systematically go on sale, and then disappear from the data, this can lead to a sharply declining price index (e.g., for women's dresses) unless the product that exits is linked with a new comparable product (next season's similar women's dress). To be precise, we drop ob-

servations when they are flagged as on temporary sale and are not observed with a regular price afterwards. In contrast, if we observe a price for the same quote line at a later point following the sale, we will include the sale observations even if there has been a version change. In the case of a version change, we compute the effective price change by adjusting for quality as in equation (66) below.

B.3.2 Quality Adjustments

When a BLS price collector identifies a version change of a particular product (e.g., a new version of the same rain coat), they determine whether the substitution is "comparable." If they deem it to be comparable, they assess whether a quality adjustment is necessary. Specifically, the price collector uses the code CP for a comparable substitution, the code QC for a substitution that is considered comparable after quality adjustment, and SR for non-comparable substitutions. For observations that are considered QC, the analyst will record a quality adjustment factor. This information is then used in the construction of the price relative for that product.

We follow an analogous procedure. We include price relatives at the time of version changes in our index construction only if the version change is comparable (i.e., CP or QC). In the case of QC substitutions, we make use of the reported quality adjustment using the formula

$$r_{it} = \left(\frac{P_{it}}{P_{i,t-\tau} + QA_{i,t-\tau,t}} \right)^{1/\tau}, \tag{66}$$

where $QA_{i,t-\tau,t}$ is the quality adjustment entered for the substitution.

B.3.3 Aggregation

Armed with these price relatives, we first aggregate to the product category level (ELI) within each state using a simple geometric average

$$R_{j,x,t} = \prod_{i \in j,x} r_{i,t},$$

where j is an ELI and x is a state.

Finally, we aggregate the ELI price relatives $R_{j,x,t}$ within sectors in each state using a weighted geometric average

$$R_{s,x,t} = \prod \left[(R_{j,x,t})^{W_j / \sum_{m \in s,x} W_m} \right],$$

where s denotes sector, and W_j is the expenditure weight of each ELI. These sectors can be defined broadly as all of non-tradeables or even the entire non-shelter CPI. We use expenditure weights that are constant across states and time. Specifically, we use the CPI expenditure weights for 1998.

B.4 Definition of Non-Tradeables Inflation

Below we list the ELIs that we categorize as non-tradeables. We define non-tradeables in a relatively conservative manner since including tradeable goods in our definition of what constitutes a non-tradeable good can lead to attenuation in the slope of the Phillips curve (if tradeable goods are price nationally). Our definition of non-tradeables is similar to the BLS service aggregation. It differs in two ways. First, we include ELIs in the Food Away from Home category as non-tradeables. Second we exclude several ELIs in Transportation Services, Utilities, and Truck Rentals. An important example is airline tickets. These have highly variable prices and are collected using a different procedure than other services in the CPI Research Database. See Nakamura and Steinsson (2008) for more discussion of the behavior of transportation services prices.

- education services
  - college tuition and fixed fees
  - elementary and high school tuition and fixed fees
  - day care and nursery school
  - technical and business school tuition and fixed fees
- telephone services
  - main station charges
  - interstate telephone services
- food away from home
  - lunch
  - dinner
  - candy, gum, etc.
  - breakfast or brunch
  - full service meals and snacks

  - limited service meals and snacks
  - food at employee sites and schools
  - food from vending machines and mobile vendors
  - board, catered events, and other food away from home
  - beer, ale, and other alcoholic malt beverages away from home
- other personal services
  - beauty parlor services for females
  - legal fees
  - funeral expenses
  - household laundry and dry cleaning, excluding coin-operated
  - shoe repair and other shoe services
  - clothing rental
  - replacement of setting for women's rings
  - safe deposit box rental
  - ax return preparation and other accounting fees
  - care of invalids, elderly and convalescents in the home
- housing services
  - housing at school, excluding board
  - lodging while out of town
  - tenants' insurance
  - electricity
  - utility natural gas service
  - residential water and sewer service
  - garbage/trash collection
  - gardening or lawn care services
  - moving, storage, freight express

  - repair of household appliance
  - reupholstery of furniture
  - inside painting and/or papering
- medical services
  - general medical practice
  - dentures, bridges, crowns, implants
  - optometrists/opticians
  - services by other medical professionals
  - hospital room inpatient
  - nursing and convalescent home care
- recreational services
  - community antenna or cable tv
  - prerecorded - video tapes and discs
  - other entertainment services
  - pet services
  - veterinarian services
  - photographer's fees
  - film processing
  - fees for participant sports
  - admission to movies, theaters, and concerts
  - admission to sporting events
  - fees for lessons or instructions
- transportation services
  - used cars
  - truck rental
  - other vehicle rental

- painting entire automobile
- vehicle inspection
- automotive brake work
- automobile insurance
- drivers license
- local automobile registration
- vehicle tolls
- automobile service clubs
- intercity bus fare
- intercity train fare
- passenger ship fares
- intracity mass transit
- taxi fare

B.5 Definition of Tradeable Employment Shares

We follow Mian and Sufi (2014) in defining the tradeable employment share as the share associated with the following sectors: "agriculture, forestry, fishing and hunting," "mining, quarrying, and oil and gas extraction," and manufacturing (SIC sectors A, B and D; and NAICS sectors 11, 21, and 31-33). The QCEW censors data if there are fewer than three establishments in the industry-state, or if one firm constitutes more than 80 percent of industry-state employment. 5% of NAICS 3 digit state-by-industry cells are censored, while 10% of SIC 2 digit state-by-industry cells are censored. If an industry-state observation is missing or censored in a given quarter, we exclude this observation when we calculate the instrument.

Anthracite mining is discontinued after 1987 in the SIC. We drop this industry. We also drop observations from California before 1978, due to the exceptionally volatile share of agricultural employment in California during 1976-1978.

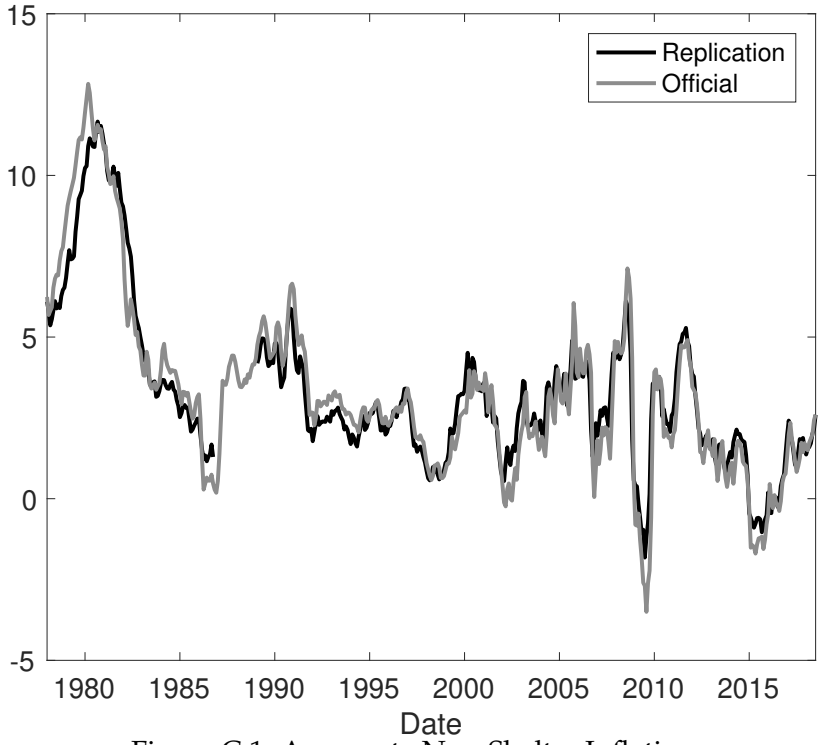


Figure C.1: Aggregate Non-Shelter Inflation

Note: The figure plots the 12-month non-shelter inflation rate for the US published by the Bureau of Labor Statistics (official) as well as the corresponding inflation rate using our methods (replication).

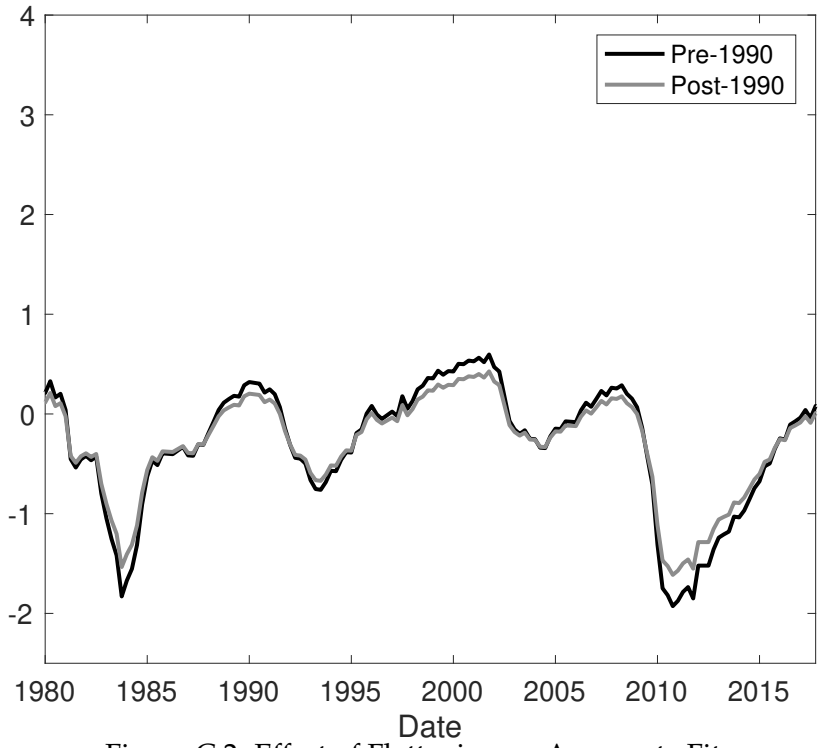


Figure C.2: Effect of Flattening on Aggregate Fit

Note: This figure plots the variation in inflation caused by changes in unemployment working through the slope of the Phillips curve according to our pre-1990 and post-1990 estimates of κ. In both cases we weight our non-shelter estimates of κ with our estimate of κ for rents.

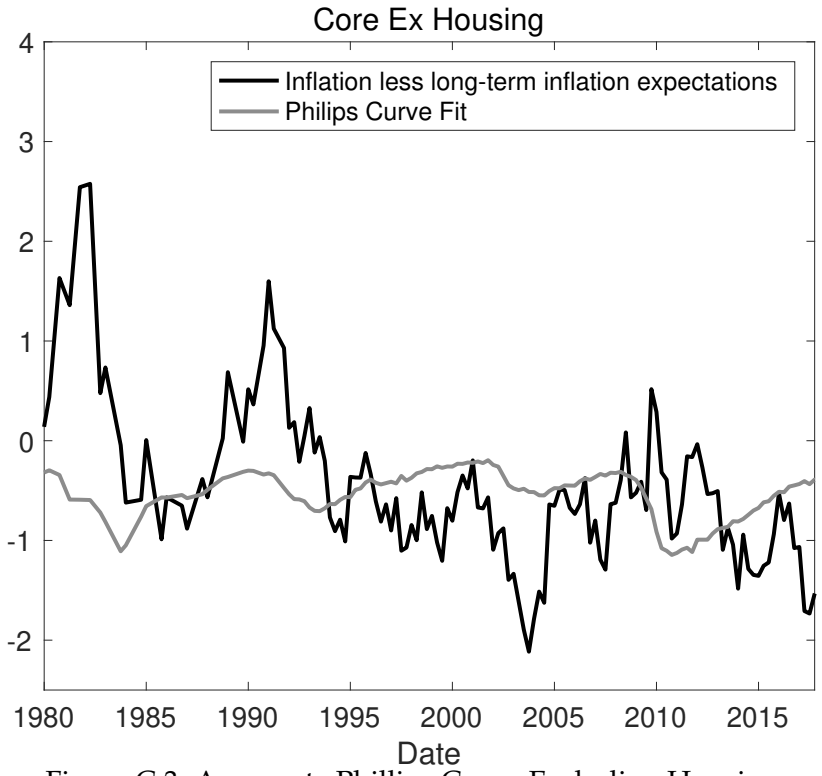


Figure C.3: Aggregate Phillips Curve Excluding Housing

Note: This figure shows the fit of the aggregate Phillips curve for core inflation excluding housing.

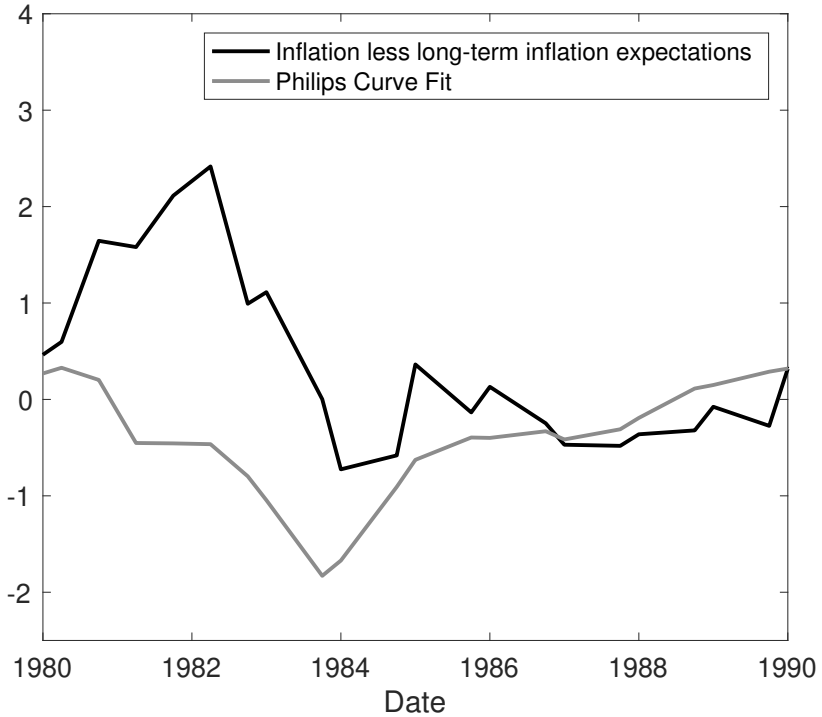


Figure C.4: Fit of the Aggregate Phillips Curve During the Volcker Disinflation

Note: This figure shows the fit of the aggregate Phillips curve for core inflation over the period 1980-1990. The gray line uses a weighted average of our pre-1990 non-shelter estimate for κ and our estimate of κ for rents.

Table C.1: First Stage Regressions with Future Sum of Unemployment and Relative Prices

	(1)	(2)	(3)	(4)
	Panel A: Future Sum of Unemployment			
Lagged Unemployment	7.029	3.661	5.477	
	(0.635)	(0.474)	(0.510)	
Lagged Tradeable Demand				-4.465
				(0.594)
Lagged Relative Price	0.181	-0.178	0.259	0.833
	(0.160)	(0.202)	(0.565)	(0.516)
	Panel B: Future Sum of Relative Price of Non-Tradeables			
Lagged Unemployment	0.520	1.713	-1.973	
	(0.889)	(0.684)	(1.096)	
Lagged Tradeable Demand				1.011
				(1.028)
Lagged Relative Price	18.572	16.895	13.081	13.710
	(0.194)	(0.288)	(1.000)	(1.283)
State Effects		✓	✓	✓
Time Effects			✓	✓

Note: This table presents results for the first stage regressions for our estimation of κ. In Panel A, the outcome is the discounted future sum of quarterly state unemployment, in percentage points, truncated at 20 quarters. In Panel B the outcome is the discounted future sum of the relative price of non-tradeables, in 100 x log points, truncated at 20 quarters. In the first three columns, the regressors are the fourth lags of unemployment and the relative price of non-tradeables, in 100 x log points. In the final column, the regressors are the fourth lags of tradeable demand and the relative price of non-tradeables. Standard errors are reported in parentheses, clustered by state. The fixed effects included for each column are reported at the bottom of the table. All regressions are unweighted and have 3323 observations.

Table C.2: Estimates of λ from Regression (17)

	No Fixed Effects	No Time Effects	Lagged Unempl.	Tradeable Demand IV
	(1)	(2)	(3)	(4)
λ	0.0010	0.0022	0.0029	0.0020
	(0.0001)	(0.0002)	(0.0009)	(0.0007)
State Effects		✓	✓	✓
Time Effects			✓	✓

Note: This table presents estimates of λ, the coefficient on the present value of relative prices from regression equation (17). The outcome variable is cumulative non-tradeable inflation over four quarters, measured in percentage points. The regressors are discounted future sums of quarterly state unemployment, in percentage points, and the relative price of non-tradeables, in 100 x log points. Each of these is truncated at 20 quarters. In the first three columns we instrument using the fourth lags of quarterly state unemployment and the relative price of non-tradeables (this is OLS for ψ). In the fourth column, we replace lagged unemployment with our tradeable demand instrument among the instruments. In all columns, we estimate λ by two-sample two stage least squares, and apply the correction to our standard errors from Chodorow-Reich and Wieland (2019). The sample period is 1978-2018. Standard errors are reported in parentheses, clustered by state. Fixed effects for each column are reported at the bottom of the table. All regressions are unweighted. The number of observations is 3323 in the first three columns with slightly fewer in the last column due to differencing.

Table C.3: Estimate of κ as Calibrated Value of β Varies

	$\beta = 0.99$	$\beta = 0.95$	$\beta = 0.90$
	(1)	(2)	(3)
κ	0.0062	0.0084	0.0116
	(0.0025)	(0.0033)	(0.0046)
State Effects	✓	✓	✓
Time Effects	✓	✓	✓

Note: This table presents estimates of κ from regression equation (17), with different calibrated values of β. The outcome variable is cumulative non-tradeable inflation over four quarters, measured in percentage points. The regressors are discounted future sums of quarterly state unemployment, in percentage points, and the relative price of non-tradeables, in 100 x log points. Both sums are truncated at 20 quarters. In all columns, we estimate κ by two-sample two stage least squares, and apply the correction to our standard errors from Chodorow-Reich and Wieland (2019). We include time and state fixed effects. The sample period is 1978-2018. Standard errors are reported in parentheses, clustered by state.

Table C.4: Estimates of κ for Different Truncation Lengths of Discounted Sums

	$T = 10$	$T = 20$	$T = 30$	$T = 40$
	(1)	(2)	(3)	(4)
κ	0.0100	0.0062	0.0044	0.0051
	(0.0038)	(0.0025)	(0.0019)	(0.0021)
State Effects	✓	✓	✓	✓
Time Effects	✓	✓	✓	✓

Note: This table presents estimates of κ from regression specification (17), for different truncation lengths of the discounted sums on the right-hand-side. The outcome variable is cumulative non-tradeable inflation over four quarters, measured in percentage points. The regressors are the present values of quarterly state unemployment, in percentage points, and the relative price of non-tradeables, in 100 x log points. We vary to truncation point for these sums between T=10 and t=40 across the columns in the table. We include time and state fixed effects. The sample period is 1978-2018. Standard errors are reported in parentheses, clustered by state. We use a two sample two stage least squares regression and apply the correction to our standard errors from Chodorow-Reich and Wieland (2019).

Table C.5: Slope of the Regional Phillips Curve: Rents

	No Fixed Effects	No Time Effects	Lagged Unempl.
	(1)	(2)	(3)
	Panel A: Estimates of κ from equation (17)		
κ	0.0074	0.0179	0.0243
	(0.0006)	(0.0014)	(0.0053)
	Panel B: Estimates of ψ from equation (19)		
ψ	0.268	0.356	0.603
	(0.041)	(0.044)	(0.124)
State Effects		✓	✓
Time Effects			✓

Note: The table presents estimates of ψ, and κ for rents. The outcome variable is the state-level annual rent inflation rate, measured in percentage points from the American Community Survey for the years 2001 to 2017 that we gathered from IPUMS USA . In Panel A, the regressor of interest is the discounted future sum of annual state unemployment, measured in percentage points. In Panel B, the regressor of interest is lagged state unemployment, measured in percentage points. We estimate κ by two-sample two stage least squares, and apply the correction to our standard errors from Chodorow-Reich and Wieland (2019). Standard errors are reported in parentheses and clustered by state. Controls for each column are reported at the bottom of the table. All regressions are unweighted.

References

ADAM, K. AND M. PADULA (2011): "Inflation Dynamics and Subjective Expectations in the United States," *Economic Inquiry*, 49, 13–25.

ADAO, R., M. KOLESAR, AND E. MORALES (2019): "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 134, 1949–2010.

ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2020): "Transparency in Structural Research," *Journal of Business and Economic Statistics*, 38, 711–722.

ANGELETOS, G.-M. AND C. LIAN (2018): "Forward Guidance without Common Knowledge," *American Economic Review*, 108, 2477–2512.

BABB, N. AND A. K. DETMEISTER (2017): "Nonlinearities in the Phillips Curve for the United States: Evidence Using Metropolitan Data," Finance and Economics Discussion Series 2017-070. Washington: Board of Governors of the Federal Reserve System.

BALL, L. AND S. MAZUMDER (2011): "Inflation Dynamics and the Great Recession," *Brookings Papers on Economic Activity*, 2011, 337–381.

——— (2019): "A Phillips Curve with Anchored Expectations and Short-Term Unemployment," *Journal of Money, Credit and Banking*, 51, 111–137.

BARNICHON, R. AND G. MESTERS (2019): "Identifying Modern Macro Equations with Old Shocks," *Quarterly Journal of Economics*, forthcoming.

BARTIK, T. J. (1991): "Who Benefits from State and Local Economic Development Policies?" WE Upjohn Institute for Employment Research.

BERAJA, M., E. HURST, AND J. OSPINA (2019): "The Aggregate Implications of Regional Business Cycles," *Econometrica*, 87, 1789–1833.

BERNANKE, B. S. (2007): "Inflation Expectations and Inflation Forecasting," speech given at the NBER Summer Institute, July 2007.

BIANCHI, F. AND C. ILUT (2017): "Monetary/Fiscal Policy Mix and Agents' Beliefs," *Review of Economic Dynamics*, 26, 113–139.

BLANCHARD, O. (2016): "The Phillips Curve: Back to the '60s?" *American Economic Review: Papers and Proceedings*, 106, 31–34.

BUREAU OF LABOR STATISTICS (1982): *BLS Handbook of Methods*, no. v. 1-2 in BLS Handbook of Methods, U.S. Department of Labor, Bureau of Labor Statistics.

CALVO, G. A. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

CARVALHO, C., S. EUSEPI, E. MOENCH, AND B. PRESTON (2021): "Anchored Inflation Expectations," Working Paper, PUC-Rio.

CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2011): "Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins," *American Economic Review*, 101, 471–75.

CHODOROW-REICH, G. AND J. WIELAND (2019): "Secular Labor Reallocation and Business Cy-

cles," *Journal of Political Economy*, forthcoming.

CHRISTIANO, L., M. EICHENBAUM, AND M. TRABANDT (2015): "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, 7, 110–167.

COGLEY, T. AND A. M. SBORDONE (2008): "Trend inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips curve," *American Economic Review*, 98, 2101–26.

COIBION, O. AND Y. GORODNICHENKO (2012): "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120, 116–159.

——— (2015a): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105, 2644–78.

——— (2015b): "Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, 7, 197–232.

COIBION, O., Y. GORODNICHENKO, AND R. KAMDAR (2018): "The Formation of Expectations, Inflation, and the Phillips curve," *Journal of Economic Literature*, 56, 1447–91.

CRUMP, R. K., S. EUSEPI, MARC GIANNONI, AND A. SAHIN (2019): "A Unified Approach to Measuring u*," *Brookings Papers on Economic Activity*, 2019, 143–214.

DEL NEGRO, M. (1998): "Aggregate Risk Sharing Across US States and Across European Countries," Working Paper, Yale University.

DEL NEGRO, M., M. P. GIANNONI, AND F. SCHORFHEIDE (2015): "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics*, 7, 168–196.

DEL NEGRO, M., M. LENZA, G. E. PRIMICERI, AND A. TAMBALOTTI (2020): "What's Up With the Phillips Curve," *Brookings Papers on Economic Activity*, forthcoming.

DUPRAZ, S., E. NAKAMURA, AND J. STEINSSON (2020): "A Plucking Model of the Business Cycle," Working Paper, University of California, Berkeley.

ERCEG, C. J. AND A. T. LEVIN (2003): "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics*, 50, 915–944.

FEENSTRA, R. C., P. LUCK, M. OBSTFELD, AND K. N. RUSS (2018): "In search of the Armington elasticity," *Review of Economics and Statistics*, 100, 135–150.

FITZGERALD, T. J., C. JONES, M. KULISH, AND J. P. NICOLINI (2020): "Is There a Stable Relationship between Unemployment and Future Inflation?" Working Paper, Federal Reserve Bank of Minneapolis.

FITZGERALD, T. J. AND J. P. NICOLINI (2014): "Is There a Stable Relationship between Unemployment and Future Inflation? Evidence from US Cities," Working Paper, Federal Reserve Bank of Minneapolis.

FRIEDMAN, M. (1968): "The Role of Monetary Policy," *American Economic Review*, 58, 1–17.

FUHRER, J. C. AND G. MOORE (1995): "Inflation Persistence," *Quarterly Journal of Economics*, 110, 127–159.

GABAIX, X. (2020): "A Behavioral New Keynesian Model," *American Economic Review*, 110, 2271–2327.

GALI, J. (2008): *Monetary Policy, Inflation, and the Business Cycle*, Princeton, NJ: Princeton University Press.

GALI, J. AND M. GERTLER (1999): "Inflation Dynamics: A Structural Econometrics Analysis," *Journal of Monetary Economics*, 44, 195–222.

GALI, J., M. GERTLER, AND J. D. LOPEZ-SALIDO (2005): "Robustness of the Estimates of the Hybrid New Keynesian Phillips Curve," *Journal of Monetary Economics*, 52, 1107–1118.

GERTLER, M. AND J. LEAHY (2008): "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, 116, 533–572.

GILCHRIST, S., R. SCHOENLE, J. SIM, AND E. ZAKRAJŠEK (2017): "Inflation dynamics during the financial crisis," *American Economic Review*, 107, 785–823.

GOODFRIEND, M. AND R. G. KING (2005): "The Incredible Volcker Disinflation," *Journal of Monetary Economics*, 52, 981–1015.

GORDON, R. (1982): "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment," in *Workers, Jobs, and Inflation*, ed. by M. N. Baily, Washington, DC: Brookings Institute, 89–158.

GREENWOOD, J., Z. HERCOWITZ, AND G. W. HUFFMAN (1988): "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 78, 402–417.

HOOPER, P., F. S. MISHKIN, AND A. SUFI (2019): "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or is It Just Hibernating?" NBER Working Paper No. 25792.

IRELAND, P. N. (2007): "Changes in the Federal Reserve's Inflation Target: Causes and Consequences," *Journal of Money, Credit and Banking*, 39, 1851–1882.

ITSKHOKI, O. AND D. MUKHIN (2017): "Exchange rate disconnect in general equilibrium," Tech. rep., National Bureau of Economic Research.

JORGENSEN, P. L. AND K. J. LANSING (2019): "Anchored Inflation Expectations and the Flatter Phillips Curve," Working Paper, Federal Reserve Bank of San Francisco.

KILEY, M. T. (2015a): "An Evaluation of the Inflationary Pressure Associated with Short- and Long-Term Unemployment," *Economics Letters*, 137, 5–9.

——— (2015b): "Low Inflation in the United States: A Summary of Recent Research," FEDS Notes, Board of Governors of the Federal Reserve System.

KLENOW, P. J. AND O. KRYVTSOV (2008): "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation," *Quarterly Journal of Economics*, 123, 863–904.

MAVROEIDIS, S., M. PLAGBORG-MØLLER, AND J. H. STOCK (2014): "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 52, 124–88.

MCLEAY, M. AND S. TENREYRO (2019): "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual*, 34, 199–255.

MIAN, A. AND A. SUFI (2014): "What Explains the 2007–2009 Drop in Employment?" *Econometrica*, 82, 2197–2223.

MISHKIN, F. S. (2007): "Inflation Dynamics," *International Finance*, 10, 317–334.

NAKAMURA, E. AND J. STEINSSON (2008): "Five Facts About Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464.

——— (2014): "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 104, 753–92.

——— (2018): "Identification in Macroeconomics," *Journal of Economic Perspectives*, 32, 59–86.

NAKAMURA, E., J. STEINSSON, P. SUN, AND D. VILLAR (2018): "The Elusive Costs of Inflation: Price Dispersion during the US Great Inflation," *Quarterly Journal of Economics*, 133, 1933–1980.

NEFTÇI, S. (1984): "Are Economic Time Series Asymmetric over the Business Cycle?" *Journal of Political Economy*, 92, 307–328.

NGUYEN, H. (2014): "Demand-Driven Propagation: Evidence from the Great Recession," Working Paper, World Bank.

NISHIZAKI, K. AND T. WATANABE (2000): "Output-Inflation Trade-Off at Near-Zero Inflation Rates," *Journal of the Japanese and International Economies*, 14, 304–326.

OBSTFELD, M. AND K. ROGOFF (1983): "Speculative Hyperinflations in Maximizing Models: Can We Rule Them Out?" *Journal of Political Economy*, 91, 675–687.

PHELPS, E. S. (1967): "Phillips Curves, Expectations of Inflaiton and Optimal Unemployment over Time," *Economica*, 34, 254–281.

PHILLIPS, A. W. (1958): "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," *Economica*, 25, 283–299.

POOLE, R., F. PTACEK, AND R. VERBRUGGE (2005): "Treatment of Owner-Occupied Housing in the CPI," Working Paper, Bureau of Labor Statistics.

ROBERTS, J. M. (1995): "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit and Banking*, 27, 975–984.

ROTEMBERG, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 60, 1187–1211.

ROTEMBERG, J. J. AND M. WOODFORD (1997): "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 12, 297–346.

RUBBO, E. (2020): "Networks, Phillips Curves, and Monetary Policy," Working Paper, Harvard University.

RUDD, J. AND K. WHELAN (2005): "New Tests of the New-Keynesian Phillips Curve," *Journal of Monetary Economics*, 52, 1167–1181.

RUGGLES, S., S. FLOOD, S. FOSTER, R. GOEKEN, J. PACAS, M. SCHOUWEILER, AND M. SOBEK (2021): "IPUMS USA: Version 11.0 [dataset] Minneapolis, MN: IPUMS, 2021." .

SAMUELSON, P. AND R. SOLOW (1960): "Analytical Aspects of Anti-Inflation Policy," *American Economic Review*, 50, 177–194.

SARGENT, T. J. (1982): "The Ends of Four Big Inflations," in *Inflation: Causes and effects*, University of Chicago Press, 41–98.

SBORDONE, A. M. (2002): "Prices and Unit Labor Costs: A New Test of Price Stickiness," *Journal*

of Monetary Economics, 49, 265–292.

SICHEL, D. (1993): "Business Cycle Asymmetry: A Deeper Look," *Economic Inquiry*, 31, 224–236.

STOCK, J. H. AND M. W. WATSON (2007): "Why Has US Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 39, 3–33.

——— (2019): "Slack and Cyclically Sensitive Inflation," NBER Working Paper No. 25987.

TAYLOR, J. B. (1979): "Staggered Contracts in a Macro Model," *American Economic Review*, 69, 108–113.

——— (1980): "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 88, 1–24.

WOODFORD, M. (1998): "Doing without Money: Controlling Inflation in a Post-Monetary World," *Review of Economic Dynamics*, 1, 173–219.

——— (2003): *Interest and Prices*, Princeton, NJ: Princeton University Press.