

NBER WORKING PAPER SERIES

THE SLOPE OF THE PHILLIPS CURVE:
EVIDENCE FROM U.S. STATES

Jonathon Hazell
Juan Herreño
Emi Nakamura
Jón Steinsson

Working Paper 28005
<http://www.nber.org/papers/w28005>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2020

We would like to thank Thomas 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 Olivier Blanchard, Yuriy Gorodnichenko, Pierre-Olivier Gourinchas, John Williams, Regis Barnichon 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Jonathon Hazell, Juan Herreño, Emi Nakamura, and Jón Steinsson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Slope of the Phillips Curve: Evidence from U.S. States
Jonathon Hazell, Juan Herreño, Emi Nakamura, and Jón Steinsson
NBER Working Paper No. 28005
October 2020
JEL No. E30

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 Phillips curve is very flat and was very flat 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 re-inflation 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 since the 1990s is mostly due to long-run inflationary expectations becoming more firmly anchored.

Jonathon Hazell
Department of Economics
Princeton University
214 Julis-Rabinowitz Center
Princeton, NJ 08544
jadhazell@gmail.com

Juan Herreño
Columbia Business School
Uris Hall
3022 Broadway
New York, NY 10027
jdh2181@columbia.edu

Emi Nakamura
Department of Economics
University of California, Berkeley
685 Evans Hall
Berkeley, CA 94720
and NBER
enakamura@berkeley.edu

Jón Steinsson
Department of Economics
University of California, Berkeley
671 Evans Hall
Berkeley, CA 94720
and NBER
jsteinsson@berkeley.edu

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:

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

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 inflationary expectations are adaptive: $\beta E_t \pi_{t+1} = \pi_{t-1}$ in equation (1). This yields the accelerationist Phillips curve:

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

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).

The insensitivity of inflation to changes in unemployment over the past few decades has led many economists to suggest that the Phillips curve has disappeared—or is “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 then gave way to “missing reflation” in the late 2010s as unemployment fell to levels



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. Author's calculations. The figure replicates Figure 1 from Stock and Watson (2019).

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




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. 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. The inflation expectations we use before 1990 can be found at: <https://www.philadelphiafed.org/-/media/research-and-data/real-time-center/survey-of-professional-forecasters/historical-data/additional-cpie10.xls?la=en>

regime in the United States that caused the rapid fall in long-run inflation expectations we observe 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

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

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 co-vary 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 Babb and Detmeister (2017) and Hooper, Mishkin, and Sufi (2019) 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 an 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.

Our cross-sectional estimates indicate that the Phillips curve is very flat and was very flat even during the 1980s: a one percentage point increase in unemployment lowered inflation over time by a mere 0.46 percentage points in the 1980s. This implies that the 5 percentage points increase in unemployment during the Volcker disinflation lowered inflation by only about 2.3 percentage points. 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 about 1.8. This contrasts sharply with empirical specifications that make use of time series variation: a specification without time fixed effects yields a 45 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.

Figure 2 shows that core CPI inflation fell by about 6 percentage points during the Volcker disinflation (as measured by the BLS’s research series which employs modern methods back in

time). Our estimates indicate that only a little more than 1/3 of this fall was due to the rise in unemployment over this period. Figure 2 shows that long-run inflationary expectations fell by about 4 percentage points from 1981 to 1987, 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.

Our estimates of the slope of the Phillips curve imply essentially no “missing disinflation” during the Great Recession or “missing reflation” in the late 2010s or late 1990s. We conclude that the stability of inflation since 1990 is due to long-run inflationary 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 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.¹

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”

¹For 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.

instrument building on insights from Nguyen (2014). This instrument is based on the idea that 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.² 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 inflationary 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), Geerolf (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 through the lens of fully specified DSGE models (see, e.g., Del Negro et al., 2015; Christiano et al., 2015; Gilchrist et al., 2017).

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.

²Much 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% out of this expenditure weight is 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 inflationary expectations in determining aggregate inflation. To this end, we solve equation (1) forward to get

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

where $\omega_t \equiv E_t \sum_{j=0}^{\infty} \beta^j (\kappa u_{t+j}^n + \nu_{t+j})$. This equation illustrates how inflation at time t is determined by the path of unemployment out into the infinite future.³ 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

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

Assuming that shocks to u_t^n 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

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

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

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

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

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

³Sbordone (2002) and Rudd and Whelan (2005) develop approaches to estimating the Phillips Curve on aggregate data using versions of equation (4).

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.⁴

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 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 inflationary expecta-

⁴Equations (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 $E_t u_{t+\infty}$ is a constant independent of long-run inflation expectations. But the solved-forward Phillips curve has a $E_t\pi_{t+\infty}$ term with a coefficient of one, i.e., it takes the form of equation (6) as opposed to equation (4).




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.

tions 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.⁵

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.⁶ It is also sensitive to whether the expectations data used come from the SPF or from the Michigan Survey of Consumers (Coibion and Gorodnichenko, 2015b) emphasize and 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 fixed effects “difference out” changes in long run inflation expectations. The model also illustrates the importance of using non-tradeable inflation to estimate the regional Phillips curve.

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.

⁵These 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.

⁶We discuss this in more detail in appendix B.1.

The population of the home region is denoted by ζ . 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.⁷ 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 consumption of a composite consumption good, N_{Ht} is 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

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

where φ is the household’s Frisch elasticity of labor supply, σ is 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 tradeables C_{Ht}^T and non-tradeables C_{Ht}^N given by

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

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

⁷In other words, we are considering an economy in the cashless limit (Woodford, 1998, 2003).

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

$$C_{Ht}^T = \left[\tau_{Ht}^H \frac{1}{\eta} C_{Ht}^{TH} \frac{\eta-1}{\eta} + \tau_{Ht}^F \frac{1}{\eta} C_{Ht}^{TF} \frac{\eta-1}{\eta} \right]^{\frac{\eta}{\eta-1}}, \quad (8)$$

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-produced 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 home bias in tradeable consumption, so $\tau_{Ht}^H = \tau_{Ht}^F = \zeta$.

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

$$\int_0^1 C_{Ht}^N(z) P_{Ht}^N(z) dz + \int_0^1 C_{Ht}^{TH}(z) P_{Ht}^T(z) dz + \int_0^1 C_{Ht}^{TF}(z) P_{Ft}^T(z) dz + E_t [M_{Ht,t+1} B_{H,t+1}]$$

$$\leq B_{Ht} + W_{Ht}N_{Ht} + \int_0^1 \Xi_{Ht}^N(z)dz + \int_0^1 \Xi_{Ht}^T(z)dz$$

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, and $\Xi_{Ht}^N(z)$ and $\Xi_{Ht}^T(z)$ are the profits of non-tradeable and tradeable firms producing varieties z , respectively, 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.

Labor is immobile across regions. Within each region, there is a single labor market. Workers in region H receive a nominal wage W_{Ht} . Let P_{Ht} denote the lowest cost of purchasing a unit of the composite consumption good C_{Ht} . Household optimization regarding the trade-off between current consumption and current labor supply yields the following labor supply curve:

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

where 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

$$\chi N_{Ht}^{\varphi-1} = \frac{W_{Ht}}{P_{Ht}}. \quad (9)$$

Household optimization regarding the trade-off between current consumption and consumption in the next period yields the following consumption Euler equation:

$$\beta R_t^n E_t \left[\frac{u_c(C_{H,t+1}, N_{H,t+1})}{u_c(C_{Ht}, N_{Ht})} \frac{P_{Ht}}{P_{H,t+1}} \right] = 1. \quad (10)$$

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 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:

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

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

where P_{Ht} is a price index that gives the minimum cost of purchasing a unit of C_{Ht} and 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. Utility maximization, furthermore, implies the following demand curves for each of the varieties of goods produced in the economy:

$$C_{Ht}^N(z) = C_{Ht}^N \left(\frac{P_{Ht}^N(z)}{P_{Ht}^N} \right)^{-\theta} \quad C_{Ht}^{TH}(z) = C_{Ht}^{TH} \left(\frac{P_{Ht}^T(z)}{P_{Ht}^T} \right)^{-\theta} \quad C_{Ht}^{TF}(z) = C_{Ht}^{TF} \left(\frac{P_{Ft}^T(z)}{P_{Ft}^T} \right)^{-\theta} \quad (11)$$

The cost minimizing price indexes are given by

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

$$P_{Ht} = \left[\phi_N P_{Ht}^N{}^{1-\eta} + \phi_T \tau_{Ht}^H P_{Ht}^T{}^{1-\eta} + \phi_T \tau_{Ht}^F P_{Ft}^T{}^{1-\eta} \right]^{\frac{1}{1-\eta}}. \quad (12)$$

As we noted above, the problem of the foreign household is analogous. We therefore refrain from describing it in detail here. However, for simplicity, we do not allow for tradeable demand shocks to foreign tradeable consumption as we did 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

$$Y_{Ht}^N(z) = Z_{Ht}^N N_{Ht}^N(z), \quad (13)$$

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 acts to maximize its value

$$E_t \sum_{j=0}^{\infty} M_{Ht,t+j} [P_{H,t+j}^N(z) Y_{H,t+j}^N(z) - W_{H,t+j} N_{H,t+j}^N(z)].$$

Its demand is given by

$$Y_{Ht}^N(z) = \zeta C_{Ht}^N \left(\frac{P_{Ht}^N(z)}{P_{Ht}^N} \right)^{-\theta}.$$

The firm must satisfy demand. This implies that production must satisfy the constraint

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

Firm z takes the wage W_{Ht} as given. Optimal choice of labor by the firm implies that

$$W_{Ht} = S_{Ht}^N(z) Z_{Ht}^N, \quad (14)$$

where $S_{Ht}^N(z)$ is the firm's nominal marginal cost, i.e. the Lagrange multiplier on its output constraint. Firm z can reoptimize its price with probability $1 - \alpha$ as in Calvo (1983). With probability α it must keep its price unchanged. Optimal price setting by firm z in periods when it can change its price implies

$$\sum_{k=0}^{\infty} \alpha^k E_t \left[M_{Ht,t+k} Y_{H,t+k}^N(z) \left(P_{Ht}^{N*}(z) - \frac{\theta}{\theta-1} S_{H,t+k}^N(z) \right) \right] = 0, \quad (15)$$

where $P_{Ht}^{N*}(z)$ is the price the firm chooses if it has the opportunity to reset its price in period t . Intuitively, the firm sets its price equal to a constant markup over a weighted average of current and expected future marginal cost.

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

$$Y_{Ht}^T(z) = Z_{Ht}^T N_{Ht}^T(z)$$

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

Firm z maximizes its value given by

$$E_t \sum_{j=0}^{\infty} M_{Ht,t+j} \left[P_{H,t+j}^T(z) Y_{H,t+j}^T(z) - W_{H,t+j} N_{H,t+j}^T(z) \right].$$

Demand for the output of firms in the tradeable sector comes from both the home and foreign

regions. Firm z 's demand is thus given by

$$Y_{Ht}^T(z) = (\zeta C_{Ht}^{TH} + (1 - \zeta) C_{Ft}^{TH}) \left(\frac{P_{Ht}^T(z)}{P_{Ht}^T} \right)^{-\theta}.$$

The firm's optimal choice of labor implies that

$$W_{Ht} = S_{Ht}^T(z) Z_{Ht}^T(z).$$

where $S_{Ht}^T(z)$ is the firm's nominal marginal cost. The tradeable goods firm also have an opportunity to change their price with probability $1 - \alpha$. Optimal choice of a new reset price at these times implies

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

The problems of foreign firms are analogous to those of home firms. For this reason, we do not detail them here.

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

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

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 $\bar{\pi}_t$.⁸ Since

⁸Prior work that allows for a time-varying inflation target includes Stock and Watson (2007), Ireland (2007) and Cogley and Sbordone (2008).

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).

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 yields the following regional Phillips Curve for the inflation of non-tradeable goods:

$$\pi_{Ht}^N = \beta E_t \pi_{H,t+1}^N - \kappa \hat{u}_{Ht} - \lambda \hat{p}_{Ht}^N + \nu_{Ht}^N, \quad (16)$$

and aggregate Phillips Curve for overall inflation:

$$\pi_t = \beta E_t \pi_{t+1} - \kappa \hat{u}_t + \nu_t, \quad (17)$$

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.1.

Equations (16) and (17) 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.2, 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 (16) and (17). 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.3. 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 (16) and (17) is the presence of the relative price of non-tradeables term $\lambda \hat{p}_{Ht}^N$ in equation (16). This term implies that inflation in the non-tradeables sector will be lower the higher is the relative price of non-tradeables. The reason for this 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 (16) — forward to obtain

$$\pi_{Ht}^N = -E_t \sum_{j=0}^{\infty} \beta^j (\kappa \tilde{u}_{H,t+j} + \lambda \hat{p}_{H,t+j}^N) + E_t \pi_{t+\infty} + \omega_{Ht}^N, \quad (18)$$

where $\tilde{u}_{Ht} = u_{Ht} - E_t u_{H,t+\infty}$ and $\omega_{Ht}^N = E_t \sum_{j=0}^{\infty} \beta^j \nu_{H,t+j}^N$. For expositional simplicity, we now introduce an approximation to the model presented above. We assume that both \tilde{u}_{Ht} and \hat{p}_{Ht}^N follow AR(1) processes with autocorrelation coefficients equal to ρ_u and ρ_{pN} , respectively. In this case, equation (18) simplifies to

$$\pi_{Ht}^N = -\psi \tilde{u}_{Ht} - \delta \hat{p}_{Ht}^N + E_t \pi_{t+\infty} + \omega_{Ht}^N, \quad (19)$$

where $\psi = \kappa / (1 - \beta \rho_u)$ and $\delta = \lambda / (1 - \beta \rho_{pN})$.

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 $E_t \pi_{t+\infty}$ term in equation (19) — is constant across regions.⁹ This implies that we can adopt an empirical specification that replaces the $E_t \pi_{t+\infty}$ term in equation (19) with time fixed effects:

$$\pi_{it}^N = \alpha_i - \psi u_{it} - \delta p_{it}^N + \gamma_t + \varepsilon_{it}, \quad (20)$$

where π_{it}^N , u_{it} , and p_{it}^N are non-tradeable inflation, unemployment, and the relative price of non-tradeable goods, respectively, for region i at time t , α_i denotes a set of region fixed effects, γ_t denotes a set of time fixed effects, and ε_{it} is a regression residual.

Intuitively, while short-run inflation expectations ($E_t \pi_{t+1}$) will differ across regions due to differences in their economic circumstances, long-run inflation expectations ($E_t \pi_{t+\infty}$) are independent of the current business cycle. They are determined solely 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

⁹Our empirical specification actually allows for the possibility of a constant difference in inflation rates across regions (e.g., California becoming increasingly more expensive than Kansas for ever). Such a constant difference would be picked up by the state fixed effects.

fixed effects. As we have discussed earlier in the paper, changes in the monetary regime over time that yield variation in long-run inflation expectations may confound estimates of the slope of the Phillips curve using aggregate variation.

In principle, researchers can control for inflation expectations when estimating the slope of the Phillips curve. This procedure, if successful, obviates the need to control for changes in the monetary regime with time fixed effects in a panel setting. Controlling well for inflation expectations may, however, be quite difficult. One strand of the literature makes a rational expectations assumption and instruments for realized future inflation using lagged variables. Mavroidis et al. (2014) emphasize that this approach suffers from a weak instruments problem. Another strand of the literature uses survey measures of inflation expectations. But different measures differ substantially suggesting that they are imperfect indicators of inflation expectations.

A curious feature of the 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 (16) and (20) provides a simple explanation for this discrepancy. The slope coefficient in equation (16) is κ , while the slope coefficient in equation (20) is $\psi = \kappa/(1 - \beta\rho_u)$. Since unemployment is quite persistent, $\psi \gg \kappa$. However, our analysis highlights that under the simplifying assumption that unemployment follows an AR(1) it is relatively straightforward to convert estimates of ψ into estimates of κ . For this, one simply needs estimates (or assumptions) for ρ_u and β .

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 $E_t\pi_{t+1}$ 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.

The difference between κ and ψ arises due to the different ways equations (16) and (20) capture the effects of expected future unemployment on current inflation. In equation (16), 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 (20) 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.¹⁰

An additional advantage of estimating a specification such as equation (20) rather than equation (16) 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. Furthermore, estimates of equation (16) rely quite heavily on the exact timing of inflation and inflation expectations implied by the New Keynesian Phillips curve.

3.4 Relaxing Full Information Rational Expectations

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. Under this assumption, one can show that manipulating the aggregate Phillips curve — equation (17) — yields

$$\pi_t = -\frac{\kappa}{1 - \rho_{\bar{u}}^F \beta} \tilde{u}_t + F_t \pi_{t+\infty} + \tilde{\omega}_t, \quad (21)$$

where F_t is 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 $\rho_{\bar{u}}^F$ is agents' subjective belief about the autoregressive coefficient governing the persistence of fluctuations in unemployment. Notice that if $\rho_{\bar{u}}^F < \rho_{\bar{u}}$, the Phillips Curve is less forward looking than the rational expectations Phillips curve. Rational expectations is the special case where $\rho_{\bar{u}}^F = \rho_{\bar{u}}$ and $F_t \pi_{t+\infty} = E_t \pi_{t+\infty}$. Coibion and Gorodnichenko (2012, 2015a) provide evidence consistent 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.

¹⁰McLeay 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.

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 visiting outlets to collect prices. The database does not include the rent prices used to construct the shelter component of the database. 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

$$r_{i,t} = \left(\frac{P_{i,t}}{P_{i,t-\tau}} \right)^{1/\tau} . \quad (22)$$

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 $r_{i,t} = \dots = r_{i,t-\tau+1}$ 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.¹¹ 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 weights from the Consumer Expenditure Survey (CEX) from 1998 to perform this aggregation.

Our empirical analysis focuses on non-tradeables but we also construct state-level price in-

¹¹See the appendix to Nakamura and Steinsson (2008) for a list of the ELIs used in the construction of the CPI.

dexes 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). Appendix B.4 provides a detailed list of which ELI categories we classify as non-tradeable.

We find that there is much more variability 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 contrast, 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.¹²

Our method for calculating state-level price indexes aims to approximate the non-shelter price index published by the BLS. Figure 4 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.

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. Industry-state employment data are

¹²In the present draft, we drop Arizona due to anomalous trends that we have not been able to investigate due to Covid-19 related access restrictions at the BLS.

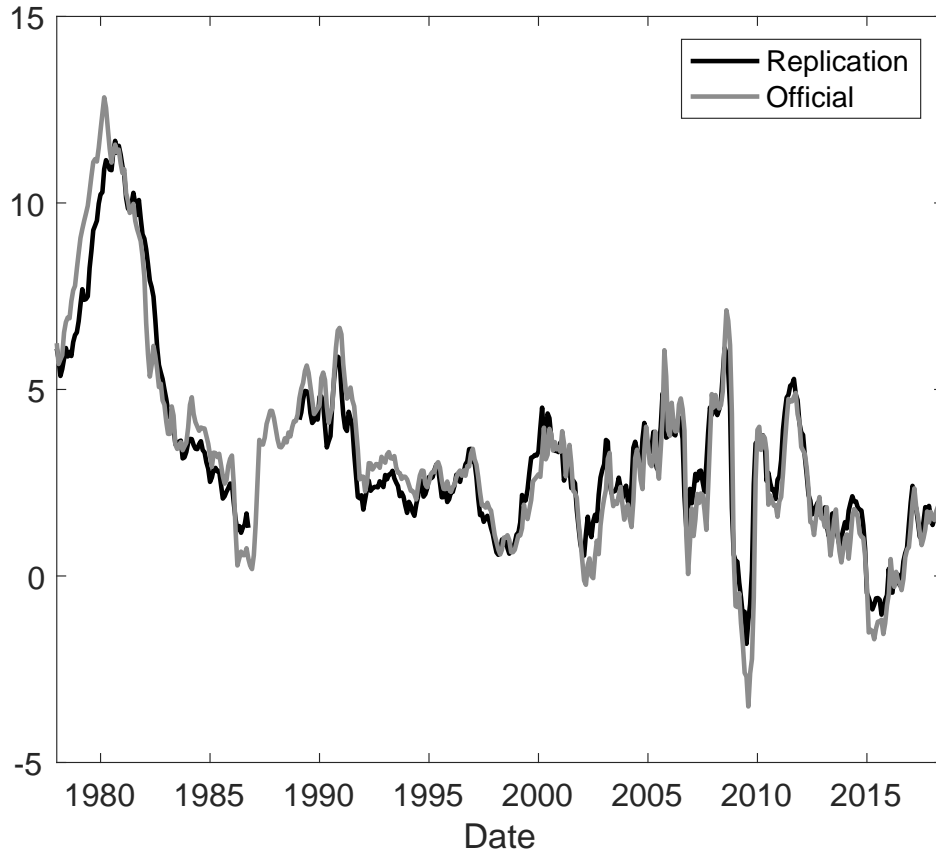


Figure 4: Aggregate Non-Shelter Inflation

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).

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. Finally, 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.¹³

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

¹³Using the X-11 algorithm for seasonal adjustment yields virtually identical results.

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 assign it zero employment in that quarter.¹⁴

5 Empirical Results

We now turn to our empirical results. We focus on estimating a regional Phillips curve for non-tradeables for the reasons discussed in section 3.2. Our empirical specification is a slight variation on equation (20):

$$\pi_{it}^N = \alpha_i + \gamma_t - \psi u_{i,t-4} - \delta p_{i,t-4}^N + \varepsilon_{it}. \quad (23)$$

Here $\pi_{it}^N = p_{it}^N - p_{i,t-4}^N$ is non-tradeable inflation over the previous 12 months, $u_{i,t-4}$ is the unemployment rate lagged by four quarters, and $p_{i,t-4}^N$ is the relative price of non-tradeable goods also lagged by four quarters. Studying inflation over four quarters allows us to reduce measurement error and eliminate seasonality. We use lagged unemployment as a regressor, instead of current unemployment, for consistency with previous studies such as Ball and Mazumder (2019).

Our panel data approach implies that we are relying on cross-state variation in unemployment to identify the slope of the Phillips curve. Figure 5 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.

As we have emphasized throughout the paper, ψ in equation (23) (and equation (3)) is not the same object as κ in equation (1). Making the simplifying assumption that the unemployment rate follows an AR(1), we have derived a simple formula for the relationship between these two concepts: $\psi = \kappa / (1 - \rho_u \beta)$, where ρ_u is the first order autoregressive coefficient for unemployment. 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

¹⁴We did not find any transcription errors in the QCEW at the state level. Chodorow-Reich and Wieland (2019) discuss such errors at the county-industry level for the QCEW. 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 5: Regional Business Cycles

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

use a more direct approach to estimating the appropriate scaling factor for converting our estimate of ψ into an estimate of κ . We estimate the following regression

$$\sum_{j=0}^T \beta^j u_{i,t+j} = \zeta u_{it} + \delta p_{it}^N + \alpha_i + \gamma_t + \epsilon_{it}. \quad (24)$$

We can then combine the estimate of ζ from this regression with our estimate of ψ from equation (23) to get an estimate of κ . In appendix A.4, we show that

$$\kappa = \frac{\psi}{4\zeta}, \quad (25)$$

under relatively mild assumptions about the dynamics of the unemployment rate. The appearance of 4 in the denominator accounts for time aggregation—the outcome variable in regression equation (23) is cumulative inflation over four quarters, but our benchmark model is specified at quarterly frequency. We discuss our choice of the truncation length T below. We calculate

standard errors for κ under the simplifying assumption that ψ is random but ζ is known.

The presence of supply shocks and cost push shocks (u_t^n and η_t in equation (1)) is an important challenge to the identification of the slope of the Phillips curve. We present results for two alternative identifying assumptions regarding supply shocks. Our first approach is to estimate equation (23) by OLS. The identifying assumption here is that when state A experiences a boom or bust relative to other state, it does not systematically experience non-tradeable supply shocks relative to 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.

Our second approach 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 demand for non-tradeable goods in states that produce those tradeable goods. For example, an increase in oil prices will differentially affect the income of people in Texas (and other oil producing states). As a result, Texans will differentially increase their demand for non-tradeables (such as restaurant meals). Building on this idea, we construct a “tradeable demand spillovers” instrument as

$$\text{Tradable Demand}_{i,t} = \sum_x \bar{S}_{x,i} \times \log S_{-i,x,t}, \quad (26)$$

where $\bar{S}_{x,i}$ is the average employment share of industry x in state i over time, and $\log S_{-i,x,t}$ is the national employment share 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 $\log S_{-i,x,t}$ 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.¹⁵

5.1 Main Results

Table 1 presents estimates of ψ , ζ , and κ based on OLS regressions of equations (23) - (25) for our full sample period of 1978-2018. Let’s start by considering the estimates of ψ in Panel A. When we estimate equation (23) without fixed effects, our estimate has “wrong” sign, i.e., higher unem-

¹⁵In a related approach, McLeay and Tenreyro (2019) use identified demand shocks from government spending to estimate the slope of the regional Phillips Curve.

Table 1: Slope of the Regional Phillips Curve: OLS

	No Fixed Effects	No Time Effects	Baseline	No Rel. Price Term
	(1)	(2)	(3)	(4)
<i>Panel A: Estimates of ψ from equation (23)</i>				
ψ	-0.103 (0.044)	0.017 (0.048)	0.112 (0.056)	0.155 (0.066)
<i>Panel B: Estimates of ζ from equation (24)</i>				
ζ	10.06 (0.72)	7.48 (0.75)	8.99 (0.48)	9.02 (0.46)
<i>Panel C: Estimates of κ from equation (25)</i>				
κ	-0.0026 (0.0011)	0.0006 (0.0016)	0.0031 (0.0016)	0.0043 (0.0018)
State Effects		✓	✓	✓
Time Effects			✓	✓
$p_{i,t-4}^N$	✓	✓	✓	

Note: The table presents OLS estimates of ψ , ζ and κ . In Panel A, the outcome variable is cumulative non-tradeable inflation over four quarters, measured in percentage points. The regressor of interest is the fourth lag of quarterly state unemployment, measured in percentage points. In panel B, the outcome variable is the present value of unemployment, truncated at 20 quarters. The regressor of interest is quarterly state unemployment. In panel C, we report estimates of $\kappa = \psi/4\zeta$. In all panels, the sample period is 1978-2018. Standard errors are reported in parentheses. These are two-way clustered by state and time. Controls for each column are reported at the bottom of the table. All regressions are unweighted. The number of observations is 4490 in panel A, 4400 in panel B, and 3916 in panel C.

ployment is associated with higher rather than lower inflation ($\psi = -0.103$). Adding state fixed effects raises the estimate of ψ to 0.017. This implies that states with persistently high unemployment had lower average inflation rates over our sample period. Adding time fixed effects further raises the estimate of ψ to 0.112. This implies that supply shocks that result in “stagflation” were an important driver of inflation at the national level over sample period. Finally, dropping the relative price of non-tradeables control raises our estimate of ψ modestly to 0.155.

Panel B of Table 1 presents estimates of the scaling factor ζ from equation (24). For the baseline case with state and time fixed effects and including the control for the relative price of non-tradeables, our estimate of ζ is 8.99. We truncate the discounted sum of future unemployment rates on the left-hand-side of equation (24) at $T = 20$ quarters. We have explored truncating at other horizons. If we truncate at $T = 10$, we get substantially smaller estimates of ζ . This reflects the fact that the shorter sum does not capture the full dynamic effect of changes in unemployment.

If we, instead, truncate at $T = 30$, we get quite similar values for ζ .¹⁶

In Panel C of Table 1, we then combine our estimates of ψ and ζ using equation (25) to arrive at estimates of κ . Focusing on the baseline case in column (3), we divide our estimate of $\psi = 0.112$ by 4×8.99 to arrive at an estimate for κ of 0.0031. Intuitively, ψ captures the effect not only of current unemployment, but also future unemployment, while κ is an estimate of the effect of current unemployment alone.

Table 2 presents IV estimates of ψ , ζ , and κ for three different versions of our tradeable demand spillovers instrument: the level of tradeable, a three-year difference, and a five-year difference. These IV specifications yield substantially larger estimates of both ψ and κ than OLS. The IV estimates of ψ range from 0.339 to 0.525. The IV estimates of ζ are also larger, ranging from 12.3 to 14.04. Combining these yields estimates of κ ranging from 0.0060 to 0.0105. The fact that our IV estimates of κ are substantially larger than our estimates of κ based on OLS suggests that IV strategy excludes variation due to supply shocks that confounds our OLS estimates. We take the estimate using the three-year difference instrument of 0.0075 as our baseline IV estimate of κ .¹⁷

5.2 Subsample Analysis

We next analyze to what extent the Phillips curve has flattened over our sample period. In particular, we investigate to what extent the Phillips curve was steeper during the period of the Volcker disinflation than in subsequent years. Table 3 presents OLS estimates of ψ and κ for the periods 1978-1990 and 1991-2018. We present these estimates for a specification with time fixed effects and a specification without time fixed effects. Both 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). This specification yields a sharp drop in the estimated values of ψ and κ between the early part of the

¹⁶If we truncate at $T = 40$ or higher values, we also get smaller values of ζ . For these very high values of T , the discounted sum on the left-hand-side of equation (24) starts getting substantially influenced by the subsequent business cycle since business cycles during our sample period last less than a decade. Our interpretation of this is that our sample is too short to reliably estimate equation (24) with values of T as high as 40. We have also explored modelling the unemployment rate as an AR(p) to avoid the truncation above, and we have estimated an impulse response for unemployment using a lag-augmented local projection (Montiel Olea and Plagborg-Møller, 2020). Modelling the unemployment rate as an AR(p) yields results that are highly sensitive to the chosen lag length. In particular, when we estimate such a model with many lags (say 8), our estimates indicate that the unemployment rate is non-stationary. In contrast, the impulse response we estimate using a lag-augmented local projection does not suggest a unit root. These results suggest that AR models are unable to accurately capture the dynamics of the unemployment rate.

¹⁷The level instrument has stronger trends than the differenced instruments and results based on this instrument are more sensitive to small changes in the regression specification. This leads us to favor results based on the differenced instruments.

Table 2: Slope of the Regional Phillips Curve: IV

<i>Instrument:</i>	Level	3Y Diff	5Y Diff
	(1)	(2)	(3)
<i>Panel A: Estimates of ψ from equation (23)</i>			
ψ	0.525 (0.184)	0.369 (0.142)	0.339 (0.124)
1st. Stage F Stat	28.7	30.7	45.5
<i>Panel B: Estimates of ζ from equation (24)</i>			
ζ	12.50 (1.87)	12.30 (2.24)	14.04 (1.50)
1st. Stage F Stat	23.4	37.8	58.7
<i>Panel C: Estimates of κ from equation (25)</i>			
κ	0.0105 (0.0037)	0.0075 (0.0029)	0.0060 (0.0022)

Note: The table presents IV estimates of ψ , ζ and κ . Tradeable demand is defined in equation (26). We instrument for unemployment with several instruments based on tradeable demand. Column (1) instruments with the level of tradeable demand. Column (2) uses the three-year difference of tradeable demand. Column (3) uses the five-year difference. In Panel A, the outcome variable is cumulative non-tradeable inflation over four quarters, measured in percentage points. The regressor is the fourth lag of quarterly state unemployment, measured in percentage points. In panel B, the outcome variable is the present value of unemployment, truncated at 20 quarters. The regressor is quarterly state unemployment. In panel C, we report estimates of $\kappa = \psi/4\zeta$. All specifications include state and time fixed effects and a control for the relative price of non-tradeables. The sample period is 1978-2018. Standard errors are reported in parentheses. These are two-way clustered by state and time. All regressions are unweighted. The number of observations in the first column is 4490 in panel A, 4400 in panel B, and 3916 in panel C. The number of observations is slightly smaller in the second and third columns due to the differencing in the instrument.

sample and the later part of the sample. For the pre-1990 sample, ψ is estimated to be 0.449, while κ is estimated to be 0.0150. In sharp contrast, for the post-1990 sample, ψ is estimated to be 0.009 and κ is estimated to be 0.0003. The difference across samples is roughly a factor of 50. 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 ψ and κ fall only modestly between the early part of the sample and the later part of the sample. For the pre-1990 sample, ψ is estimated to be 0.198 and κ is estimated to be 0.0055. For the post-1990 sample, ψ is estimated to be 0.090 and κ is estimated to be 0.0025. The difference across samples is a little less than a factor of two and is marginally statistically significant.

As we emphasize in section 2, estimates of the Phillips curve based on time-series variation —

Table 3: Has the Phillips Curve Flattened? – OLS

	Without Time Fixed Effect		With Time Fixed Effect	
	Pre-1990 (1)	Post-1990 (2)	Pre-1990 (3)	Post-1990 (4)
<i>Panel A: Estimates of ψ from equation (23)</i>				
ψ	0.449 (0.073)	0.009 (0.043)	0.198 (0.106)	0.090 (0.057)
<i>Panel B: Estimates of κ from equation (25)</i>				
κ	0.0150 (0.0024)	0.0003 (0.0014)	0.0055 (0.0029)	0.0025 (0.0016)

Note: The table presents OLS estimates of ψ and κ for subsamples. Columns (1) and (3) present results for the sample period 1978-1990, and columns (2) and (4) for the sample period 1991-2018. The specifications in columns (1)-(2) include state fixed effects and control for the relative price of non-tradeables. The specifications in columns (3)-(4) include time fixed effects as well. In Panel A, the outcome variable is cumulative non-tradeable inflation over four quarters, measured in percentage points. The regressor is the fourth lag of quarterly state unemployment, measured in percentage points. Standard errors are reported in parentheses. These are two-way clustered by state and time. In panel B, we report estimates of $\kappa = \psi/4\zeta$. To estimate κ , we fix the value of ζ while ψ varies by subsample. Columns (1)-(2) use the estimate of ζ from column (2) of Table 1, estimated over the full sample 1978-2018. Columns (3)-(4) use the estimate of ζ from column (3) of Table 1, again estimated over the full sample. All regressions are unweighted.

such as the estimates without time fixed effects in Table 3 — are likely to be heavily influenced by time-series variation in long-run inflation expectations $E_t\pi_{t+\infty}$. In contrast, the specifications in Table 3 that include time fixed effects difference out the influence of long-run inflation expectations. The results in Table 3 therefore suggest that the apparent flattening of the Phillips curve in the time series is really due to inflationary expectations becoming more firmly anchored over time. In the early part of the sample, inflationary expectations shifted a great deal and these shifts were negatively correlated with the unemployment rate, which meant that shifts in inflationary expectations masqueraded as a steep Phillips curve. The cross-sections results in columns (3) - (4) of Table 3 reveal that in fact the Phillips curve has always been quite flat (or at least since 1978).

Figure 6 provides a visual representation of the results in Table 3. 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 (but not time fixed effects). 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.




Figure 6: 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.

Contrast this with the right panel in Figure 6. 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.

Tables 4 reports an analogous set of results to the results in columns (3) - (4) of Table 3 but based on IV regressions using the tradeable demand instruments discussed above. As with the full sample results, using our IV strategy over subsamples yield larger estimates of ψ and κ than does OLS. However, the time pattern of the results are quite similar to those in Table 3. When we use the of three-year difference in tradeable demand as our instrument, the estimate of κ for the period prior to 1990 is 0.0106. This estimate falls to 0.0059 for the post-1990 sample. As in the case of OLS, the flattening of the Phillips curve is a factor of about 1.8. The results when the five-year difference in tradeable demand is used as an instrument are similar. The modest flattening of the Phillips curve that we find over our sample 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).

Table 4: Has the Phillips Curve Flattened? – IV

<i>Instrument:</i>	3Y Diff		5Y Diff	
	Pre-1990 (1)	Post-1990 (2)	Pre-1990 (3)	Post-1990 (4)
<i>Panel A: Estimates of ψ from equation (23)</i>				
ψ	0.522 (0.212)	0.292 (0.185)	0.422 (0.206)	0.332 (0.155)
1st. Stage F Stat	14.1	23.9	7	39
<i>Panel B: Estimates of κ from equation (25)</i>				
κ	0.0106 (0.0043)	0.0059 (0.0038)	0.0075 (0.0037)	0.0059 (0.0028)

Notes: the table presents IV estimates of ψ and κ by subsample. Columns (1) and (3) present results for the sample period 1978-1990, columns (2) and (4) for the sample period 1991-2018. We instrument for unemployment with several instruments based on tradeable demand. Columns (1)-(2) instruments with the three year difference of tradeable demand, lagged by four quarters. Columns (3)-(4) uses the five-year difference of tradeable demand. All columns include state and time fixed effects and control for the relative price of non-tradeables. In Panel A, the outcome variable is cumulative non-tradeable inflation over four quarters, measured in percentage points. The regressor of interest is the fourth lag of quarterly state unemployment, measured in percentage points. Standard errors are reported in parentheses. These are two-way clustered by state and time. In panel B, we report estimates of $\kappa = \psi/4\zeta$. To estimate κ , we fix the value of ζ while ψ varies by subsample. Columns (1)-(2) use the estimate of ζ from column (2) of Table 2, estimated over the full sample 1978-2018. Columns (3)-(4) use the estimate of ζ from column (3) of Table 2, again estimated over the full sample. All regressions are unweighted.

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 5 reports three such estimates from Rotemberg and Woodford (1997), Galí (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 rigidities assumed and the degree of real rigidities assumed. The values for κ used in these papers ranged from about an order of magnitude larger than our estimated value to a value roughly equal to our estimated value. The main difference between Galí'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. Galí'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

Table 5: 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	0.0075

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.

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 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). Our first step is to estimate a scaling factor ζ for the aggregate unemployment rate analogous to the scaling factor we use in our cross-sectional analysis—equations (24) and (25). We do this using the following specification

$$\sum_{j=0}^T \beta^j \tilde{u}_{t+j} = \zeta \tilde{u}_t + \alpha + \epsilon_t. \quad (27)$$

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. We run this regression for the sample period 1979Q4-2017Q4. This yields an estimate of ζ for

Table 6: Slope of the Regional Phillips Curve: Rents

	No Fixed Effects	No Time Effects	Baseline
	(1)	(2)	(3)
<i>Panel A: Estimates of ψ from equation (23)</i>			
ψ	0.268 (0.041)	0.356 (0.044)	0.603 (0.124)
<i>Panel B: Estimates of ζ from equation (24)</i>			
ζ	10.06 (0.72)	7.46 (0.76)	9.03 (0.46)
<i>Panel C: Estimates of κ from equation (25)</i>			
κ	0.0067 (0.0010)	0.0119 (0.0015)	0.0167 (0.0034)
State Effects		✓	✓
Time Effects			✓

Note: The table presents OLS estimates of ψ , ζ and κ for rents. In Panel A, the dependent variable is the state-level annual rent inflation rate, measured in percentage points from the American Community Survey for the years 2001 to 2017. The regressor of interest is the lag of annual state unemployment, measured in percentage points. In panel B, the outcome variable is the present value of unemployment, truncated at 20 quarters. The regressor of interest is quarterly state unemployment. In panel C, we report estimates of $\kappa = \psi/4\zeta$. Standard errors are reported in parentheses. These are two-way clustered by state and time. Controls for each column are reported at the bottom of the table. All regressions are unweighted. The number of observations is 850 in panel A, 4400 in panel B, and 850 in panel C.

aggregate variation in the unemployment rate of 6.16 with a standard error of 0.87.

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. Table 6 presents estimates of ψ and κ 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 0.0167. This estimate of κ is roughly two and a half times larger than our post-1990 IV estimate of κ for non-housing non-tradeable goods reported in Table 4 and about twice as large as our full-sample IV estimate reported in Table 2 (using the three-year difference instrument). We account for this difference below by taking a weighted average of our κ estimate for non-tradeables and this κ estimate for housing services.

Using equation (27), we can rewrite equation (6) as

$$\pi_t - E_t \pi_{t+\infty} = -\kappa \zeta \tilde{u}_t + \omega_t. \quad (28)$$

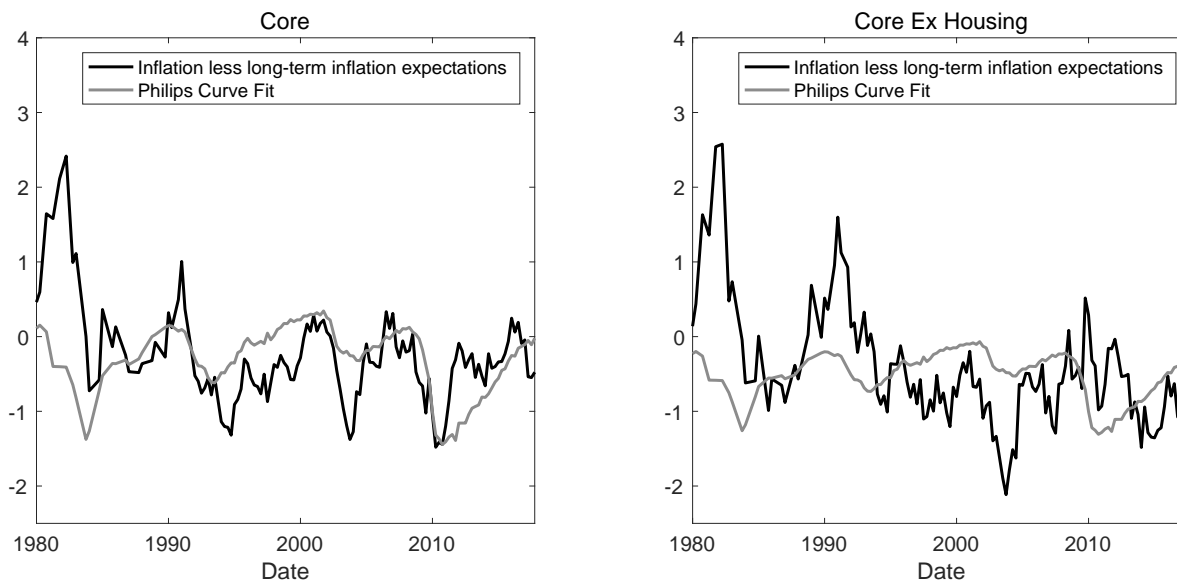


Figure 7: Aggregate Phillips Curve and Housing: Predicted vs. Fit

Note: The figure shows the fit of the aggregate Phillips curve for core inflation (left panel) and core inflation excluding shelter (right panel). The fit curves use the estimates for κ from the IV specifications, and convert them to estimates of ψ using the persistence of cyclical unemployment. For the case of core inflation including housing, we compute a weighted average of ψ using our estimates for non-tradeable goods excluding housing, and the coefficient from rents from Table 6.

To assess whether our cross-sectional estimates of κ implies a missing disinflation during the Great Recession and missing reflation during the late 2010s and the later 1990s, Figure 7 plots the left-hand side of equation (28) against the first term on the right-hand side of equation (28) using our full-sample estimates of κ and ζ from above. The left panel in Figure 7 does this for core inflation, while the right panel does it for core inflation excluding shelter. For both indices, we use the 10 year ahead inflation expectation for the CPI gathered from the Survey of Professional Forecasters as our measure of long-term inflation expectations.

Consider first the left panel in Figure 7. The black line is core CPI inflation using the BLS's research series less 10-year inflation expectations from the Survey of Professional Forecasters. The gray line is the demand-induced variation in inflation predicted by our estimates. Specifically, we multiply the unemployment gap — U.S. unemployment rate less the CBO's estimate of the natural rate of unemployment — by our aggregate estimate of ζ (6.16) times a weighted average of our IV estimate of κ for non-shelter (0.0081) and our estimate of κ for shelter (0.0167).¹⁸ 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. The disinflation during the Great Recession and reflation during the 2010s was about 0.5 percentage point smaller than

¹⁸We use the shelter and non-shelter expenditure weights in the core CPI. These are 0.42 and 0.58, respectively.

our κ would imply. The rise in inflation during the second half of the 1990s was very comparable to what our estimate implies. The fall in inflation during the 1991 and 2001 recessions somewhat exceed what our estimates imply. Overall, the fitted value explains inflation dynamics for the post-1985 period fairly well, even though the slope of the Phillips curve is estimated using the full sample. By this metric, there is very little missing disinflation or missing reflation. These findings echo the results of Ball and Mazumder (2019).

A more substantial deviation arises between the actual and fitted values is for the pre-1985 period, when actual inflation 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. The reason for this finding is straightforward. Figure 3 shows clearly that core inflation was *above* not below long-term inflation expectations during the Volcker disinflation, despite the high unemployment during this period. A natural interpretation of this phenomenon is the presence of adverse supply shocks, for example, associated with the oil price shocks associated with the Iranian Revolution and the Iran-Iraq War.¹⁹ More generally, the deviations between the black line and the gray line in Figure 3 indicate that supply shocks played a non-trivial role in inflation dynamics throughout our sample period.

It is important to recognize, however, that a disproportionate share of the systematic variation in inflation and the fitted value predicted by our model comes from the housing services (rent) component of the index. The right panel of Figure 7 is analogous to the left panel except that the inflation series is core inflation excluding housing services and the κ used in the construction of the grey line does not put weight on our estimated κ for housing services. We see that core inflation excluding housing services varies much less systematically than core inflation including housing services.

Finally, we consider the quantitative significance of the flattening of the Phillips curve that we estimate. Figure 8 compares the right-hand side of equation (28) using our pre-1990 κ estimate and using our post-1990 κ estimate. In both cases, we take a weighted average of our non-shelter κ and our κ for shelter. Since we do not have a pre-1990 κ for shelter, we rescale our post-1990 κ for shelter assuming the same flattening occurred for shelter as non-shelter. These calculations yield an overall pre-1990 $\kappa = 0.01867$ and an overall post-1990 $\kappa = 0.01039$. Applying our ζ for

¹⁹An alternative explanation is some form of intrinsic inflation persistence arising other sources than inflation expectations.




Figure 8: How Much Flatter Can a Flat Phillips Curve Get?

Note: In this figure we compute the variation in inflation caused by the $-\kappa\zeta\tilde{u}_t$ term in the Phillips curve using our estimates for κ and ζ . We report to cases, using our pre-1990 and post-1990 estimates of κ . We combine estimates of κ for the non-shelter component of goods and services from Table 4 with our estimate of κ for shelter. We translate our estimates for κ into estimates of ψ at the aggregate level using the persistence of the the aggregate unemployment gap. See the text for more detail.

aggregate variation in the unemployment gap (6.16), we get a pre-1990 $\psi = 0.4605$ and a post-1990 $\psi = 0.2563$. Using these values, Figure 8 shows that the roughly five percentage point increase in the unemployment gap in the early 1980s causes a roughly 2.3 percentage point reduction in inflation over time. Had the Phillips curve instead been as flat as we estimate in for the post-1990 period, this same increase in unemployment would have led to a 1.3 percentage point reduction in inflation. The message of the figure is that a very flat Phillips curve can't get much flatter in levels even if it flattens substantially in proportional terms.

6 Conclusion

This paper provides new estimates of the slope of the Phillips curve? We estimate that the Phillips curve is very flat, and was very flat 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 governed 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 Theory Appendix

A.1 Model Derivation

This subsection presents the complete set of equations characterizing the log-linear dynamics of the model, and derives them.

A.1.1 Log-Linearized Equations of the Model

- Parameters:

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

$$- \kappa = \lambda \varphi^{-1}$$

$$- \lambda = (1 - \alpha) (1 - \alpha \beta) / \alpha$$

$$- \mu = \theta / (\theta - 1)$$

$$- \tau_H^H = \tau_F^H = \zeta$$

- The law of motion for tradeable demand is as follows. Define $\hat{\xi}_{Ht} = \log \tau_{Ht} - \log \tau_H$ and $\hat{\xi}_{Ft} = \log \tau_{Ft} - \log \tau_F$. Then

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

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:

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

and $\hat{p}_{Ht}^T \equiv P_{Ht}^T / P_{Ht} - 1$ is the percentage deviation of the price of home produced tradeables, relative to the home consumer price level, from its steady state value of 1.

- The home Euler equation is:

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

- The Backus-Smith condition is:

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

- The foreign non-tradeable Phillips Curve is:

$$\pi_{Ft}^N = \beta E_t \pi_{F,t+1}^N + \kappa \hat{n}_{Ft} - \lambda \hat{p}_{Ft}^N + \nu_{Ft}^N$$

- The foreign tradeable Phillips Curve is:

$$\pi_{Ft}^T = \beta E_t \pi_{F,t+1}^T + \kappa \hat{n}_{Ft} - \lambda \hat{p}_{Ft}^T + \nu_{Ft}^T$$

- Definitions of inflation:

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

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

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

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

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

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

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

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

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

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

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

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

- The home resource constraint in the tradeable sector is:

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

- The foreign resource constraint in the tradeable sector is:

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

- Define total labor in the non-tradeable and tradeable sectors, in the home region, 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. 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

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

- 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

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

We simulate the model in Dynare.

A.1.2 Regional Phillips Curve Derivation

One can rewrite equation (15) around the zero inflation and balanced trade steady state as

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

where $MC_{H,t+k}^N(z) = S_{H,t+k}^N(z)/P_{H,t+k}^N$ is real marginal cost in the non-tradeable sector. A first order expansion of equation (29) around the zero inflation and balanced trade steady state yields

$$p_{Ht}^{N*}(z) - p_{H,t-1}^N = (1 - \alpha\beta) \sum_{k=0}^{\infty} (\alpha\beta)^k E_t \left[\widehat{mc}_{H,t+k}^N(z) - (p_{H,t+k}^N - p_{H,t-1}^N) \right], \quad (30)$$

where $mc_H^N = -\mu$ and $\mu = \log(\theta/(\theta-1))$. Rearranging equation (30) yields

$$p_{Ht}^{N*}(z) - p_{H,t-1}^N = \alpha\beta E_t [p_{H,t+1}^{N*}(z) - p_{Ht}^N] + (1 - \alpha\beta) \widehat{mc}_{Ht}^N + \pi_{Ht}^N. \quad (31)$$

Then note that to a first order approximation around the zero inflation and balanced trade steady state, we have

$$\pi_{Ht}^N = (1 - \alpha) (p_{Ht}^{N*}(z) - p_{H,t-1}^N). \quad (32)$$

Substituting equations (31) and (32) yields

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

where

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

By an analogous series of steps, we have

$$\pi_{Ht}^T = \beta E_t \pi_{H,t+1}^T + \lambda \widehat{mc}_{Ht}^T, \quad (33)$$

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

Then from equation (14) we have

$$mc_{Ht}^N = w_{Ht} - p_{Ht}^N - z_{Ht}^N \quad (34)$$

$$= (w_{Ht} - p_{Ht}) + (p_{Ht} - p_{Ht}^N) - z_{Ht}^N \quad (35)$$

$$= \varphi^{-1}n_{Ht} + (p_{Ht} - p_{Ht}^N) - z_{Ht}^N \quad (36)$$

where in the third line we use the labor supply curve (9) and the definition of home CPI (12).

Therefore

$$\begin{aligned} \pi_{Ht}^N &= \beta E_t \pi_{H,t+1}^N + \lambda \widehat{mc}_{Ht}^N \\ &= \beta E_t \pi_{H,t+1}^N + \lambda (\varphi^{-1} \hat{n}_{Ht} - \hat{p}_{Ht}^N - z_{Ht}^N) \\ &= \beta E_t \pi_{H,t+1}^N + \lambda \varphi^{-1} \hat{n}_{Ht} - \lambda \hat{p}_{Ht}^N - \lambda z_{Ht}^N \\ &= \beta E_t \pi_{H,t+1}^N + \kappa \hat{n}_{Ht} - \lambda \hat{p}_{Ht}^N + \nu_{Ht}^N \end{aligned} \quad (37)$$

where $\nu_{Ht}^N = -\lambda z_{Ht}^N$, $\kappa = \lambda \varphi^{-1}$. Equation (37) is the regional non-tradeable Phillips Curve.

Next, from equation

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

and

$$\begin{aligned} mc_{Ht}^T &= w_{Ht} - p_{Ht}^T - z_{Ht}^T \\ &= (w_{Ht} - p_{Ht}) + (p_{Ht} - p_{Ht}^T) - z_{Ht}^T \\ &= \varphi^{-1}n_{Ht} + (p_{Ht} - p_{Ht}^T) - z_{Ht}^T \end{aligned}$$

again using the labor supply curve (9). Therefore

$$\begin{aligned} \pi_{Ht}^T &= \beta E_t \pi_{H,t+1}^T + \lambda \widehat{mc}_{Ht}^T \\ &= \beta E_t \pi_{H,t+1}^T + \lambda (\varphi^{-1} \hat{n}_{Ht} - \hat{p}_{Ht}^T - z_{Ht}^T) \\ &= \beta E_t \pi_{H,t+1}^T + \kappa \hat{n}_{Ht} - \lambda \hat{p}_{Ht}^T + \nu_{Ht}^T \end{aligned}$$

where $\nu_{Ht}^T = -\lambda z_{Ht}^T$.

A.1.3 Aggregate Phillips Curve Derivation

Therefore aggregate non-tradeable inflation $\pi_t^N = \zeta \pi_{Ht}^N + (1 - \zeta) \pi_{Ft}^N$ satisfies

$$\begin{aligned}\pi_t^N &= \zeta (\beta E_t \pi_{H,t+1}^N + \kappa \hat{n}_{Ht} - \lambda \hat{p}_{Ht}^N + \nu_{Ht}^N) \\ &\quad + (1 - \zeta) (\beta E_t \pi_{F,t+1}^N + \kappa \hat{n}_{Ft} - \lambda \hat{p}_{Ft}^N + \nu_{Ft}^N) \\ &= \beta E_t \pi_{t+1}^N + \kappa \hat{n}_t + \nu_t^N - \lambda [\zeta \hat{p}_{Ht}^N + (1 - \zeta) \hat{p}_{Ft}^N]\end{aligned}$$

and aggregate tradeable inflation satisfies

$$\begin{aligned}\pi_t^T &= \zeta (\beta E_t \pi_{H,t+1}^T + \kappa \hat{n}_{Ht} - \lambda \hat{p}_{Ht}^T + \nu_{Ht}^T) \\ &\quad + (1 - \zeta) (\beta E_t \pi_{F,t+1}^T + \kappa \hat{n}_{Ft} - \lambda \hat{p}_{Ft}^T + \nu_{Ft}^T) \\ &= \beta E_t \pi_{t+1}^T + \kappa \hat{n}_t + \nu_t^T - \lambda [\zeta \hat{p}_{Ht}^T + (1 - \zeta) \hat{p}_{Ft}^T].\end{aligned}$$

Then aggregate inflation satisfies

$$\begin{aligned}\pi_t &= \phi_N \pi_t^N + \phi_T \pi_t^T \\ &= \phi_N (\beta E_t \pi_{t+1}^N + \kappa \hat{n}_t + \nu_t^N - \lambda [\zeta \hat{p}_{Ht}^N + (1 - \zeta) \hat{p}_{Ft}^N]) \\ &\quad + \phi_T (\beta E_t \pi_{t+1}^T + \kappa \hat{n}_t + \nu_t^T - \lambda [\zeta \hat{p}_{Ht}^T + (1 - \zeta) \hat{p}_{Ft}^T]) \\ &= \beta E_t \pi_{t+1} + \kappa \hat{n}_t + \nu_t \\ &\quad - \lambda ([\zeta \hat{p}_{Ht}^N + (1 - \zeta) \hat{p}_{Ft}^N] + [\zeta \hat{p}_{Ht}^T + (1 - \zeta) \hat{p}_{Ft}^T])\end{aligned}$$

where $\nu_t \equiv \phi_N \nu_t^N + \phi_T \nu_t^T$. Then it is easy to verify that

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

so aggregate inflation satisfies

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

A.1.4 Steady State Equations

At the steady state the real wage is

$$\frac{W}{P} = \frac{\theta - 1}{\theta}.$$

Note that, although the price level is not stationary and does not have a well defined value at the steady state, the real wage is stationary, so we omit time subscripts on the real wage.

Then

$$C_H^N = \phi_N C_H$$

$$C_H^{TH} = \phi_T \tau_H^H C_H$$

$$C_H^{TF} = \phi_T \tau_H^F C_H$$

and

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

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

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

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

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

$$\implies N_H^T = \zeta \phi_T C,$$

using $C_F = C_H = C$.

A.1.5 Deriving the Other Log-Linearized Equations

From now on, we set supply shocks to zero: the only shock is to tradeable demand for home production. The $dll[\cdot]$ operator denotes taking the log deviation from the steady state.

The home Euler equation is

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

Then log-linearising yields

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

$$\begin{aligned}
&\Rightarrow E_t \left[\left[\frac{u_{cc}C}{u_c} \hat{c}_{H,t+1} + \frac{u_{cn}N}{u_c} \hat{n}_{H,t+1} \right] - \left[\frac{u_{cc}C}{u_c} \hat{c}_{Ht} + \frac{u_{cn}N}{u_c} \hat{n}_{Ht} \right] - \pi_{H,t+1} \right] = -\hat{r}_t^n \\
&\Rightarrow \frac{u_{cc}C}{u_c} \hat{c}_{Ht} + \frac{u_{cn}N}{u_c} \hat{n}_{Ht} = E_t \left[\frac{u_{cc}C}{u_c} \hat{c}_{H,t+1} + \frac{u_{cn}N}{u_c} \hat{n}_{H,t+1} \right] + \hat{r}_t^n - E_t \pi_{H,t+1} \\
&\Rightarrow \hat{c}_{Ht} + \frac{u_{cn}N}{u_{cc}C} \hat{n}_{Ht} = E_t \left[\hat{c}_{H,t+1} + \frac{u_{cn}N}{u_{cc}C} \hat{n}_{H,t+1} \right] + \frac{u_c}{u_{cc}C} (\hat{r}_t^n - E_t \pi_{H,t+1}) \\
&\Rightarrow \hat{c}_{Ht} + \frac{u_{cn}}{u_{cc}} \hat{n}_{Ht} = E_t \left[\hat{c}_{H,t+1} + \frac{u_{cn}}{u_{cc}} \hat{n}_{H,t+1} \right] + \frac{u_c}{u_{cc}C} (\hat{r}_t^n - E_t \pi_{H,t+1}).
\end{aligned}$$

Next, we have

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

where we use the steady state condition from labor supply that

$$\chi N^{\varphi^{-1}} = \mu^{-1} \quad \mu = \frac{\theta}{\theta-1}.$$

Then

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

Therefore

$$\begin{aligned}
\hat{c}_{Ht} - \frac{u_{cc}\mu^{-1}}{u_{cc}}\hat{n}_{Ht} &= E_t \left[\hat{c}_{H,t+1} - \frac{u_{cc}\mu^{-1}}{u_{cc}}\hat{n}_{H,t+1} \right] + \left(\frac{u_{cc}C}{u_c} \right)^{-1} (\hat{r}_t^n - E_t\pi_{H,t+1}) \\
\implies \hat{c}_{Ht} - \mu^{-1}\hat{n}_{Ht} &= E_t [\hat{c}_{H,t+1} - \mu^{-1}\hat{n}_{H,t+1}] - \sigma \left(1 - \mu^{-1} (1 + \varphi^{-1})^{-1} \right) (\hat{r}_t^n - E_t\pi_{H,t+1}) \\
\implies \hat{c}_{Ht} - \mu^{-1}\hat{n}_{Ht} &= E_t [\hat{c}_{H,t+1} - \mu^{-1}\hat{n}_{H,t+1}] - \sigma_c (\hat{r}_t^n - E_t\pi_{H,t+1})
\end{aligned}$$

where $\sigma_c = \sigma \left(1 - \mu^{-1} (1 + \varphi^{-1})^{-1} \right)$. Solving forward the Euler equation yields

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

Similarly, for foreign households we have

$$\begin{aligned}
\hat{c}_{Ft} - \mu^{-1}\hat{n}_{Ft} &= -\sigma_c E_t \sum_{j=0}^{\infty} \hat{r}_{t+j}^n - \sigma_c p_{Ft} \\
\implies \hat{c}_{Ht} - \mu^{-1}\hat{n}_{Ht} + \sigma_c p_{Ht} &= \hat{c}_{Ft} - \mu^{-1}\hat{n}_{Ft} + \sigma_c p_{Ft} \\
\implies \hat{c}_{Ht} - \mu^{-1}\hat{n}_{Ht} &= \hat{c}_{Ft} - \mu^{-1}\hat{n}_{Ft} + \sigma_c (p_{Ft} - p_{Ht}),
\end{aligned}$$

which is the Backus-Smith condition.

Next, we have the consumer demand functions and prices

$$\begin{aligned}
\hat{c}_{Ht}^N &= \hat{c}_{Ht} - \eta (p_{Ht}^N - p_{Ht}) \\
\hat{c}_{Ht}^T &= \hat{\xi}_{Ht} + \hat{c}_{Ht} - \eta (p_{Ht}^T - p_{Ht}^T) \\
\hat{c}_{Ft}^T &= \hat{\xi}_{Ft} + \hat{c}_{Ht} - \eta (p_{Ft}^T - p_{Ht}^T)
\end{aligned}$$

and

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

$$\hat{c}_{Ft}^{TH} = \hat{c}_{Ft} - \eta (p_{Ht}^T - p_{Ft})$$

$$\hat{c}_{Ft}^{TF} = \hat{c}_{Ft} - \eta (p_{Ft}^T - p_{Ft}).$$

Note that

$$\tau_{Ht} + \tau_{Ft} = 1$$

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

$$\implies \hat{\xi}_{Ft} = -\frac{\tau_H^H}{\tau_H^F} \hat{\xi}_{Ht}$$

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

Also

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

and similarly

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

Note that without supply shocks

$$Y_{Ht}^N = N_{Ht}^N$$

$$Y_{Ht}^T = N_{Ht}^T$$

$$Y_{Ft}^N = N_{Ft}^N$$

$$Y_{Ft}^T = N_{Ft}^T.$$

Then we have

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

that is, total labor supplied by households in the home region equals total labor demanded by firms (the ζ term reflects that N_{Ht} is household labor supply). So,

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

and

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

Aggregate employment is

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

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

where N is steady state household labor supply, equal across the two regions at the symmetric steady state. We have market clearing conditions in the non-tradeable sector

$$N_{Ht}^N = \zeta C_{Ht}^N$$

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

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

and

$$N_{Ft}^N = \zeta C_{Ft}^N$$

$$\implies \hat{n}_{Ft}^N = \hat{c}_{Ft}^N$$

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

and in the tradeable sector we have

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

and so

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

Similarly

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

We also have

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

A.2 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 expenditure 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:

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

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

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

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

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

yields

$$\pi_{Ht} = \beta E_t \pi_{H,t+1} - (\phi_N + \phi_T \tau_H^H) \kappa \hat{u}_{Ht} - \lambda (\phi_N \hat{p}_{Ht}^N + \phi_T \tau_H^H \hat{p}_{Ht}^T) - \phi_T \tau_H^F \kappa \hat{u}_{Ft} - \lambda \phi_T \tau_H^F \hat{p}_{Ft}^T.$$

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

$$\pi_{Ft} = \beta E_t \pi_{F,t+1} - (\phi_N + \phi_T \tau_F^F) \kappa \hat{u}_{Ft} - \lambda (\phi_N \hat{p}_{Ft}^N + \phi_T \tau_F^F \hat{p}_{Ft}^T) - \phi_T \tau_F^H \kappa \hat{u}_{Ht} - \lambda \phi_T \tau_F^H \hat{p}_{Ht}^T.$$

Subtracting the second of these last two equations from the first (and using the fact that $\tau_H^H =$

$\tau_F^H = \zeta$) yields

$$\pi_{Ht} - \pi_{Ft} = \beta (E_t \pi_{H,t+1} - E_t \pi_{F,t+1}) - \phi_N \kappa (\hat{u}_{Ht} - \hat{u}_{Ft}) - \phi_N \lambda (\hat{p}_{Ht}^N - \hat{p}_{Ft}^N). \quad (38)$$

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.3 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:

$$\pi_{Ht}^N = \beta E_t \pi_{H,t+1}^N - \kappa \hat{u}_{Ht} + \lambda \sigma^{-1} \hat{c}_{Ht} - \lambda \hat{p}_{Ht}^N + \nu_{Ht}^N, \quad (39)$$

and the aggregate Phillips Curve under separable preferences:

$$\pi_t = \beta E_t \pi_{t+1} - \kappa \hat{u}_t + \lambda \sigma^{-1} \hat{c}_t + \nu_t. \quad (40)$$

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 (40) 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.4 Time Aggregation

Here, we derive equation (25). In particular, we show how time aggregation results in the factor of 4 showing up in the denominator of that equation. Recall that our empirical specification involves cumulative inflation over four quarters, while our model is written in terms of quarterly inflation.

Consider the non-tradeable regional Phillips Curve—equation (18) from the main text:

$$\pi_{Ht}^N = -E_t \sum_{j=0}^{\infty} \beta^j (\kappa \tilde{u}_{H,t+j} + \lambda \hat{p}_{H,t+j}^N) + E_t \pi_{t+\infty}, \quad (41)$$

where for simplicity we have set the supply shock ω_{Ht}^N equal to zero. We start by approximating \hat{p}_{Ht}^N by a driftless AR(1) with autoregressive coefficient $\rho_{\hat{p}_H^N}$. This yields

$$\pi_{Ht}^N = -\kappa E_t \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} - \frac{\lambda}{1 - \beta \rho_{\hat{p}_H^N}} \hat{p}_{Ht}^N + E_t \pi_{t+\infty}$$

or equivalently

$$p_{Ht}^N - p_{H,t-1}^N = -\kappa E_t \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} - \frac{\lambda}{1 - \beta \rho_{\hat{p}_H^N}} \hat{p}_{Ht}^N + E_t \pi_{t+\infty}.$$

We can write this equation out for four consecutive periods

$$\begin{aligned} p_{Ht}^N - p_{H,t-1}^N &= -\kappa E_t \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} - \frac{\lambda}{1 - \beta \rho_{\hat{p}_H^N}} \hat{p}_{Ht}^N + E_t \pi_{t+\infty} \\ p_{H,t-1}^N - p_{H,t-2}^N &= -\kappa E_{t-1} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-1} - \frac{\lambda}{1 - \beta \rho_{\hat{p}_H^N}} \hat{p}_{H,t-1}^N + E_{t-1} \pi_{t+\infty} \\ p_{H,t-2}^N - p_{H,t-3}^N &= -\kappa E_{t-2} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-2} - \frac{\lambda}{1 - \beta \rho_{\hat{p}_H^N}} \hat{p}_{H,t-2}^N + E_{t-2} \pi_{t+\infty} \\ p_{H,t-3}^N - p_{H,t-4}^N &= -\kappa E_{t-3} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-3} - \frac{\lambda}{1 - \beta \rho_{\hat{p}_H^N}} \hat{p}_{H,t-3}^N + E_{t-3} \pi_{t+\infty} \end{aligned}$$

Summing the preceding four equations together yields

$$\begin{aligned}
p_{Ht}^N - p_{H,t-4}^N &= -\kappa(E_t \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} + E_{t-1} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-1} \\
&\quad + E_{t-2} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-2} + E_{t-3} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-3}) \\
&\quad - \frac{\lambda}{1 - \beta\rho_{\hat{p}^N}} (\hat{p}_{Ht}^N - \hat{p}_{H,t-1}^N - \hat{p}_{H,t-2}^N - \hat{p}_{H,t-3}^N) \\
&\quad + E_t \pi_{t+\infty} + E_{t-1} \pi_{t+\infty} + E_{t-2} \pi_{t+\infty} + E_{t-3} \pi_{t+\infty}.
\end{aligned}$$

Taking expectations at time $t - 4$ then yields

$$\begin{aligned}
E_{t-4} p_{Ht}^N - p_{H,t-4}^N &= -\kappa(E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} + E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-1} \\
&\quad + E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-2} + E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-3}) \\
&\quad - \frac{\lambda}{1 - \beta\rho_{\hat{p}^N}} E_{t-4} [\hat{p}_{Ht}^N - \hat{p}_{H,t-1}^N - \hat{p}_{H,t-2}^N - \hat{p}_{H,t-3}^N] \\
&\quad + 4E_{t-4} \pi_{t+\infty}.
\end{aligned}$$

Adding and subtracting p_{Ht}^N yields

$$\begin{aligned}
p_{Ht}^N - p_{H,t-4}^N &= -\kappa(E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} + E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-1} \\
&\quad + E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-2} + E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-3}) \\
&\quad - \frac{\lambda}{1 - \beta\rho_{\hat{p}^N}} (\rho_{\hat{p}^N}^4 + \rho_{\hat{p}^N}^3 + \rho_{\hat{p}^N}^2 + \rho_{\hat{p}^N}) \hat{p}_{H,t-4}^N \\
&\quad + 4E_{t-4} \pi_{t+\infty} - (E_{t-4} p_{Ht}^N - p_{Ht}^N). \tag{42}
\end{aligned}$$

Comparing equation (42) to equation (23), both equations have cumulative inflation over four quarters on the left hand side. We can see that:

1. The third line of equation (42) corresponds to the $\hat{p}_{H,t-4}^N$ term in equation (23).
2. The first term in the fourth line of equation (42) corresponds to the time fixed effect in equation (23)

3. The second term in the fourth line of equation (42) is an expectational error, uncorrelated with $u_{H,t-4}$ by the law of iterated expectations.

This implies that ψ in equation (23) is given by

$$\begin{aligned} \psi = \kappa \text{Cov} & \left(E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} + E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-1} \right. \\ & \left. + E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-2} + E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-3}, \tilde{u}_{H,t-4} \right), \end{aligned} \quad (43)$$

where the covariance is conditional on the variables in the third and fourth lines of equation (42).

To simplify this expression, we make the following simplifying assumption:

$$E_{t-4} \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j} = \sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-4}.$$

This approximation is accurate if $\sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j}$ has similar dynamics to a driftless random walk. We present empirical evidence supporting this in section A.4.1 below. With this approximation, equation (43) simplifies to

$$\psi = 4\kappa \text{Cov} \left[\sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-4}, \tilde{u}_{H,t-4} \right].$$

Finally, we make the approximation that

$$\sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j-4} = \sum_{j=0}^T \beta^j \tilde{u}_{H,t+j-4}$$

for a sufficiently large T . In this case, we have that

$$\kappa = \frac{\psi}{4\zeta},$$

since ζ is the coefficient from a regression of $\sum_{j=0}^T \beta^j \tilde{u}_{H,t+j-4}$ on $\tilde{u}_{H,t-4}$.

A.4.1 Studying the Dynamics of the Present Value of Unemployment

The relationship between κ , ζ , and ψ derived above made the simplifying assumption that $\sum_{j=0}^{\infty} \beta^j \tilde{u}_{H,t+j}$ follows a driftless random walk. We can assess the accuracy of this assumption

Table A.1: Is the Present Value of State Unemployment a Random Walk?

	No Fixed Effects	No Time Effects	Baseline
	(1)	(2)	(3)
Lagged Present Value	0.997 (0.003)	0.996 (0.005)	0.994 (0.003)
State Effects		✓	✓
Time Effects			✓

Note: We regress $\sum_{j=0}^T \beta^j \tilde{u}_{i,t+j}$ on its lag, where $\tilde{u}_{i,t}$ is unemployment in state i in quarter t . The sample period is 1978-2018. We set $T = 20$. Unemployment is in percentage points. The regression is unweighted. Standard errors are in parentheses. These are two-way clustered by date and state. The number of observations is 7378.

by running the following regression

$$\sum_{j=0}^T \beta^j \tilde{u}_{i,t+j} = \alpha_i + \gamma_t + \theta \sum_{j=0}^T \beta^j \tilde{u}_{i,t+j-1} + \epsilon_{i,t}.$$

As in the main text, we truncate the infinite sum at $T = 20$. Table A.1 presents results for this regression. The coefficient on the first lag is very close to 1 in all specification. This is the case, despite the Nickell (1981) downward bias for autoregressive coefficients in finite samples.

A.4.2 Procedure for Obtaining Correct Standard Errors for κ

Next, we outline a simple procedure to obtain correct the standard errors for κ , using two stage least squares. The reduced form equation is equation (23) from the main text:

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

The first stage equation is equation (24) from the main text:

$$\sum_{j=0}^T \beta^j u_{i,t+j-4} = \zeta u_{i,t-4} + \delta p_{i,t-4}^N + \alpha_i + \gamma_t + \nu_{i,t-4}.$$

The second stage equation, adapted from equation (42), is

$$\begin{aligned}
\pi_{it} &= -\kappa(E_{t-4} \sum_{j=0}^T \beta^j u_{i,t+j} + E_{t-4} \sum_{j=0}^T \beta^j u_{i,t+j-1} \\
&\quad + E_{t-4} \sum_{j=0}^T \beta^j u_{i,t+j-2} + E_{t-4} \sum_{j=0}^T \beta^j u_{i,t+j-3}) \\
&\quad - \delta \hat{p}_{H,t-4}^N + \alpha_i + \gamma_t + \eta_{it} \\
&= -4\kappa \sum_{j=0}^T \beta^j u_{i,t+j-4} - \delta \hat{p}_{H,t-4}^N + \alpha_i + \gamma_t + \eta_{it},
\end{aligned}$$

where π_{it} is cumulative inflation over four quarters. So, -4κ can be readily estimated by two stage least squares, yielding the correct standard errors. In practice, to obtain a value of κ , one can divide the outcome variable by -4.

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

$$\pi_t - E_t \pi_{t+\infty} = \alpha + \psi \tilde{u}_{t-4} + \epsilon_t, \quad (44)$$

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: Slope of the Aggregate Phillips Curve

	Pre-1990	Post-1990
	(1)	(2)
Core CPI	0.796 (0.120)	0.111 (0.027)
Median CPI	0.386 (0.136)	0.250 (0.032)
Shelter CPI	1.624 (0.350)	0.396 (0.050)
PCE	0.416 (0.078)	0.034 (0.021)
Core less Shelter CPI	0.221 (0.103)	-0.084 (0.028)
Core CPI RS	0.182 (0.108)	0.150 (0.028)

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); mort-




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.

gage 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.

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 observations 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 (45) 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

$$r_{it} = \left(\frac{P_{it}}{P_{i,t-\tau} + QA_{i,t-\tau,t}} \right)^{1/\tau}, \quad (45)$$

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
- tax 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

References

- ADAM, K. AND M. PADULA (2011): "Inflation Dynamics and Subjective Expectations in the United States," *Economic Inquiry*, 49, 13–25.
- 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.
- 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.
- CHODOROW-REICH, G. AND J. WIELAND (2019): "Secular Labor Reallocation and Business Cycles," *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.

- 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.
- 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.
- 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.
- GEEROLF, F. (2019): "The Phillips Curve: A Relation between Real Exchange Rate Growth and Unemployment," Working Paper, UCLS.
- 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.
- JORGENSEN, P. L. AND K. J. LANSING (2019): "Anchored Inflation Expectations and the Flatter Phillips Curve," Working Paper, Federal Reserve Bank of San Francisco.
- 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.
- MONTIEL OLEA, J. L. AND M. PLAGBORG-MOLLER (2020): "Local Projection Inference is Simpler and More Robust Than You Think," Working Paper, Columbia University.
- 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.
- 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.
- NICKELL, S. (1981): "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 49, 1417–1426.
- 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 Inflation 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. AND M. WOODFORD (1997): "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 12, 297–346.
- RUDD, J. AND K. WHELAN (2005): "New Tests of the New-Keynesian Phillips Curve," *Journal of Monetary Economics*, 52, 1167–1181.
- 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.
- 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.