Is There a Stable Relationship between Unemployment and Future Inflation?[*]

Terry Fitzgerald Callum Jones Mariano Kulish
Juan Pablo Nicolini

August 2022

Abstract

Evaluating the stability of the Phillips curve using aggregate data is challenging, due to the bias that endogenous monetary policy imparts on estimated Phillips curve coefficients. We argue that regional data can be used to identify the structural relationship between unemployment and inflation. Our analysis using city and state-level data from 1977-2017, is consistent with the notion that both the reduced form and the structural parameters of the Phillips curve are, to a substantial degree, quite stable.

Keywords: Endogenous monetary policy; Stability of the Phillips curve.
JEL classifications: E52, E58.

[*]This is substantially revised version of Fitzgerald and Nicolini (2014). We would like to thank Marco Bassetto, V.V. Chari, Jordi Galí, Juan Herreño, Constantino Hevia, Patrick Kehoe, Ellen McGrattan, Jón Steinsson, Martin Sola, Pedro Teles, Silvana Tenreyro, and Mike Woodford for helpful discussions, and Han Gao, Brian Holtemeyer, and Claire Hou for outstanding research assistance. We also thank numerous seminar participants for comments. The views expressed herein are those of the authors and not necessarily those of the Board of Governors of the Federal Reserve System, the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Fitzgerald: Federal Reserve Bank of Minneapolis. Terry.Fitzgerald@mpls.frb.org. Jones: Federal Reserve Board. callum.j.jones@frb.gov. Kulish: University of Sydney. mariano.kulish@sydney.edu.au. Nicolini: Federal Reserve Bank of Minneapolis and Universidad Di Tella. juanpa@minneapolisfed.org.

1 Introduction

We revisit the empirical debate on the stability of the Phillips curve over time, using data from the United States. Our main innovation is the use of state-level data for that purpose. There are two principal reasons for this strategy. The first is that if a central bank responds to shocks with the purpose of maintaining inflation close to some target, aggregate data may be largely uninformative as to the existence of a stable relationship between unemployment and future inflation. The second is that as monetary policy responds to aggregate shocks only, state-level shocks can be used to identify the key parameters.

Nakamura and Steinsson (2014) used regional data to identify the fiscal multiplier. We borrow their idea to address the bias brought about by the problem of endogenous policy in a Phillips curve model. Our empirical analysis using state- and city-level data is consistent with Phillips curve parameters that are quite stable over time.

The notion that endogenous policy may introduce an estimation bias is an old one and has been applied in many contexts, including in models with Phillips curves. We revisit this point in a very simple model in which a Phillips curve relationship is assumed to be true. We also assume that the central bank optimally sets monetary policy so as to fully stabilize inflation, and we show that aggregate data alone cannot be used to identify the Phillips curve featured by the model. More generally, if the central bank has a dual mandate, identification is possible, but if the policy rule is misspecified, the estimates of the Phillips curve will be biased.

To motivate the empirical exercises that are the core of the paper, we use the same model to show how regional data can be used to identify the relationship between unemployment and future inflation. The main insight is that as monetary policy reacts only to aggregate shocks, region-specific variation can be used to uncover the true relationship between inflation and unemployment.[1] The analysis with state-level data is consistent with the notion that the rela-

[1] We thank Narayana Kocherlakota for raising this question to us during a 2012 policy briefing at the Minneapolis Fed.

tionship between inflation and unemployment has remained quite stable since the 1970s in the US.

The empirical analysis is done in two complementary ways. First, in Section 3 we study reduced form relationships between inflation and unemployment. We do so in order to address the literature that, as in Atkeson and Ohanian (2001), has criticized Phillips curve models using reduced form models. We first document that, as is well known, the estimated reduced form parameter using aggregate data does exhibit substantial variation over time. We then show that when using state-level data, as suggested by the theory, the estimate of the reduced form coefficient is quite stable over time. This is so, even though we compare the period of high and unstable inflation (1977–1985) with the subsequent decades, in which inflation was much lower and stable.

Whatever structural model is behind the time series properties of inflation and unemployment, the reduced form results are consistent with the notion that controlling for aggregate shocks – including monetary policy shocks – a 1 percentage point reduction in unemployment is typically associated with a 0.2 to 0.3 percentage point increase in inflation. More importantly, this behavior seems to have been stable over time.

Using regional data eliminates policy induced biases, but it brings in other complications. One must properly account for the heterogeneity across regions, including heteroskedastic shocks that may render the aggregation problem complicated. Our benchmark analysis is based on strong homogeneity assumptions, and some robustness exercises will be provided.

Perhaps the most important caveat is that in order to remove the effect of the aggregate, one needs to use regional inflation as a deviation from national inflation. But while inflation is a purely nominal variable, the difference between state level inflation and aggregate inflation is the relative price between the non-traded goods at the state level and the goods that are traded across states. Thus, our estimated reduced form coefficient could reflect Balassa-Samuelson-like effects for instance, a purely real phenomenon, unrelated to the Phillips curve.

The natural way to address these caveats is to estimate a structural model that allows for a variety of other real shocks that can account for movements in the relative price between tradables and non-tradables. The model can also be estimated for separate subset of regions, that are similar among them, so as to evaluate the role of heterogeneity.

This motivates our second exercise, discussed in Section 4, in which we set up a structural model with frictions in the setting of prices and wages. We first show that regional data can be used to circumvent the endogenous policy bias also in the structural model. We then show that the estimated Calvo parameter for prices using state-level data is very stable over time. The analysis does detect some statistically significant instability in the wage Calvo parameter. We do argue, however, that when translated to either the slope of the Phillips curve or the implied frequency of wage changes, the difference is of little economic significance. We also show in this section that, in contrast, when aggregate data alone is used, the estimates are sensitive to the sample period and the assumptions regarding the monetary policy rule.

Our results imply a value of about seven to eight months for the average duration of price contracts and an average duration of between five and seven months for wage contracts. Both estimates are in line with the micro evidence on nominal frictions, as we discuss in Section 4.

The structural model makes clear that while the reduced form slope is related to the parameters that govern the slope of the New Keynesian Phillips curve, it is a different object. The stability of the reduced-form slope estimate is of interest not just because it suggests the possible stability of the structural parameters but because it is model free and can serve both as a reference to evaluate any structural model and to address policy questions.[2]

When translated into slopes of the Phillips curves, our estimates are in line with the ones in McLeay and Tenreyro (2020) who use partial information methods to directly estimate the New Keynesian Phillips curve.

However, our results are quite different from the ones obtained in Hazell,

[2]We have used the estimated structural model to simulate data to conduct reduced form regressions. We find that our estimated structural model gives rise to a quite similar reduced form slope as that computed using the data. Our estimated structural model is thus consistent with the reduced form evidence.

Herreno, Nakamura and Steinsson (2020), who estimate the New Keynesian Phillips curve model using reduced form regressions similar to ours. Using theory, they obtain a mapping between the reduced form coefficients and properties of the unemployment rate, to the slope of the Phillips curve.[3] They also find a stable Phillips curve, but their estimates imply a much lower response of inflation to a 1% increase in unemployment. A full analysis of the main drivers of these differences is beyond the scope of this paper, since our model does not include a search block where unemployment can be jointly analyzed with inflation.[4]

We would like to finish this introduction by highlighting that while we view our attempts at controlling for some of the difficulties that using state level data to identify aggregate parameters as a few first steps, we believe that a key contribution of the paper is to highlight properties of cross-sectional data which suggest its potential to address identification issues. Obtaining a precise estimate of the slope of the Phillips curve does require further research.

The paper is organized as follows. Section 2 briefly discusses related papers. In Section 3, we first show in a simple theory how endogenous monetary policy can blur the true structural relationship in the aggregate. We also show how this is not the case for the regional data, since regional variation can be used to identify the true structural parameters. We then run the regressions implied by the theory, using data from 27 metropolitan statistical areas (MSAs) in the US from 1976 to 2018. As we show, the regressions are consistent with the notion of a reduced form Phillips curve that has remained stable over time. In Section 4, we estimate a full New Keynesian model separately on state and aggregate data. We find that the estimates of the structural parameters that govern the frequency of price and wage adjustments are found to be quite stable over time when using state-level data, echoing the reduced form findings. By contrast, the estimates using aggregate data are sensitive to the time period used in estimation.

[3] They approximate the process of unemployment as an exogenous autoregressive process that they estimate from the data.

[4] We are currently developing a model with search frictions and heterogeneous regions to address this issue.

2 Background on the Phillips Curve and Related Literature

The hypothesis of an exploitable Phillips curve remains a controversial subject. For example, Atkeson and Ohanian (2001) (henceforth AO) show that the empirical relationship between current aggregate unemployment and inflation growth is highly unstable over the period 1960–2000 in the US. They forcefully argue this point by showing that a naive prediction rule for inflation that simply uses past inflation is systematically better than empirical Phillips curves at forecasting inflation. A natural interpretation of their results follows from the observation that the period covered by the analysis includes changes in the policy regime. Thus, the corresponding shift in parameters is evidence that the relationship is not structural, an unavoidable corollary of the Lucas critique.

Even the most extreme defender of the New Keynesian paradigm would agree with the notion that the Calvo parameter is not invariant to *any* policy regime change. The quantitative question we pursue is whether the Calvo parameters can be safely assumed to be policy invariant – and therefore not subject to the Lucas critique – *given* the policy regime changes actually experienced by the US in the postwar era. The evidence in this paper points towards a positive answer to that question. Our interpretation of the evidence in AO, therefore, is that the instability over time of the estimated relationship using aggregate data is the result of policy changes, along the lines discussed in Sargent (1999). Under this interpretation, the evidence in AO is uninformative regarding the true relationship between current unemployment and future inflation.

Recent events have again put into question the stability of the Phillips curve relationship. The "flattening" of the Phillips curve has been debated at length, fed by the strong changes in unemployment rates in the United States during the 2008–2009 recession and the subsequent recovery, with little sign of inflation rates responding to those movements. A series of papers addressing this issue followed the policy debate.[5]

These criticisms exhibit two main characteristics. First, aggregate data are used in the analysis. This is problematic since, as mentioned above, a bias arises when monetary policy endogenously responds to shocks, as preceding literature discussed in detail below has forcefully argued. Second, these criticisms are based, albeit most of the time implicitly, on the behavior of reduced form parameters over time, which makes addressing the identification problem hard.[6] The paper of AO represents a concrete example, and its virtue is that it is explicit regarding the nature of the exercise. But arguing that the stagflation of the '70s represents evidence of an unstable Phillips curve, as many do, also entails a reduced form discussion, and so does arguing that the "missing deflation" in 2009 and 2010 and the subsequent "missing inflation" represent evidence of a flattening of the Phillips curve. So, while many times we will directly compare our results with a particular interpretation of AO, it should be understood that our results speak to a broader literature, which evaluates the stability of the Phillips curve in its structural form as well.

Our empirical exploration using state-level data is consistent with the notion that the slopes of price and wage Phillips curves in a standard New Keynesian model are roughly invariant to the policy regimes experienced in the US since 1977, the first year for which we have data.

These results suggest an alternative interpretation of the data used by proponents of the "shifting Phillips curve," first developed by Hazell, Herreno, Nakamura and Steinsson (2020). According to this interpretation, the changes over time in the correlation between unemployment and inflation observed in aggregate data could result from changes in the policy regime followed by the Federal Reserve over the period. The evidence in AO is compatible with a change in the policy rule that started somewhere in the '80s. And the stagflation of the '70s is the result of a monetary policy that made inflation persistently higher, at a time in which the economy was undergoing a recession.[7]

[5] See Krugman (2015); Blanchard (2016); and, for a recent survey of the literature, Hooper, Mishkin and Sufi (2019).

[6] There are a few exceptions, such as Coibion and Gorodnichenko (2015).

[7] See Gao, Kulish and Nicolini (2020) for an interpretation along these lines.

This rather brief account of the recent history of US monetary policy evolved in an economy in which the frequency of price and wage changes remained quite stable over time – at least, so says our state-level analysis.

As mentioned above, the notion that endogenous policy makes identification of structural parameters problematic is an old idea. It dates at least to the work of Samuelson and Solow (1960) and Kareken and Solow (1963). It has since then been applied in several contexts by Brainard and Tobin (1968), Goldfeld and Blinder (1972), Worswick (1969), Peston (1972), and Goodhart (1989). These papers show that if policy reacts to the state of the economy, the relationship in the aggregate data can be blurred by the policy rule.

Haldane and Quah (1999) were the first to apply it to a modern New Keynesian model. They assume that the central bank has a dual mandate and optimally chooses policy. They show that the estimated relationship is a function of the relative weight that the central bank puts on inflation. Their analysis was followed by Mishkin (2007); Carlstrom, Fuerst and Paustian (2009); and Edge and Gurkaynak (2010), among others.

As mentioned above, our work is heavily inspired by Nakamura and Steinsson (2014), who used regional data to identify the fiscal multiplier. We apply their idea to the estimation of Phillips curve model. Our strategy, spelled out in the working paper version of this paper (see Fitzgerald and Nicolini, 2014) has since been used by Kiley (2015), Babb and Detmeister (2017), Leduc and Wilson (2017), and more recently by Levy (2019), Hooper, Mishkin and Sufi (2019), McLeay and Tenreyro (2020), and Hazell, Herreno, Nakamura and Steinsson (2020).

Beraja, Hurst and Ospina (2019) and Jones, Midrigan and Philippon (2022) use state- and aggregate-level data together to identify the parameters of structural New Keynesian models; however, those papers do not speak to the issue we address – namely, the stability of Calvo price and wage parameters over time. We additionally use information on prices at the MSA level in estimating the New Keynesian model.

Finally, Nishizaki and Watanabe (2000) is an empirical paper that uses regional data from Japan to control for aggregate shocks and examine the stability of the Phillips curve. While it does not discuss the role of endogenous policy, one of its exercises is similar to the one that we perform in our reduced form analysis, discussed next.

3 Reduced Form Analysis

In this section, we use a reduced form representation to guide some simple regression analysis. The main reason to do so is that a sizeable share of the literature addressing the stability of the Phillips curve has framed the discussion in reduced form terms, as discussed in detail in Section 2.

Consider an economy composed of a continuum of geographically separated regions that potentially exhibit price frictions. All regions use the same unit of account and face the same monetary policy. Let $\pi_t(s), u_t(s)$ represent regional inflation and unemployment for region s. Assume also that the equilibrium solution in each region can be characterized by the following dynamic system:

$$\begin{aligned} \pi_{t+1}(s) &= b\pi_t(s) + cu_t(s) + di_t + eX_t(s) + \varepsilon^{\pi}_{t+1}(s) + \xi^{\pi}_{t+1} \quad (1) \\ u_{t+1}(s) &= b'\pi_t(s) + c'u_t(s) + d'i_t + e'X_t(s) + \varepsilon^{u}_{t+1}(s) + \xi^{u}_{t+1}, \quad (2) \end{aligned}$$

where $\varepsilon^j_t(s)$ and ξ^j_t, for $j = u, \pi$, are the regional and aggregate shocks; i_t is the interest rate determined by monetary policy, to be discussed below; and $X_t(s)$ is a vector that allows for the inclusion of control variables in the regression analysis that follows. We call the dynamic system defined by (1) and (2) the reduced form of some structural model. The vector $X_t(s)$ is introduced to allow for control variables in the regression analysis that follows. To simplify the algebra, we now set $X_t(s) = 0$ for all t, s.

We assume that the underlying structural model is such that all shocks have zero unconditional means and regional shocks are independent of the aggregate shock. The terms di_t and $d'i_t$ describe the effect of monetary policy on the system. The timing indicates that the monetary authority decides on policy before observing the $t+1$ shocks.

Letting $\varphi(s)$ be state weights with $\int_0^1 \varphi(s)\ \mathrm{d}s = 1$, the aggregates are:

$$\begin{aligned}\pi_{t+1} &= \int_0^1 \varphi(s)\pi_{t+1}(s)\ \mathrm{d}s \\ u_{t+1} &= \int_0^1 \varphi(s)u_{t+1}(s)\ \mathrm{d}s.\end{aligned}$$

We then obtain the following relationship between the aggregate variables:

$$\pi_{t+1} = b\pi_t + cu_t + di_t + \xi^{\pi}_{t+1} \tag{3}$$

$$u_{t+1} = b'\pi_t + c'u_t + d'i_t + \xi^{u}_{t+1}. \tag{4}$$

The focus of this section is the ability to identify and estimate the parameters of the reduced form equations (3) and (4).

A particular example of a structural model that delivers a reduced form like the one described above will be discussed in the next section, where we also estimate its structural parameters. But the system defined by (3) and (4) is compatible with many other models. In particular, as we show in the Online Appendix, this reduced form is also consistent with a simple old Keynesian model essentially identical to the one presented in Taylor (1999) and discussed in Cochrane (2011). As we show there, under this interpretation, the coefficient c in (3) can be associated with the slope of a NAIRU Phillips curve.

The stability over time of parameter c in equation (3), particularly across different monetary policy regimes, has been the focus of much discussion in the literature. In particular, the natural interpretation of the analysis in Atkeson and Ohanian (2001) is that the estimate of c obtained using aggregate data is unstable over time. We now address this issue.

3.1 Exogenous Policy

To fix ideas, assume first that the monetary authority follows an exogenous constant interest rate policy. Then, taking differences in (3), equilibrium inflation evolves as

$$\pi_{t+1} - \pi_t = b\left(\pi_t - \pi_{t-1}\right) + c\left(u_t - u_{t-1}\right) + \left(\xi^{\pi}_{t+1} - \xi^{\pi}_{t}\right). \tag{5}$$

Under this policy, standard econometric techniques should suffice to identify the parameter c.

Figure 1 shows the rolling coefficient for c that results in estimating an equation (5) using inflation and unemployment data for the US from 1975 to 2017. In estimation, we use a series of controls, detailed in the Appendix. We estimate the equation using both headline and core inflation, which explains why we have two solid lines in the figure. Specifically, for each of the two measures of inflation, we first estimate the coefficient c in equation (5) using semiannual data from the first semester of 1975 to the second semester of 1995.[8] The resulting point estimate is then plotted in Figure 1 as the value for the second semester of 1995. We then repeated the estimation, but using data starting and ending one semester after; plotted the point estimate for the first semester of 1996; and reproduced the steps moving forward. Each point in the series thus represents the point estimate of c for a sample size that starts 20 years before and ends at that point. The dotted lines represent 90% confidence intervals.

The figure makes clear how the point estimate for c depends on the sample period. For instance, when we use headline inflation, the first estimate is very close to -1, but it decreases over time to become zero by the end of the sample. A similar but less drastic change is apparent for the estimates using core inflation. The picture explains why using a Phillips curve like (5) estimated using aggregate data would perform poorly as an out-of-sample forecasting device. This explains the exercise in Atkeson and Ohanian (2001).

To the extent that policy is exogenous, Figure 1 offers evidence that is inconsistent with a stable value for c in this model. But our take is different: as policy is not exogenous, the evidence provided in Figure 1 is in itself un-

[8] We use semiannual data because the frequency for which we have regional data is semiannual. We also used a few controls, as explained in the Online Appendix. The results without controls, also reported in the Online Appendix, are very similar.

Figure 1: Coefficient from Rolling 20-Year Regression, Aggregate Level

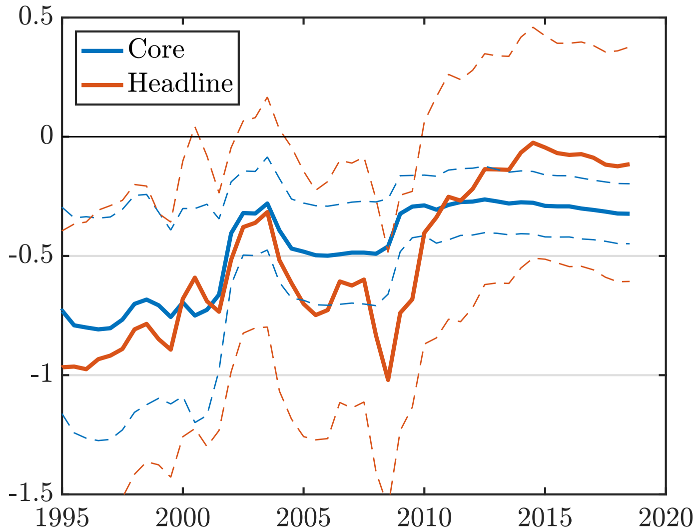

informative regarding the value of the reduced form parameter c. We address this issue next.

3.2 Endogenous Policy

We now assume the central bank has a mandate to stabilize inflation. We also assume the central bank knows the model economy. Specifically, it solves the following policy problem:

$$\min_{i_t} \frac{1}{2} E_t \left[\pi_{t+1} - \pi^*_{t+1} \right]^2 ,$$

given π_t, u_t, and the solution for aggregate inflation (3) . The target for inflation is given by π^*_{t+1} and is part of the policy rule. The objective function is defined as the time$-t$ expectation of the deviation of next period inflation relative to the target. Implicit in this way of writing the problem is the assumption that the central bank chooses policy before observing time $t+1$ shocks.

As shown in the Appendix, the optimal policy rule[9] is

$$i_t^{Opt} = \frac{1}{d}\left[\pi^*_{t+1} - \left(b\pi_t + cu_t + E_t\xi^\pi_{t+1}\right)\right], \tag{6}$$

so the equilibrium value for inflation is given by

$$\pi_{t+1} = \pi^*_{t+1} + \xi^\pi_{t+1} - E_t\xi^\pi_{t+1}. \tag{7}$$

Inflation in equilibrium therefore equals the target plus a forecasting error that, by definition, is orthogonal to any variable in the central bank's information set at time t. In particular, inflation is independent of all the model parameters. This is the consequence of a central bank that knows the model of the economy and uses it to design policy so as to stabilize a specific target.[10] A direct implication of this observation is that if the central bank's only objective is to stabilize inflation and it uses a model that describes the economy well, the behavior of inflation in equilibrium is completely uninformative regarding the underlying model that determines inflation. It should be obvious by now that this property is independent of the model that determines inflation, as long as the central bank knows it.

The behavior of equilibrium inflation depends on the behavior of the target, π^*_{t+1}, which is not necessarily observable. To gain further insight, we next consider two specifications. Consider first the case of a constant inflation target, so $\pi^*_t = \pi^*$ for all t. Then, taking differences in (7),

$$\pi_{t+1} - \pi_t = \left(\xi^\pi_{t+1} - E_t\xi^\pi_{t+1}\right) - \left(\xi^\pi_t - E_{t-1}\xi^\pi_t\right),$$

so current unemployment would be related to the change in inflation to the extent that the forecast error $(\xi^\pi_t - E_{t-1}\xi^\pi_t)$ affects unemployment u_t. But if an estimate of the change in inflation that is different from zero is obtained, it is unrelated to the direct effect of unemployment on future inflation, or c.

[9] We show in the Online Appendix that with this policy rule, there is a unique solution. See also Cochrane (2011) for a discussion of determinacy in models of this type.

[10] As mentioned in Section 2, this insight is not new. See the literature quoted therein.

Assume next that

$$
\begin{aligned}
\pi_t^* &= \pi_{t-1}, \text{ if } \pi_{t-1} \in [\pi_{\min}, \pi^{\max}] \\
\pi_t^* &= \pi^{\max}, \text{ if } \pi_{t-1} > \pi^{\max} \\
\pi_t^* &= \pi_{\min}, \text{ if } \pi_{t-1} < \pi_{\min}.
\end{aligned}
\tag{8}
$$

This case corresponds to a central bank that establishes a range for the target and, to the extent that current inflation is within the bands, wants to keep inflation equal to the previous period. As long as the target remains within the band, $\pi_{t+1}^* = \pi_t$, then

$$\pi_{t+1} - \pi_t = \xi_{t+1}^{\pi} - E_t \xi_{t+1}^{\pi},$$

so inflation follows a random walk. In this case, current unemployment–or, for that matter, any variable in the information set at time t–should not help predict inflation growth. In this case, no forecasting rule for inflation could beat a random walk. As shown in the Online Appendix, the reduced form (3) and (4) are consistent with a simple NAIRU-type model. Therefore, such a model, coupled with the assumption that the central bank stabilizes inflation around a target as defined in (8), generates equilibrium observations that are fully consistent with the result that a random walk is good predictor for inflation, as in AO. The example also rationalizes the difficulty the literature encountered in its attempts at developing trustworthy forecasting models for inflation, as explained in Stock and Watson (2009). In the next section we explain why state-level data can be used to tackle the endogeneity problem.

3.3 State-Level Data Regressions

We now show how to estimate the reduced form parameters exploiting the fact that regional variables' deviations from the national average will not be correlated with policy.

We first replace the optimal policy (6) into the solution for inflation in each region (1) and obtain

$$\pi_{t+1}(s) = \pi^*_{t+1} + b\left(\pi_t(s) - \pi_t\right) + c\left(u_t(s) - u_t\right) + \varepsilon^\pi_{t+1}(s) + \xi^\pi_{t+1} - E_t\xi^\pi_{t+1}. \quad (9)$$

Notice that by exploiting state-level deviations from the national average, the effect of policy does not enter the solution.

In order to estimate equation (9), we need to take a stand on the evolution over time of the target for inflation. In what follows, we consider an agnostic specification. Thus, we define a time dummy and run

$$\pi_{t+1}(s) = D_t + b\left(\pi_t(s) - \pi_t\right) + c\left(u_t(s) - u_t\right) + \varepsilon^\pi_{t+1}(s) + \left(\xi_{t+1} - E_t\xi_{t+1}\right). \quad (10)$$

The time dummy is naturally interpreted as an estimate of the inflation target for each period in the context of this simple model.[11] Of course, it could also pick up other aggregate shocks, like mark up-shocks or shocks to energy prices. Finally, as we did in the case of aggregate data, we use a series of controls that we briefly discuss below.

3.4 Results

In this section, we show the results using CPI inflation and unemployment data for 27 metropolitan statistical areas in the US. For many MSAs and periods, the lowest frequency for the data is semiannual, so we used that frequency to construct the database. The price data for MSAs are available only as non-seasonally adjusted, so we compute yearly changes. In our regressions we define $u_t(s)$ as the period t unemployment rate for MSA s and $\pi_{t+1}(s)$ as the inflation rate over the following year (i.e., $CPI_{t+2}(s)/CPI_t(s)$). We use headline as a measure of inflation, for which we have data since 1977.[12]

There are a few issues that we need to address in order to clarify the way we will interpret the estimated parameters of equation (10). Our first inter-

[11] In the working paper version of this paper (Fitzgerald and Nicolini, 2014), we discuss more specific assumptions that lead to alternative formulations for the regression. We also compare the results of those regressions with this agnostic strategy.

[12] The Online Appendix describes this data set in detail.

pretation will be based on our use of system (3) and (4) as representing purely a reduced form of an unspecified structural model. As such, the estimates provide information only on such a reduced form and lack any additional interpretation. For that purpose, a simple OLS regression suffices, and the only relevant question is if the estimate of the coefficient c is stable over time.

A second possibility is to interpret the system (3) and (4) as a reduced form of a NAIRU (old) Keynesian model. Under that interpretation, the coefficient c approximates the estimate of the slope of the NAIRU Phillips curve, as we show in the Online Appendix. However, for the OLS estimator to be unbiased, it is necessary that unemployment be uncorrelated with the shock, $\varepsilon^{\pi}_{t+1}(s) + \xi^{\pi}_{t+1} - E_t\xi^{\pi}_{t+1}$. The second component, being a forecast error, presents no difficulty. However, if the region-specific shock is autocorrelated over time, there will be a bias. In that case, it may be important to use instrumental variables. To this end, we will also report two-stage least-squares (2SLS) results in what follows. We have no natural instrument, but since the problem arises only if the regional shocks are autocorrelated, using lagged values of the unemployment rate would naturally reduce the bias. Thus, we use lagged values of the unemployment rate in the first stage. As further justification for this interpretation, one can analyze the estimates of the autocorrelation of the errors. We do so in the working paper version of this paper (Fitzgerald and Nicolini, 2014), where we show that there is no strong evidence of autocorrelation being a major issue in our preferred specification.

The variable $u_t(s)$ and u_t represents deviations from the natural rate of unemployment. To allow for the possibility that the natural rate of unemployment differs across MSAs, we introduce a region fixed effect in the regressions. In addition, to control for potential heteroscedasticity, we compute the statistical tests using standard errors that are clustered at the MSA level. All tests results are uniformly stronger if we do not cluster the errors. Finally, we use a series of regional controls that may correlate with shocks affecting local economic conditions, like inflation expectations and government expenditures or temperature and precipitations, as well as lagged values of both inflation and unemployment. In particular, we add in some specifications the Bartik-type

Table 1: Regressions with Headline Inflation

	1977-2018	1977-1984	1985-1990	1991-2000	2001-2010	2011-2018
	A. Headline Inflation, OLS, without Controls					
c	-0.28^{**}	-0.31^{**}	-0.41^{**}	-0.31^{**}	-0.24^{**}	-0.21^{**}
	(0.04)	(0.11)	(0.11)	(0.05)	(0.06)	(0.07)
Overall R^2	0.88	0.83	0.69	0.45	0.70	0.51
Obs	2059	381	288	492	536	362
	B. Headline Inflation, 2SLS, without Controls					
c	-0.27^{**}	-0.39^{**}	-0.29^{*}	-0.46^{**}	-0.21^{**}	-0.24^{**}
	(0.04)	(0.12)	(0.15)	(0.13)	(0.08)	(0.08)
Overall R^2	0.88	0.79	0.71	0.39	0.70	0.51
Obs	2055	377	288	492	536	362
	C. Headline Inflation, 2SLS, with Controls					
c	-0.33^{**}	-0.50^{**}	-0.45^{**}	-0.45^{**}	-0.28^{**}	-0.28^{*}
	(0.05)	(0.19)	(0.14)	(0.10)	(0.10)	(0.13)
Overall R^2	0.88	0.76	0.65	0.40	0.70	0.54
Obs	1933	327	288	484	532	362

Standard errors in parentheses
* significance at 5% level, ** significance at 1% level

instrument to control for government spending constructed by McLeay and Tenreyro (2020). A detailed explanation of the controls, as well as a detailed discussion of how robust the results are is presented in the Online Appendix.

Table 1 provides estimates for the coefficient c in regression (10). Results are reported for OLS and 2SLS without and with controls.[13] We present results for the whole period first and then for five sub-periods. The first sub-period is chosen to contain the years of rising inflation and the Volcker stabilization. The second sub-period contains the rest of the decade until 1990. We take these two to be the ones with policy regimes that differ from the rest of the sample.

The point estimate for c using the whole period is close to -0.3 for the three

[13]We report the estimates for all other parameters in the Online Appendix.

Figure 2: Coefficient from Rolling 20-Year Regression, State Level

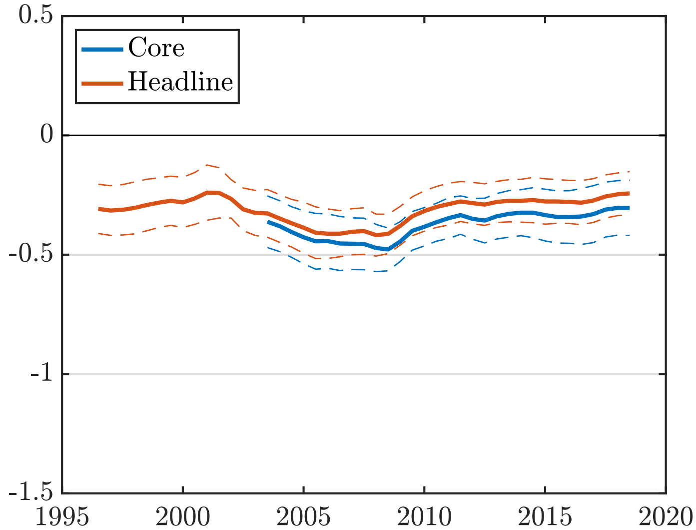

specifications and very precisely estimated. In addition, the point estimate is similar for all the sub-periods and are all statistically significant. In fact, for all specifications and almost all sub-periods, the point estimate is within one standard deviation of -0.3.

As further evidence of the stability of the estimated coefficient, we show in Figure 2 an exercise like the one presented in Figure 1, but using state-level data to run the rolling regressions, rather than aggregate data. In this case, it takes two pictures (Figures 1 and 2) to be worth a thousand words.

In the working paper version of this paper (Fitzgerald and Nicolini, 2014) and its appendix, we performed several additional exercises. We first explored the possibility that results would be driven by a few MSAs so that other geographic issues could affect the results. We also checked if the overlapping nature of our data is important. We finally explored the extent to which autocorrelation of the errors could be an issue, given that we use lagged values of unemployment differentials as an instrument. As we show there, the autocorrelation of unemployment differentials is lower than for unemployment. We showed our results to be quite robust to all these concerns.

These results can be thought of as consistent with an old Keynesian structural model; they thereby relate to the criticism of Atkeson and Ohanian and others. But they can be interpreted as reduced form regressions from the perspective of current structural New Keynesian models. One may therefore wonder the extent to which the results of this section speak to the stability of the frequency of price and wage adjustment in structural New Keynesian models. This is a natural question to raise, since the coefficients of reduced form solutions are functions of the parameters of the corresponding structural model. Thus, we now estimate a simplified version of the the state-level structural model of Jones, Midrigan and Philippon (2022).

4 Structural Model

We now move beyond linear reduced forms and estimate an economy with Calvo-type rigidities in prices and wages. We use our estimation results to evaluate the stability of the parameters over time. As discussed in Section 2, the assumptions in Calvo are not to be understood as invariant to any policy regime change. The question we address is whether those parameters have been stable across the monetary regime changes that have prevailed in the US since 1977, the first year for which we have state-level data.

We employ the simplest framework, which forms the basis of numerous models in the literature. Thus, we use as a starting point the standard three-equation New Keynesian model. In adapting that model to a series of geographically separated units in which local shocks can move local pricing and employment decisions that are different than those for the country as a whole, we do need to extend that basic popular model to allow for tradable and non-tradable goods. This is the only deviation from the standard textbook example of the New Keynesian model with price and wage frictions. As the model is a variation of a standard small scale New Keynesian model, we describe the main features below and we relegate the full description to the Appendix.

4.1 Model Description

The economy consists of a continuum of ex ante identical islands. These islands form a monetary union and trade with one another. Consumers on each island derive utility from the consumption of a final good and from leisure:

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta_t(s) \left[\log(c_t(s)) - \frac{\eta_t^n(s)}{1+\nu} n_t(s)^{1+\nu} \right],$$

where s indexes the island, $c_t(s)$ is consumption, $n_t(s)$ is labor supplied, $\beta_t(s)$ is a preference shock, and $\eta_t^n(s)$ is a labor disutility shock. The structure of the shock processes is described below.

The final good $y_t(s)$ is assembled using inputs of non-traded goods $y_t^N(s)$ and traded goods $y_t^M(s,j)$ imported from island j:

$$y_t(s) = \left(\omega^{\frac{1}{\sigma}} y_t^N(s)^{\frac{\sigma-1}{\sigma}} + (1-\omega)^{\frac{1}{\sigma}} \left(\int_0^1 y_t^M(s,j)^{\frac{\kappa-1}{\kappa}} \mathrm{d}j \right)^{\frac{\kappa}{\kappa-1}\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

where ω determines the share of non-traded goods, σ is the elasticity of substitution between non-traded and traded goods, and κ is the elasticity of substitution across varieties of traded goods. Letting $p_t^N(s)$ and $p_t^M(s)$ denote the inputs' corresponding prices, the price of the final good on an island is

$$p_t(s) = \left(\omega p_t^N(s)^{1-\sigma} + (1-\omega) \left(\int_0^1 p_t^M(j)^{1-\kappa} \mathrm{d}j \right)^{\frac{1-\sigma}{1-\kappa}} \right)^{\frac{1}{1-\sigma}}. \qquad (11)$$

Notice that in the particular case of $\omega = 0$, there are only traded goods and the consumption basket in each location is the same as in the aggregate, in which case inflation in each state is the same as in the aggregate and the model collapses to the simple textbook three-equation model. Thus, the only innovation of our model is to allow for non-traded goods at the state level, which in turns explains why inflation at the regional level may differ from the aggregate.

The production technologies we use are standard in both the monetary and the trade literatures. In particular, we model non-traded goods and traded export goods $y_t^X(s)$ on each island as CES composites of varieties k of differentiated intermediate inputs with an elasticity of substitution ϑ:

$$y_t^N(s) = \left(\int_0^1 y_t^N(s,k)^{\frac{\vartheta-1}{\vartheta}} \mathrm{d}k\right)^{\frac{\vartheta}{\vartheta-1}}$$
$$y_t^X(s) = \left(\int_0^1 y_t^X(s,k)^{\frac{\vartheta-1}{\vartheta}} \mathrm{d}k\right)^{\frac{\vartheta}{\vartheta-1}}.$$

The production of the varieties of non-traded goods and the varieties of traded exports on each island is linear in labor:

$$y_t^N(s,k) = z_t^N(s) n_t^N(s,k)$$
$$y_t^X(s,k) = z_t^X(s) n_t^X(s,k),$$

where $z_t^N(s)$ and $z_t^X(s)$ are productivity shocks.

Nominal frictions affect this economy in a standard way. Individual producers of non-tradable and tradable intermediate goods are subject to Calvo price adjustment frictions–parameterized by λ_p, the probability that a firm cannot reset its price in a given period. The evolution of non-tradable prices thus evolves according to

$$\hat{p}_t^N(s) = \lambda_p \hat{p}_{t-1}^N(s) + (1-\lambda_p)\hat{p}_t^{N\star}(s),$$

where variables with hats denote their log-deviation from steady-state and $\hat{p}_t^{N\star}(s)$ is the optimal price that firms set in the case they are able to adjust prices. This optimal price is forward looking and solves

$$\hat{p}_t^{N\star}(s) = \beta\lambda_p \mathbb{E}_t \hat{p}_{t+1}^{N\star}(s) + (1-\beta\lambda_p)(\hat{w}_t(s) - \hat{z}_t^N(s)),$$

where $\hat{w}_t(s)$ is nominal wages. These two equations yield the price Phillips curve for non-tradable goods:

$$\hat{\pi}_t^N(s) = \beta \mathbb{E}_t \hat{\pi}_{t+1}^N(s) + \frac{(1-\beta\lambda_p)(1-\lambda_p)}{\lambda_p} \left(\hat{w}_t(s) - \hat{p}_t^N(s) - \hat{z}_t^N(s) \right),$$

so that the slope of the price Phillips curve in non-tradable inflation is $\frac{(1-\beta\lambda_p)(1-\lambda_p)}{\lambda_p}$. An analogous argument applies to the slope of the Phillips curve in tradable inflation.

Labor is immobile across states and is aggregated using a CES aggregator with an elasticity of substitution across labor varieties of ψ. Individual households supply differentiated varieties of labor that are subject to Calvo wage adjustment frictions parameterized by λ_w, the probability that a labor variety cannot reset its wage in a given period. The evolution of wages evolves according to

$$\hat{w}_t(s) = \lambda_w \hat{w}_{t-1}(s) + (1-\lambda_w)\hat{w}_t^\star(s),$$

where $\hat{w}_t^\star(s)$ is the optimal wage that unions set in the case they are able to adjust wages, which evolves according to

$$\hat{w}_t^\star(s) = \beta\lambda_w \mathbb{E}_t \hat{w}_{t+1}^\star(s) + \frac{(1-\lambda_w\beta)}{(1+\psi\nu)} \left(-\hat{\mu}_t(s) + \psi\nu\hat{w}_t(s) + \log\eta_t(s) - \log\eta + \nu\hat{n}_t(s) \right),$$

where $\hat{\mu}_t$ is the shadow value of wealth.

Applying the same logic used in deriving the price inflation Phillips curve to this optimal wage setting equation, we can derive the following wage inflation Phillips curve:

$$\hat{\pi}_t^w = \beta \mathbb{E}_t \hat{\pi}_{t+1}^w + \frac{(1-\beta\lambda_w)(1-\lambda_w)}{\lambda_w} \frac{1}{(1+\psi\nu)} \left(-\hat{\mu}_t(s) - \hat{w}_t(s) + \log\eta_t(s) - \log\eta + \nu\hat{n}_t(s) \right).$$

At the aggregate level, monetary policy is set using a Taylor rule when the ZLB does not bind. The nominal interest rate i_t responds to its lag with weight α_r; deviations of inflation π_t from target $\bar{\pi}$ with weight α_π; deviations of output y_t from the flexible-price level of output y_t^F, with weight α_y; and the growth rate of the output gap with weight α_x:

$$1 + i_t = (1 + i_{t-1})^{\alpha_r} \left[(1 + \bar{i}) \left(\frac{\pi_t}{\bar{\pi}} \right)^{\alpha_\pi} \left(\frac{y_t}{y_t^F} \right)^{\alpha_y} \right]^{1-\alpha_r} \left(\frac{y_t}{y_{t-1}} / \frac{y_t^F}{y_{t-1}^F} \right)^{\alpha_x} \exp(\varepsilon_t^i),$$

The following shocks drive fluctuations in the model. At the state level, we have shocks to the rate of time preference of individual households, to the household's disutility from work, to productivity, and to non-tradable productivity.[14] At the aggregate level we also have shocks to the rate of time preference of individual households, labor disutility, and aggregate productivity, in addition to shocks to the interest rate rule ε_t^i and the aggregate price Phillips curve (via standard markup shocks).[15]

The model in Jones, Midrigan and Philippon (2022) has households that also derive utility from the consumption of housing goods, which must be used as collateral for household borrowing. These features allow them to capture better the relative state-level data around the Great Recession described in Mian and Sufi (2011, 2014). In robustness exercises, we add these realistic features to our model and show in the Online Appendix that our results are very robust to this extension.

4.2 Estimation Strategy

We use Bayesian methods, as is common in the literature. Our estimation on state-level data for 51 states over the period 1977 to 2017, however, is not standard: inflation data do not exist for around half of the 51 states in our panel. And the inflation series that are available are observed at only a biannual frequency, whereas the remaining state-level observables are observed annually. So, to rely on as much data as possible, we estimate the state-level model on an unbalanced mixed-frequency panel. To the best of our knowledge, the use of an unbalanced mixed-frequency panel in the estimation of a structural model is new in the literature. We describe the estimation in

[14] In robustness exercises, we also allow for shocks to the household's preference for housing and the loan-to-value borrowing constraint (or credit shocks).

[15] The Online Appendix contains a full description of the model.

more detail below.

Approach To capture the period of zero nominal interest rates, we use a piecewise linear approximation as proposed in Guerrieri and Iacoviello (2015), Jones (2017), and Kulish, Morley and Robinson (2017). Under this approximation, the reduced form solution of our model has a time-varying VAR representation:

$$\mathbf{x}_t = \mathbf{J}_t + \mathbf{Q}_t \mathbf{x}_{t-1} + \mathbf{G}_t \epsilon_t,$$

where $\mathbf{x}_t$ collects the state and aggregate endogenous variables and ϵ_t collects the state and aggregate shocks. The time-varying coefficient matrices $\mathbf{J}_t$, $\mathbf{Q}_t$, and $\mathbf{G}_t$, arise because of the non-linearities induced by the ZLB. In the particular case of $\omega = 0$, the vector $\mathbf{x}_t$ includes the current values for the aggregate shocks as well as inflation – which is the same across states – the output gap – which may be different across states, owing to local shocks and the immobile labor force – and the nominal interest rate.

Following Jones, Midrigan and Philippon (2022), we separate the state-level variables from the aggregate variables. We decompose the vector of variables for each island s, expressed in log-deviations from the steady state as $\mathbf{x}_t(s)$, into a component due to state s's dependence on its own history $\mathbf{x}_{t-1}(s)$ and its own shocks $\epsilon_t(s)$ and a component encoding the state-level dependence on aggregate variables:

$$\mathbf{x}_t(s) = \underbrace{\mathbf{Q}\mathbf{x}_{t-1}(s) + \mathbf{G}\epsilon_t(s)}_{\text{state-level component}} + \underbrace{\widetilde{\mathbf{J}}_t + \widetilde{\mathbf{Q}}_t \mathbf{x}^*_{t-1} + \widetilde{\mathbf{G}}_t \epsilon^*_t}_{\text{aggregate component}}. \tag{12}$$

The coefficient matrices that appear in the aggregate component, $\widetilde{\mathbf{J}}_t$, $\widetilde{\mathbf{Q}}_t$, and $\widetilde{\mathbf{G}}_t$, are time-varying because of the non-linearities induced by the ZLB. The vector $\mathbf{x}^*_t$ which contains the aggregate variables evolves as:

$$\mathbf{x}^*_t = \mathbf{J}^*_t + \mathbf{Q}^*_t \mathbf{x}^*_{t-1} + \mathbf{G}^*_t \epsilon^*_t. \tag{13}$$

Here, ϵ^*_t are the aggregate shocks. Given this structure of our model, letting $\bar{\mathbf{x}}^*_t = \int \mathbf{x}_t(s)\mathrm{d}s$ denote the economy-wide average of the island-level variables, the deviation of island-level variables from their economy-wide averages, $\hat{\mathbf{x}}_t(s) = \mathbf{x}_t(s) - \bar{\mathbf{x}}^*_t$, is a time-invariant function of island-level variables alone:

$$\hat{\mathbf{x}}_t(s) = \mathbf{Q}\hat{\mathbf{x}}_{t-1}(s) + \mathbf{G}\epsilon_t(s), \tag{14}$$

where we use the assumption that island-level shocks have zero mean in the aggregate, that is, $\int \epsilon_t(s)\mathrm{d}s = 0$. We make explicit also that a key assumption we make in (12) in order to arrive at (14) is that the parameters across states are the same (that is, that the coefficient matrices $\mathbf{Q}$ and $\mathbf{G}$ for the state-level components are not state-specific).

The use of deviations of state-level observables from aggregates in estimation is crucial for our study. This is because by removing the dependence of state-level outcomes on aggregate variables, the nominal interest rate drops out from the reduced form just as it did in the reduced form analysis of Section 3.3 that led to specification (10). Equation (14) therefore circumvents, as (10) did, the problem of having to rely on aggregate data to estimate the Phillips curve in the presence of endogenous and possibly time-varying policy at the aggregate level.[16] This argument mirrors the one made in the reduced form analysis in Section 3.3, where subtracting aggregate optimal policy from the solution for state-level inflation removes aggregate quantities.[17]

In the particular case in which consumption is composed only of tradable goods ($\omega = 0$), the final goods price (11) – and therefore inflation – is the same in every state, and the deviation from the aggregate is equal to zero in every state. In this case, even with local state shocks moving the output gap, a representation like (14) would fail to identify the Calvo price parameter, as there would be no relative variation in state-level inflation data.

[16]Another advantage of representation (14) is that we can overcome the curse of dimensionality associated with all 51 states' dependence on the time-varying aggregate structure, which would otherwise make our estimation computationally infeasible.

[17]More formal arguments can be found in the literature. As mentioned in Section 2, Haldane and Quah (1999) were the first to show that endogenous policy leads to biases in estimating New Keynesian models. A simple and very elegant argument is presented in McLeay and Tenreyro (2020).

The Online Appendix contains all remaining details of our Bayesian estimation procedure, including the construction of the model's likelihood function and the priors of the estimated parameters.

Data We estimate the model separately using state-level data and aggregate data. We use a panel of employment, nominal output, wages, and inflation in the cross section of 51 US states from 1977 to 2017.[18] We use aggregate data from 1977 to 2015 on employment, output, wages, inflation, and the Fed Funds rate.[19] We construct these data in a similar way to the state-level data. We also use sequence of expected durations of the ZLB between 2009 and 2015 from the Blue Chip Financial Forecasts survey from 2009 to 2010 and the New York Federal Reserve's Survey of Primary Dealers from 2011 to 2015 (see Kulish, Morley and Robinson, 2017).

4.3 Estimation Results

The key objects of the estimated structural model that we focus on are the two Calvo parameters. We thus discuss our results regarding λ_p and λ_w first. This formal statistical analysis allows us to discuss the extent to which the parameters of interest are statistically stable over time. However, in order to get a sense of the extent to which any statistical difference brings about relevant economic differences, we also discuss the implications of our results regarding two transformations of the Calvo parameters. The first is to convert the Calvo parameters into slopes of the corresponding price and wage Phillips curves. This is important, since those slopes are the relevant objects governing the dynamics of the system. The second is to convert the Calvo parameters into frequency of price changes by firms and wage changes by unions in the

[18] See the Online Appendix for details of data availability across states and time, how we construct our series, and for details of how we use data series observed at different frequencies. For the robustness check where we include housing and household debt, we extend the set of observables to household debt and house prices in robustness checks. In this case, we can only estimate the model from 1999 to 2017, given data availability.

[19] We extend the sample to 1965 onwards in robustness exercises reported in Table 6 and in the Online Appendix. We also extend the set of observables to include household debt and house prices, as reported in the Online Appendix.

Table 2: Posterior Distributions, Relative State Data Only

	1977 to 2017			1977 to 1998			1999 to 2015		
Parameter	Mode	5%	95%	Mode	5%	95%	Mode	5%	95%
λ_p	0.59	0.57	0.61	0.65	0.63	0.67	0.64	0.62	0.65
λ_w	0.41	0.39	0.43	0.58	0.56	0.61	0.47	0.44	0.49

model. This not only provides us with an alternative metric but also allows us to compare our implied estimates with the micro estimates found in the literature.

In light of the previous discussion, we first report in Table 2 the posterior distributions of the Calvo parameters λ_p and λ_w estimated using state-level data only. The remaining structural parameters for all estimations are reported in the Online Appendix, including all prior specifications. The first panel of Table 2 reports the results of the estimation for the entire sample, 1977 to 2017. We find that the Calvo parameter for prices is 0.59 at the posterior mode, and the Calvo parameter for wages is 0.41 at the posterior mode. The posterior distributions for both parameters are very tight around their respective modes, with 90% of the mass concentrated in barely 4 basis points.

The second and third panels of Table 2 report the results for two sub-samples, the first covering the 1977 to 1998 period and the second covering the 1999 to 2015 period.[20] As the table makes clear, the estimates for the Calvo price parameter are very close to each other and to the estimate for the overall sample. Both of them are also tightly estimated, with a 90% probability interval of 4 and 3 basis points, respectively. The estimates for the Calvo wage parameter present signs of instability. The estimate for the second sub-sample is close to the estimate for the overall sample and also very precisely estimated

[20]The natural way would be to split the sample equally, choosing 1997 as the break year. However, we will check the robustness of the estimates to a model that additionally uses household debt during the buildup and subsequent bust around the financial crisis, as emphasized in Jones, Midrigan and Philippon (2022). As the debt data at the state level start in 1999, we chose to start the second sub-sample in that year. We also choose to end the second subsample in 2015, the same ending date as in our aggregate sample. We detrend the state-level data separately over each corresponding subsample.

– a 90% probability interval of 5 basis points. However, the estimate for the first sub-sample (0.56) is higher than the estimate for the overall sample (0.41), with a probability interval of 5 basis points.

Table 3 shows the Calvo parameters of the same model estimated of aggregate data alone. We also report the estimated Taylor rule parameters. In estimating the model with aggregate data, there is no reason to restrict the estimation to a start date in 1977. However, in order to make a comparison of the results with the ones in Table 2, we use the exact same periods as in there. We explore and report a larger sample period for the aggregate data estimation below.

Before turning to the discussion of the estimated Calvo parameters, notice that the estimated coefficients of the Taylor rule vary substantially across the two sub-periods. How these different policy regimes may affect the estimates is discussed below.

Regarding the values for the Calvo parameters over the full sample, note first that the difference with the ones estimated using state-level data is striking: the mode of the Calvo price parameter is 0.92 (compared with 0.59 in Table 2), while for the Calvo wage parameter, the mode is 0.84 (compared with 0.41 in Table 2).[21]

The sample size of the aggregate data is substantially shorter than the size of the panel used in the state-level analysis. In spite of that, the Calvo price parameter is quite precisely estimated, with a 90% probability band of 4 basis points. The case of the wage Calvo parameter is slightly less precise, with a corresponding value of 8 basis points. In comparing the differences between the estimates of the two different sub-samples we see differences (8 basis points for the Calvo price and 5 basis points for the Calvo wage parameter), but they are smaller than those for the Calvo wage parameter in using state-level data (11 basis points).

These rather small differences in the estimated Calvo parameters across

[21]The finding that wages are more flexible at the state level compared with the aggregate-level has already been pointed out in Beraja, Hurst and Ospina (2019) and in Jones, Midrigan and Philippon (2022). We find that observation applies also to prices.

Table 3: Posterior Distributions, Aggregate Data Only

	1977 to 2015			1977 to 1998			1999 to 2015		
Parameter	Mode	5%	95%	Mode	5%	95%	Mode	5%	95%
				Calvo Parameters					
λ_p	0.92	0.90	0.94	0.85	0.79	0.90	0.93	0.90	0.94
λ_w	0.84	0.80	0.88	0.90	0.87	0.93	0.85	0.80	0.89
				Taylor Rule Parameters					
α_r	0.80	0.73	0.84	0.63	0.31	0.78	0.81	0.73	0.87
α_p	2.38	1.98	2.86	2.03	1.51	2.98	1.62	1.15	2.76
α_x	0.46	0.36	0.62	1.92	0.91	1.95	0.17	0.12	0.23
α_y	0.28	0.21	0.37	0.07	0.01	0.24	0.24	0.19	0.33

the two sub-periods using aggregate data mask much larger differences in the implied slopes of the Phillips curves, which have been the elasticities focused on in the literature. Just as in standard New Keynesian models, the slope of the Phillips curve in our model is a non-linear function of the Calvo parameter, given by

$$\text{slope}_k = \frac{(1-\beta\lambda_k)(1-\lambda_k)}{\lambda_k}, \quad k \in \{p, w\}. \tag{15}$$

A quick inspection of (15) reveals that a change in λ_k from 0.9 to 0.95, say, implies a more drastic change in the Phillips curve slope than a change in λ_k from, say, 0.6 to 0.65.

With this non-linearity in mind, we map the implied Calvo price and wage estimates to the slopes of the Phillips curves in Table 4 to get a sense of what our estimates for the Calvo price and wage parameters imply for the slopes of their respective Phillips curves.[22] As expected, the implied slopes vary considerably depending on whether we use the state-level estimates or the aggregate ones. Our state-level estimate for the whole sample of λ_p implies a slope of 0.26. The aggregate estimates give a much flatter slope, closer to

[22] The slope of the curves may involve other parameters from preferences or technology. But the term (15) is typically found in the formulas for the slopes (see Galí, 2008).

0.01, consistent with New Keynesian models estimated with aggregate data in the literature.

The slope of the price Phillips curve implied by our whole state-level sample estimate of the Calvo price parameter, very close to 0.3, is statistically indistinguishable from the estimate of the preferred specification in McLeay and Tenreyro (2020). The point estimate they report is 0.379 with a standard deviation of 0.052 (see column 4 of Table 3 on page 273). They also use MSA-level data, but they use a limited information approach, a somewhat different sample, and different observables than we do. Our estimate of a relatively steep Phillips curve slope is also close to the findings of Barnichon and Mesters (2020) who employ monetary shocks as instrumental variables in an alternative limited information approach. In the structural estimation literature, our state-level estimate is essentially the same as the calibrated value used by Martin and Philippon (2017) who find, in the context of a New Keynesian model, that a Phillips curve slope coefficient of 0.3 works well in capturing relative changes in macroeconomic variables in the eurozone.

However, as we emphasized in Section 2, our estimated slopes are substantially higher than the ones reported in Hazell, Herreno, Nakamura and Steinsson (2020), who use a very different approach than ours.

But the key finding we want to emphasize is how the estimates of the implied slope of the Phillips curves change across sub-periods. As expected from the previous discussion, there are no substantial differences across subperiods in the estimation of the slopes for the price Calvo parameters using state-level data. But there are major differences using aggregate data. For the case of the wage Phillips curve, there are detectable differences in the implied slope using the state-level estimates. But the differences relative to the estimated slope using the whole sample are larger when using aggregate data.

This is most apparent in Figure 3, which plots the posterior distribution of the slopes implied by the posterior distribution of Calvo parameters for two sub-samples, but they are normalized to the full sample mode to aid the comparison. The distribution of Phillips curve slopes is not only significantly wider using the estimates coming from aggregate data but also vary widely

Table 4: Implied Slopes of Phillips Curve at Baseline Estimates

	1977 to 2015	1977 to 1998	1999 to 2015
A. State-Level Estimates			
Prices*	0.281	0.185	0.210
Wages†	0.874	0.308	0.597

B. Aggregate-Level Estimates

	1977 to 2017	1977 to 1998	1999 to 2015
Prices*	0.007	0.026	0.006
Wages†	0.031	0.011	0.027

*: Price Phillips curve slope is $(1-\beta\lambda_p)(1-\lambda_p)/\lambda_p$
†: Wage Phillips curve slope is $(1-\beta\lambda_w)(1-\lambda_w)/\lambda_w$

across periods.

In the case of the wage slope estimated of state-level data (bottom left panel of in Figure 3), although the distributions suggest statistically different slopes across periods, the difference is of relatively little economic significance.

To see this in a different metric, note that the Calvo parameters governing nominal rigidities in our model have a precise interpretation: the timing of price and wage adjustments are time dependent, with an average contract duration of $1/(1-\lambda_k), k \in \{p, w\}$. Thus, at the mode, these different slopes in the wage Phillips curve correspond to a frequency of wage adjustment of 2.4 quarters for the 1977 to 1998 sample and 1.9 quarters for the 1999 to 2015 sample. For the comparable estimates on aggregate data, the frequency of wage adjustment is around 10 quarters for the 1977 to 1998 sample but 6.6 quarters for the 1999 to 2015 sample. In Table 5, we present a full analysis of the mapping between Calvo parameters and frequency of price and wage changes for our estimates in Tables 2 and 3.

Table 5 highlights the close match between our state-level estimates and existing micro evidence on the frequency of price and wage changes. Because of the importance of price stickiness for aggregate dynamics, a large literature

Figure 3: Distributions of Phillips Curve Slopes

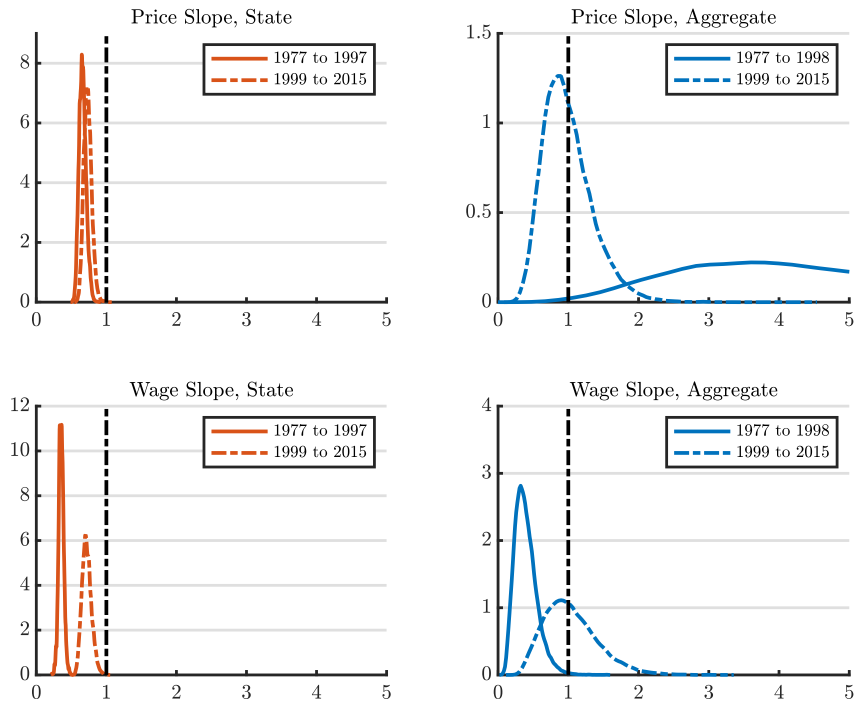

Notes: Each sub period posterior distribution of slopes is normalized by the mode of the full sample slope.

has developed that uses micro evidence to shed light on the frequency of price and wage adjustments and thus λ_p and λ_w. Our estimates are surprisingly close to those reported in these studies. For instance, Nakamura and Steinsson (2008) find average price durations of about 7 to 9 months, while our range of estimates of between 0.62 and 0.67 for the Calvo price parameter λ_p over the subsamples implies average durations between 7.9 to 9.1 months. For wages, Bihan, Montornes and Heckel (2012) find that the mean duration of a wage spell is just over 2 quarters or 6 months, using firm-level data from France. Our range of estimates, depending on the sample, of between 0.39 and 0.61 for the Calvo wage parameter λ_w implies an average duration of a wage contract of about 1.6 quarters (or just under 5 months) to 2.6 quarters (about 7.7 months).

The large differences in the distributions of the slope that emerge when

Table 5: Average Contract Durations Implied by Calvo Parameters

	State-Level			Aggregate-Level		
	Mode	5%	95%	Mode	5%	95%
A. Price Contracts (Quarters)						
Full Sample	2.4	2.3	2.5	12.7	9.6	17.6
1977 to 1998	2.9	2.7	3.0	6.8	4.8	9.6
1999 to 2017	2.7	2.6	2.9	13.7	10.2	18.1
B. Wage Contracts (Quarters)						
Full Sample	1.7	1.6	1.8	6.3	5.0	8.5
1977 to 1998	2.4	2.3	2.6	10.5	7.7	14.7
1999 to 2017	1.9	1.8	2.0	6.6	5.0	8.9

relying on aggregate data reflect changes in the monetary policy regime, according to our interpretation of the results presented so far. These differences are therefore consistent with the evidence provided in Section 3: while the reduced form parameter on state-level data was invariant to the sub-periods used for the estimation, the slopes implied by the estimates using aggregate data that depends on the policy rule changed over time. The structural estimation, however, allows us to move beyond those qualitative statements and evaluate the quantitative relevance of the key conceptual point raised by Haldane and Quah (1999): that endogenous changes in the policy regime blur the ability to estimate the structural parameters using aggregate data.

In order to do so, we show the results obtained from two exercises. In the first, we use the fact that the estimated Taylor rule parameters α_r, α_p, α_x, and α_y vary widely across the two sub-samples, as shown in Table 3. For instance, we find that the weight on the growth rate of potential output is highest in the first sub-sample of 1977 to 1998, while the weight on inflation deviations is smallest over the second sub-sample (which includes the zero lower bound period).

With this fact in mind, we repeat the estimation using aggregate data only over the full sample of 1977 to 2015, comparable with the first panel in Table 3.

But rather than jointly estimating the Taylor rule, we fix its parameters at the sub-sample estimates from Table 3. Thus, we estimate the Calvo parameters for the whole sample but fix the Taylor rule at the values estimated for the 1977 to 1998 sub-sample, as reported in the second panel of Table 3 (that is, $\alpha_r = 0.63$, $\alpha_p = 2.03$, $\alpha_x = 1.92$, and $\alpha_y = 0.07$). Then, we repeat the same estimation but fix the parameters of the Taylor rule at the values estimated for the 1999 to 2015 subsample (that is, $\alpha_r = 0.81$, $\alpha_p = 1.62$, $\alpha_x = 0.17$, and $\alpha_y = 0.24$).

These results are in Panel A of Table 6. The first column reports the estimated Calvo parameters when the Taylor rule is estimated for the full sample. These are the same as the ones reported in the first column of Table 3. We added them to aid the comparison. To avoid clutter, we also chose not to report the confidence intervals as they are similar to what was reported so far and the full results can be found in the Appendix. The second column reports the estimates when the Taylor rule is fixed at the estimated values of the first sub-period. The third column reports the estimates when fixing the Taylor rule parameters at the estimated values of the second sub-period.

In our second and final exercise, we repeat the estimation using aggregate data, but without restricting the sample period to coincide with the state-level data. The motivation to do so is the presumption that the period of increasing inflation and subsequent stabilization that the US experienced starting in the mid '60s and ending in the mid '80s was a different policy regime than the one that followed after the Volcker stabilization. That presumption leads us to estimate the model for the 1965-2005 period as well as for the sub-periods that are obtained by dividing the sample in 1985, much in the spirit of the results reported in Table 3, but without restricting the estimation to be over the same sample period than with the state-level data exercise. The results are reported in Panel B of Table 6. The bottom panel shows the estimated values for the policy rule and confirms the presumption of large variations across sub-periods.

Again, there is substantial variation over sub-periods in the estimated values for the Calvo parameters. The implications for the estimated slopes of the

Table 6: Mode of Posterior Distributions, Interaction With Policy Rules

A. Aggregate Data Only, Fixed Taylor Rule Parameters

Parameter	1977 to 2015*	1977 to 2015†	1977 to 2015‡
λ_p	0.92	0.89	0.92
λ_w	0.84	0.77	0.84

B. Aggregate Data Only, Policy Regime Periods

Parameter	1965 to 2005§	1965 to 1985§	1986 to 2005§
Calvo Parameters			
λ_p	0.86	0.72	0.92
λ_w	0.90	0.91	0.87
Taylor Rule Parameters			
α_r	0.95	0.95	0.90
α_p	5.62	6.44	2.92
α_x	0.47	0.59	0.22
α_y	0.89	0.73	0.27

$\star$: Estimated Taylor Rule with uniform priors
†: Taylor Rule parameters fixed at 1977 to 1998 estimates (see Table 3)
‡: Taylor Rule parameters fixed at 1999 to 2015 estimates (see Table 3)
§: No credit or house price series and no credit or housing preference shocks

corresponding Phillips curves are even more pronounced, which is consistent with these sub-samples capturing more clear policy regime changes (Figure 4). This figure is comparable to Figure 3 and illustrates the wide dispersion of implied slopes over the aggregate posterior distributions of λ_p and λ_w.

4.4 Regional Subsamples

One assumption up to this point has been that all states have been treated symmetrically. In particular, the parameters of the model are the same for all states. One concern with this assumption is that it implies that the response of state-level variables to shocks are the same, and that the variance of the

Figure 4: Distributions of Phillips Curve Slopes, Interaction with Policy

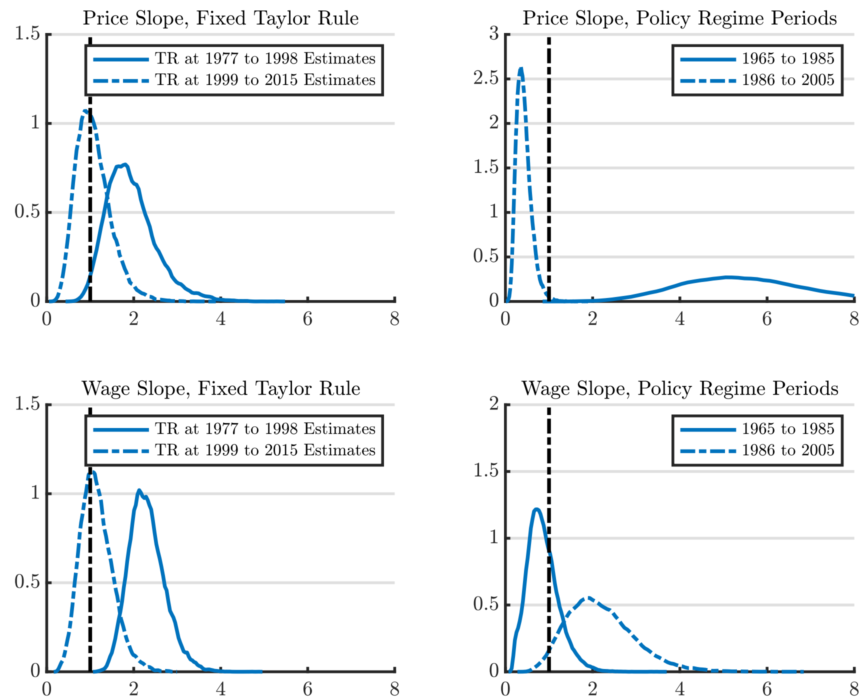

Notes: Each sub-period posterior distribution of slopes is normalized by the mode of the full sample slope.

shocks are the same for all shocks.[23]

In this section, to capture the possibility that heteroskedasticity may be impacting our inferences of the price and wage stickiness, we relax the assumption that all states have the same parameters and split the state-level data into two subsamples. Table 7 shows the posterior estimates for two subsamples of the state-level data. We order the states by their population size and first conduct an estimation using relative state-level data of the largest states, which we define as those with at least 5% of the US population each. This subgroup contains only four states, California, Florida, New York, and Texas, and in total comprise about one-third of the US population. We also conduct an estimation on the remaining 46 states, which together comprise the remaining two-thirds of the US population.

[23]We thank a referee for this comment.

Table 7: Posterior Distributions, State Groups, 1977 to 2017

	Largest States (32% Pop)			Smallest States (68% Pop)		
Parameter	Mode	5%	95%	Mode	5%	95%
λ_p	0.66	0.63	0.69	0.57	0.54	0.58
λ_w	0.69	0.64	0.72	0.40	0.39	0.43
ρ_z	0.98	0.96	0.99	0.98	0.98	0.99
ρ_n	0.64	0.57	0.72	0.96	0.95	0.97
ρ_b	0.96	0.94	0.97	0.96	0.96	0.97
ρ_z^N	0.95	0.92	0.97	0.97	0.96	0.97
σ_z	1.26	1.16	1.43	1.52	1.48	1.57
σ_n	0.40	0.24	0.57	0.03	0.03	0.04
σ_b	0.46	0.35	0.73	0.63	0.57	0.73
σ_z^N	1.12	0.98	1.30	1.44	1.37	1.49

The modal estimates of the Calvo price parameter are similar across the two subgroups (0.66 for the largest states, and 0.57 for the smallest states). The data from the larger states suggest wages are stickier than for the smaller states (and also compared to our full sample), with the posterior Calvo wage parameter estimated to be 0.69 at the mode and 0.40 for the smallest states. Translated into the average duration of contracts, however, these differences are somewhat small (9.7 months in the larger states versus 5 months in the smaller states). We find evidence that the volatility of the shocks differ across the two subsamples–for example, the volatility of TFP shocks in the largest states is about four-fifths that of the smallest states. Nonetheless, these differences in the estimated variances of the shock processes across states do not significantly influence our inferences about nominal rigidities and the estimated slopes of the Phillips curves.

5 Conclusion

The empirical literature on the stability of the Phillips curves has largely ignored the impact of endogenous monetary policy on empirical analysis of the Phillips curve. This omission has important implications: when policy is endogenous, aggregate data may be uninformative as to the existence of a stable relationship between unemployment and future inflation.

We show how regional data can be used to identify the structural relationship between unemployment and inflation. This insight guides our empirical strategy: we use city-level and state-level data from 1977 to 2017 and show that both the reduced form and the structural parameters of the Phillips curve are quite stable over time.

Our reduced form estimation is consistent with the notion that a 1 percent lower unemployment is associated with 30 basis points more of inflation, for constant values of aggregate variables, including policy.

Our structural estimation results are consistent with the assumption that prices change on average every $2^{1}/_{2}$ quarters, while wages change on average every 2 quarters. Our study is also consistent with the notion that these parameters can be safely assumed to be invariant to policy regime changes of the magnitude observed in the US since the mid '70s. These implications are in line with the findings in Alvarez, Beraja, Gonzalez-Rozada and Neumeyer (2018), which show that a model with exogenous Calvo frictions approximates very well an estimated menu-cost model, as long as inflation rates are not much higher than 10% a year.

Our analysis assumes symmetry across states. Through sub-sample analysis we have explored the sensitivity of parameter estimates across large and small states and did not find significant differences along that dimension. A full analysis along other dimensions of asymmetries, such as differences in demographic and occupational characteristics or in labor market frictions, while outside the scope of this paper, is a worthwhile avenue for future research.

References

Alvarez, Fernando, Martin Beraja, Martín Gonzalez-Rozada, and Pablo Andres Neumeyer, "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu Cost Models," *The Quarterly Journal of Economics*, 2018, *134* (1), 451–505.

Atkeson, Andrew and Lee E. Ohanian, "Are Phillips Curves Useful for Forecasting Inflation?," *Quarterly Review*, 2001, *25* (1), 2–11.

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

Barnichon, Regis and Geert Mesters, "Identifying Modern Macro Equations with Old Shocks*," *The Quarterly Journal of Economics*, 06 2020, *135* (4), 2255–2298.

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

Bihan, Herve Le, Jeremi Montornes, and Thomas Heckel, "Sticky Wages: Evidence from Quarterly Microeconomic Data," *American Economic Journal: Macroeconomics*, 2012, *4* (3), 1–32.

Blanchard, Olivier, "The Phillips Curve: Back to the '60s?," *American Economic Review*, 2016, *106* (5), 31–34.

Brainard, William C. and James Tobin, "Pitfalls in Financial Model-Building," Cowles Foundation Discussion Papers 244, Cowles Foundation for Research in Economics, Yale University 1968.

Carlstrom, Charles T., Timothy S. Fuerst, and Matthias Paustian, "Inflation Persistence, Monetary Policy, and the Great Moderation," *Journal of Money, Credit and Banking*, 2009, *41* (4), 767–786.

Cochrane, John H, "Determinacy and Identification with Taylor Rules," *Journal of Political Economy*, 2011, *119* (3), 565–615.

Coibion, Olivier and Yuriy Gorodnichenko, "Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 197–232.

Edge, Rochelle M. and Refet S. Gurkaynak, "How Useful Are Estimated DSGE Model Forecasts for Central Bankers?," *Brookings Papers on Economic Activity*, 2010, *41* (2), 209–259.

Fitzgerald, Terry and Juan Pablo Nicolini, "Is There a Stable Relationship between Unemployment and Future Inflation? Evidence from U.S. Cities," 2014.

Galí, Jordi, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Frankework*, Princeton University Press, 2008.

Gao, Han, Mariano Kulish, and Juan Pablo Nicolini, "Two Illustrations of the Quantity Theory of Money Reloaded," 2020.

Goldfeld, Stephen M. and Alan S. Blinder, "Some Implications of Endogenous Stabilization Policy," *Brookings Papers on Economic Activity*, 1972, *3* (3), 585–644.

Goodhart, C. A. E., *Money, Information and Uncertainty*, second edition ed., Macmillan International Higher Education, 1989.

Guerrieri, Luca and Matteo Iacoviello, "Occbin: A Toolkit to Solve Models with Occasionally Binding Constraints Easily," *Journal of Monetary Economics*, March 2015, *70*, 22–38.

Haldane, Andrew and Danny Quah, "UK Phillips curves and Monetary Policy," *Journal of Monetary Economics*, 1999, *44* (2), 259–278.

Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jon Steinsson, "The Slope of the Phillips Curve: Evidence from U.S. States," NBER Working Papers 28005, National Bureau of Economic Research, Inc 2020.

Hooper, Peter, Frederic S Mishkin, and Amir Sufi, "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or is It Just Hibernating?," Working Paper 25792, National Bureau of Economic Research 2019.

Jones, Callum, "Unanticipated Shocks and Forward Guidance at the Zero Lower Bound," 2017.

_ , Virgiliu Midrigan, and Thomas Philippon, "Household Leverage and the Recession," *Econometrica*, 2022, *90* (5), 2471–2505.

Kareken, John and Robert M. Solow, *Stabilization Policies*, Englewood Cliffs, New Jersey: Prentice-Hall, 1963.

Kiley, Michael T., "What Can the Data Tell Us About the Equilibrium Real Interest Rate?," Finance and Economics Discussion Series 2015-77, Board of Governors of the Federal Reserve System 2015.

Krugman, Paul, "Anchors Away (Slightly Wonkish)," The Conscience of a Liberal, New York Times December 4 2015.

Kulish, Mariano, James Morley, and Tim Robinson, "Estimating DSGE Models with Zero Interest Rate Policy," *Journal of Monetary Economics*, 2017, *88*, 35–49.

Leduc, Sylvain and Daniel J. Wilson, "Has the Wage Phillips Curve Gone Dormant?," 2017.

Levy, Antoine, "The Euro Area Phillips Curve: Identification with Regional Data," 2019.

Martin, Philippe and Thomas Philippon, "Inspecting the Mechanism: Leverage and the Great Recession in the Eurozone," *American Economic Review*, July 2017, *107* (7), 1904–37.

McLeay, Michael and Silvana Tenreyro, "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual 2019*, 2020, *34*, 199–255.

Mian, Atif and Amir Sufi, "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis," *American Economic Review*, 2011, *101* (5), 2132–2156.

_ **and** _ , "What Explains the 2007-2009 Drop in Employment?," *Quarterly Journal of Economics*, 2014, *82* (6), 2197–2223.

Mishkin, Frederic S., "Inflation Dynamics," *International Finance*, 2007, *10* (3), 317–334.

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

_ **and Jón Steinsson**, "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 2014, *104* (3), 753–792.

Nishizaki, Kenji and Tsutomu Watanabe, "Output-Inflation Trade-Off at Near-Zero Inflation Rates," *Journal of the Japanese and International Economies*, 2000, *14* (4), 304 – 326.

Peston, Maurice, "The Correlation between Targets and Instruments," *Economica*, 1972, *39* (156), 427–431.

Samuelson, Paul A. and Robert M. Solow, "Analytical Aspects of Anti-Inflation Policy," *The American Economic Review*, 1960, *50* (2), 177–194.

Sargent, Thomas J., *The Conquest of American Inflation*, Princeton University Press, 1999.

Stock, James and M. W. Watson, "Forecasting in Dynamic Factor Models Subject to Structural Instability," in Jennifer Castle and Neil Shepard, eds., *The Methodology and Practice of Econometrics: A Festschrift in Honor of David F. Hendry*, Oxford University Press, 2009, pp. 1–57.

Taylor, John B., "A Historical Analysis of Monetary Policy Rules," in John B. Taylor, ed., *Monetary Policy Rules*, University of Chicago Press, 1999, pp. 319–348.

Worswick, G. D. N., "Fiscal Policy and Stabilization in Britain," *Journal of Money, Credit and Banking*, 1969, *1* (3), 474–495.