

Has the U.S. Wage Phillips Curve Flattened? A Semi-Structural Exploration*

Jordi Galí Luca Gambetti

October 2018

Abstract

Estimates of a conventional wage Phillips curve for the U.S. economy point to a decline in its slope coefficient in recent years, as well as a shrinking role of lagged price inflation in the determination of wage inflation. We provide estimates of a *conditional* wage Phillips curve, based on a structural decomposition of wage, price and unemployment data generated by a VAR with time varying coefficients identified by a combination of long-run and sign restrictions. Our estimates show that the key qualitative findings from the unconditional reduced form regressions also emerge in the conditional evidence, suggesting that they are not driven by endogeneity problems or possible changes in the importance of wage markup shocks, though such factors may have overstated some of the actual changes.

JEL Classification No.: E24, E31

Keywords: inflation targeting, sacrifice ratio, wage rigidities, time-varying structural vector-autoregressive models.

*Prepared for the XXII Conference of the Central Bank of Chile, Santiago, October 25-26 . Galí: CREI, Universitat Pompeu Fabra and Barcelona GSE. Gambetti: Collegio Carlo Alberto, UAB and Barcelona GSE.

1 Introduction

The deep and prolonged recession triggered by the global financial crisis of 2007-2009 led to a large increase in the unemployment rate in most advanced economies. Ten years later, at the time of writing this paper, the recession has long ended, and the subsequent recoveries have brought the unemployment rate to levels close to –and in some cases even below– those at the peak of the previous expansion. In the U.S the unemployment rate increased from 4.4 percent in May 2007 to 10 percent in November 2009. Since that peak was attained, the unemployment rate has decreased –albeit at a slower pace than in earlier recoveries– down to its current level of 3.9 percent. Both movements represent, respectively, the largest increase (5.6 percentage points) and the largest decrease (6.1 percentage points) in the unemployment rate experienced by the US economy during the postwar period. Despite those wide and persistent fluctuations in unemployment, inflation has remained surprisingly stable during the same period, as illustrated in Figure 1. The previous phenomenon, often referred to in the literature as the "twin puzzle," appears to be robust to the measure of inflation and economic slack used,¹ and has also been observed in other advanced economies.²

Not surprisingly, central banks around the world have sounded the alarm in the face of that development, and with good reason.³ For one, a flattening of the Phillips curve implies a larger sacrifice ratio, and the need for more extreme policy measures in order to eliminate deviations of inflation from target. Furthermore, an outright decoupling of inflation from indicators of economic slack would call into question the inflation targeting framework widely adopted by central banks over the past decades, since that framework hinges critically on the existence of a positive relation between inflation and the level of economic activity, given that it is only through its ability to influence the latter through an appropriate setting of the interest rate and other policy instruments that

¹See, e.g. Stock and Watson (2018)

²See, e.g., Ciccarelli and Osbat (2017) for euro area evidence.

³See, e.g. Constancio (2017).

central banks can aim at controlling inflation.

In the present paper we revisit a key link of the relation between prices and economic activity, namely, the relation between wage inflation and unemployment. That empirical relation, which was the focus of Phillips' (1958) original work, is widely perceived to be at the heart of the "twin puzzle." In particular, the failure of wage inflation to respond sufficiently to the tightening of the labor market in recent years is generally viewed as one of the main factors behind the extremely accommodating monetary policies at central banks like the Federal Reserve or the ECB. Figure 2 shows a scatterplot of wage inflation and the unemployment rate to illustrate that phenomenon. We use quarterly data for the period 1964Q1-2017Q4.⁴ As discussed in Galí (2011a), the absence of a clear inverse relation between the two variables over the full sample period can be attributed to the large changes in mean inflation experienced by the US economy in the '70s ("the Great Inflation") and early '80s (the Volcker disinflation). When we restrict ourselves to the '60s and the Great Moderation period, a clear negative relation between the two variables becomes noticeable. Interestingly, that relation can be seen to become nearly flat in the years of the financial crisis and the subsequent recovery, in a way consistent with the "twin puzzle" hypothesis for price inflation.

In the present paper we seek to accomplish two tasks. First, we document changes in the U.S. wage Phillips curve, using simple reduced form regressions. Secondly, and after discussing the limitations of such a reduced form approach, we provide estimates of a *conditional* wage Phillips curve, based on a structural decomposition of wage, price and unemployment data generated by a VAR with time-varying coefficients and identified by a combination of long-run and sign restrictions.

Our main findings can be summarized as follows. First, we provide evidence of a substantial decline in the estimated coefficients on both lagged inflation and unemployment in the U.S. wage Phillips curve. Secondly, our estimates

⁴The wage inflation measure is the (centered) year-on-year change in the (log) average hourly earnings of production and nonsupervisory workers in the private sector, from the Establishment Survey.

of conditional wage Phillips curves show that such evidence is not driven by endogeneity problems or possible changes in the relative importance of shocks, though such factors may have overstated some of the changes implied by the reduced form evidence. Finally, we show that the reduced sensitivity of wage inflation to unemployment is also reflected in the estimated changes in a dynamic multiplier statistic relating the time-varying joint responses of wage inflation and unemployment to different shocks.

The paper is organized as follows. Section 2 describes the reduced form estimates of a wage Phillips curve and discusses some of its limitations. Section 3 presents our structural VAR model and describes the "semi-structural" evidence based on it. Section 4 concludes.

2 Wage Inflation and Unemployment: Reduced Form Evidence

In the present section we provide some reduced form evidence on the changing relation between wage inflation and the unemployment rate in the US economy.

The starting point of our empirical analysis is the estimation of a baseline wage inflation Phillips curve given by:

$$\pi_t^w = \alpha + \gamma \bar{\pi}_{t-1}^p - \varphi u_t + \varepsilon_t \quad (1)$$

where $\pi_t^w \equiv 400(w_t - w_{t-1})$ is (annualized) quarterly wage inflation (with w_t denoting the log nominal wage), $\bar{\pi}_t^p$ is a measure of price inflation (also annualized), and u_t is the unemployment rate. A specification similar to (1) has often been proposed and used in empirical applications.⁵

Table 1A reports the OLS estimates of γ and φ in (1) for different sample periods, using our baseline specification. We use average hourly earnings of production and nonsupervisory workers in the private sector from the Establishment Survey to construct our wage inflation measure. Our baseline price inflation measure is the (annualized) quarterly rate of change in the GDP deflator. The civilian unemployment rate is our measure of unemployment.

⁵See e.g. Blanchard and Katz (1999)

Several observations stand out. First, for the sample period before the financial crisis, the coefficients on both lagged inflation and unemployment are highly significant and with the expected sign. In particular, the estimated slope coefficient suggests that an increase of 1 percentage point in the unemployment rate is associated with a reduction of (annualized) wage inflation of about 30 basis points, given price inflation.

We uncover substantial changes within the pre-crisis sample period, however. Thus, in the 1986Q1-2007Q2 subsample (roughly corresponding to the Great Moderation) the estimated inflation coefficient becomes much smaller (though still significant), while the negative effect of unemployment on wage inflation is estimated to be about twice as large (and significantly different from the pre-86 period, as reflected in the p -value reported).

Our estimates for the period of crisis and the subsequent recovery, shown in the bottom panel of the table, point to a large decline in the sensitivity of wage inflation to the unemployment rate, though the coefficient on the latter still remains significant. While the test of equality of that coefficient between the full pre-crisis and the post-crisis samples can only reject that hypothesis at the 10 percent level (see p -value 2), the equality with the "late" pre-crisis period is rejected with very low p -values (see p -value 2). In addition, it is worth noting that the estimated coefficient on lagged inflation becomes insignificant in the more recent subsample period, suggesting a further decline in the importance of price indexation in wage setting over the past decade.⁶

Figures 3A and 3B show the time-varying estimates of φ and γ , respectively, based on a rolling OLS regression with a 32-observation window. The estimates illustrate, in a flexible way, the evidence reported earlier, namely, the consecutive steepening (during the Great Moderation) and flattening (during the financial crisis and its aftermath) of the wage Phillips curve in the US, as well as the seeming irrelevance of lagged price inflation since the mid-1980s

The qualitative findings discussed above are, for the most part, robust to

⁶The finding of a reduced importance of lagged inflation since the 1980s is consistent with Blanchard's (2016) estimates of price inflation Phillips curves.

alternative specifications of equation (1), as shown in Tables 1B (using CPI inflation), 1C (using lagged unemployment), and 1D (using year-on-year lagged price inflation).

2.1 Shortcomings

The evidence reported above, based on OLS estimates of equation (1), has several shortcomings as a measure of the sensitivity of wage inflation to variations in unemployment. Firstly, the specification of (1), while frequently found in textbooks and empirical applications, is largely *ad-hoc*.⁷

Secondly, equation (1) is unlikely to represent a structural relation, invariant to policy regime shifts. In particular, its implication of a long-run tradeoff between inflation and unemployment (as long as $\gamma < 1$, in accordance with the estimates) is hard to defend on theoretical as well as on empirical grounds.

Finally, the assumption of orthogonality between the right hand side variables and the disturbance term in (1) that would justify the use of OLS is unlikely to be satisfied in practice. As argued in Galí (2011a) through the lens of the New Keynesian model with staggered wage setting, the error term of a wage equation like (1) is likely to capture shocks to natural wage markups.⁸ Estimated DSGE models suggest that (i) those shocks are far from negligible sources of macro fluctuations, and (ii) have significant effects on both price inflation and unemployment.⁹ This observation renders OLS estimates of γ and φ inconsistent. Furthermore, changes over time in the volatility of wage markup shocks could be a source of spurious changes in the OLS estimates of those coefficients across subsample periods, giving a misleading impression of a "structural change" in the response of wage inflation to unemployment.

Below we propose and implement an empirical framework that aims at assessing possible changes in the responsiveness of wage inflation to unemployment

⁷Galí (2011a) provides some possible microfoundations for such a specification based on the staggered wage-setting model of Erceg et al. (2000) augmented with partial indexation to (lagged) price inflation, but they are consistent with (1) only under the strong assumption of an exogenous AR(1) process for the unemployment rate and a constant desired wage markup.

⁸The natural wage markup is the gap between the average real wage and the marginal rate of substitution that would prevail under flexible wages.

⁹See, e.g. Smets and Wouters (2007) and Galí et al. (2012)

in a way that overcomes (at least in principle) some of the the limitations discussed above.

3 Wage Inflation and Unemployment: Semi-Structural Evidence

In the present section we describe our empirical approach to identifying the different components of unemployment, wage inflation and price inflation, based on a structural vector autoregression model with time-varying coefficients (TVC-SVAR). Our empirical model provides a flexible specification which allows for structural changes in the relation between wage inflation and unemployment, as well as other structural changes that the U.S. economy may have experienced over the sample period considered.¹⁰ In addition, our framework makes it possible to overcome the potential endogeneity problem discussed above. Finally, our approach allows us to estimate the joint dynamics of wage inflation and unemployment in response to monetary policy interventions, and to uncover any changes over time in those dynamics.

3.1 Empirical Model

Let $\mathbf{x}_t = [\Delta(y_t - n_t), \pi_t^w, \pi_t^p, u_t, i_t^L]$ where y_t is (log) GDP, n_t denotes (log) hours of all persons in the nonfarm business sector, i_t^L is the yield on 10-year Government bonds. Price inflation and wage inflation are now defined as quarterly log first-differences of the GDP deflator and wage earnings of production and nonsupervisory workers, respectively, i.e. $\pi_t^w \equiv w_t - w_{t-1}$ and $\pi_t^p \equiv p_t - p_{t-1}$. As above, u_t denotes the civilian unemployment rate. We use a long term interest to avoid problems related to the binding zero lower bound at the end of our sample. All data are quarterly. The sample period is 1964Q1-2017Q4.

We assume that the \mathbf{x}_t admits the TVC-VAR representation

$$\mathbf{x}_t = \mathbf{A}_{0,t} + \mathbf{A}_{1,t}\mathbf{x}_{t-1} + \mathbf{A}_{2,t}\mathbf{x}_{t-2} + \dots + \mathbf{A}_{p,t}\mathbf{x}_{t-p} + \mathbf{u}_t \quad (2)$$

¹⁰These may include the change in the cyclical behavior of productivity emphasized in Galí and Gambetti (2009), or the change in monetary policy starting with Paul Volcker's tenure at the Fed (e.g. Clarida, Galí and Gertler (2000))

where $\mathbf{A}_{0,t}$ is a vector of time-varying intercepts, $\mathbf{A}_{i,t}$, for $i = 1, \dots, p$ are matrices of time-varying coefficients, and \mathbf{u}_t is a Gaussian white noise vector process with time-varying covariance matrix $\mathbf{\Sigma}_t$. We assume the reduced form innovations \mathbf{u}_t are a time-varying linear transformation of the underlying structural shocks $\boldsymbol{\varepsilon}_t$ given by

$$\mathbf{u}_t \equiv \mathbf{Q}_t \boldsymbol{\varepsilon}_t \quad (3)$$

where $\mathbb{E}\{\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t'\} = I$ and $\mathbb{E}\{\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_{t-k}'\} = 0$ for all t and $k = 1, 2, 3, \dots$. It follows that $\mathbf{Q}_t \mathbf{Q}_t' = \mathbf{\Sigma}_t$. As described in the Appendix, our approach assumes all the time-varying coefficients follow random walks with independent innovations.

Estimation is carried out as in Del Negro and Primiceri (2013).¹¹ Estimates of (2) can be used to obtain the (local) reduced form moving average (MA) representation:

$$\mathbf{x}_t = \boldsymbol{\mu}_t + \mathbf{B}_t(L) \mathbf{u}_t$$

Equation (3) can then be used to recover the structural (local) TVC-MA representation:

$$\mathbf{x}_t = \boldsymbol{\mu}_t + \mathbf{C}_t(L) \boldsymbol{\varepsilon}_t$$

where $\mathbf{C}_t(L) \equiv \mathbf{B}_t(L) \mathbf{Q}_t$, and where $\mathbf{x}_t^{ij} \equiv \mathbf{C}_t^{ij}(L) \boldsymbol{\varepsilon}_t^j$ represents the component of the i th variable associated with the j th shock. Determination of \mathbf{Q}_t requires a set of assumptions to identify the different shocks (i.e. the different elements of $\boldsymbol{\varepsilon}_t$) driving fluctuations in \mathbf{x}_t .

We identify the technology shock, following Galí (1999), as the only shock in vector $\boldsymbol{\varepsilon}_t$ to have a long run effect on labor productivity, implemented by imposing $\mathbf{C}_t^{11}(1) = 0$ for all t . In addition to a technology shock, we assume the existence of four additional shocks typically found in estimated DSGE models: (non-monetary) demand shocks, monetary policy shocks, price markup shocks, and wage markup shocks. By construction, those four shocks are restricted to have only transitory effects on labor productivity. We disentangle them through restrictions on the sign of their implied comovements between certain

¹¹We refer the reader to Galí and Gambetti (2016) for details.

variables over a four-quarter horizon after each shock.¹² Our sign restrictions are motivated by the predictions of the estimated medium-scale New Keynesian model with unemployment in Galí et al. (2012).¹³ Here are our short-run sign restrictions:

- Demand shocks (to be understood as non-monetary) are assumed to generate a positive comovement among y_t , π_t^p and i_t^L .
- Monetary policy shocks imply a positive comovement between y_t and π_t^p , but a negative comovement between each of those variables and i_t^L .
- Price markup shocks are identified as the only source of fluctuations that generates a positive comovement between π_t^p and the price markup $\mu_t^p \equiv (y_t - n_t) - (w_t - p_t)$.
- Wage markup shocks are assumed to be the only structural disturbances that generate a positive comovement between π_t^p and the unemployment rate u_t , with the latter variable interpreted as a proxy for the wage markup, following Galí (2011a,b).

Table 1 summarizes our identification strategy in a compact way.

TABLE 2. Identification

	$y_t - n_t$	π_t^w	u_t	π_t^p	μ_t^p	i_t^L
<i>Technology</i>						
<i>Demand</i>	0_∞	+	-	+	-	+
<i>Monetary Policy</i>	0_∞	+	-	+	-	-
<i>Price Markup</i>	0_∞	+/-	-/+	+	+	
<i>Wage Markup</i>	0_∞	+	+	+/-	-/+	

¹²The use of sign restrictions for identification purposes in structural VARS was pioneered by Uhlig (2005)

¹³That model is itself an extension of those in Smets and Wouters (2007) that introduces an explicit relation between the unemployment rate and the wage markup discussed in Galí (2011b).

3.2 Conditional Wage Phillips Curves

The next step in our approach consists of re-estimating the wage Phillips curve (1) using the time series for wage inflation, price inflation and unemployment, *purged of the component associated with wage markup shocks* obtained using the TVC-SVAR described above. To the extent that the error term in (1) captures fluctuations in wage markup shocks, the estimation of such a *conditional* wage Phillips curve should overcome any bias resulting from the correlation between the error term and the regressors.¹⁴

Table 3A reports estimates of coefficients γ and φ in (1) for different sample periods, using the non wage markup components of the three variables involved. As in our TVC-SVAR specification, price inflation is measured as the log first-difference of the GDP deflator. As in our baseline estimates of (1) we annualize both inflation variables before applying OLS. The estimates for the full pre-crisis period point to a smaller price inflation coefficient (0.32) and a larger (in absolute terms) unemployment coefficient (-0.55) than the unconditional estimates of Table 1A. Interestingly, when we now restrict ourselves to the Great Moderation period, we still get a smaller inflation coefficient and larger unemployment coefficient than in the full pre-crisis period, but now the differences in the estimated coefficients are much smaller (and in case of the inflation coefficient, statistically insignificant). When we turn to the crisis and recovery period, we obtain estimates of the inflation and unemployment coefficients that are smaller (in absolute value) than in the pre-crisis period. As in our unconditional estimates of Table 1A for this period, the coefficient on lagged inflation is now insignificant. The coefficient on unemployment is significantly smaller than in the pre-crisis period, though more than twice as large as its unconditional counterpart.

Figures 4A and 4B reports time-varying estimates of φ and γ , respectively, based on rolling OLS regressions with a 32-observation window, applied to the

¹⁴Our approach is similar in spirit to that of Barnichon and Mesters (2018), who estimate a New Keynesian Phillips curve for price inflation using current and lagged monetary policy shocks as instruments.

non-wage markup components of the time series involved. Note that, relative to Figures 3A and 3B, and consistent with the evidence just discussed, the coefficient on unemployment appears to be more stable over time, and to experience a smaller decline in recent years. On the other hand, the conditional rolling estimates of the inflation coefficient display a pattern very similar to their unconditional counterparts, though with slightly lower values in the 1980s.

The previous qualitative findings on the estimates of conditional wage Phillips curves are largely robust to alternative specifications, as shown in Tables 3B (using lagged unemployment), and 3C (using year-on-year lagged price inflation). An exception to this similarity is given by the estimates of the coefficient on lagged year-on-year inflation in Table 3C, which do not appear to vary significantly across sample periods.

3.3 Conditional Dynamic Multipliers

The empirical approach described in the previous subsection should have overcome one of the shortcomings of the reduced form evidence, namely, the potential biases in the OLS estimates of equation (1) resulting from the endogeneity of unemployment and inflation with respect to wage markup shocks. Yet, the estimates of conditional wage Phillips curves are still subject to another important caveat, namely, the ad-hoc specification of (1). In the present subsection we uncover possible changes over time in the relation between unemployment and wage inflation without the straitjacket of any assumed functional relation between the two variables. Instead we focus on the estimated impulse responses generated by our TVC-SVAR and trace the evolution over time of the dynamic wage inflation-unemployment multiplier, defined as a ratio of the cumulative impulse responses of those two variables to a given shock ε_t^i at different horizons:

$$\Phi_t^i(k) \equiv \frac{\sum_{k=0} \frac{\partial \pi_{t+k}^w}{\partial \varepsilon_t^i}}{\sum_{k=0} \frac{\partial u_{t+k}}{\partial \varepsilon_t^i}}$$

for $k = 0, 1, 2, \dots, 8$.

Figure 5A displays the evolution of the above dynamic multiplier conditional

on monetary policy shocks. As expected, the multiplier is always negative, suggesting that a monetary shocks tends to move wage inflation and unemployment in opposite directions. We also see that its absolute value declines with the horizon, suggesting a more persistent effect on unemployment than on wage inflation. More interestingly, however, the (absolute) size of the multiplier appears to decrease over time. This is true at all the horizons considered (up to eight quarters) but particularly so at the shortest horizons. Thus, in the early part of the sample, we see how an expansionary monetary policy shock that drove the unemployment rate down by one percentage point, simultaneously raises (quarterly) wage inflation by about 3 percentage points, implying a multiplier of -3 . That short run multiplier decreases over time in absolute value, to a level close to -1 . That finding is consistent with the evidence in previous sections pointing to a more muted response of wage inflation to fluctuations in unemployment. In contrast with that evidence, however, the present estimates suggest that such a change in responsiveness has been quite gradual, having started well before the financial crisis.

Figure 5B displays analogous evidence for (non-monetary) demand shocks. Many of the qualitative patterns observed in Figure 5A are also present here including the gradual decline in the (absolute) size of the estimated multiplier at all horizons. A different picture, however, emerges in Figure 5C, which displays the dynamic multiplier for price markup shocks, and which does not suggest any major changes over time.

The above evidence rules out a change in the relative importance of different shocks as the main or only source of any reduction in the sensitivity of wage inflation to unemployment fluctuations: a change in that relation appears to have occurred even when one conditions on specific shocks. Unfortunately, our approach cannot shed direct light on the nature of the structural change(s) that may underly the lower conditional dynamic multipliers. A greater decoupling of wage inflation from price inflation, possibly due to a stronger anchoring of inflation expectations associated with the adoption of a price stability-oriented monetary policy, and captured in our estimates of the wage Phillips curve (both

conditional and unconditional) would limit the so-called second round effects on wage inflation and dampen the response of the latter variable to any given variation in the unemployment rate, providing a possible explanation to the evidence above.¹⁵

4 Conclusions

We have started the present paper by documenting the changes in the wage Phillips curve, using simple reduced form regressions applied to aggregate data. In particular, we have provided evidence of a substantial decline in the estimated coefficients on both lagged inflation and unemployment in our wage Phillips curve.

We have also provided estimates of *conditional* wage Phillips curves, based on a structural decomposition of wage, price and unemployment data generated by a VAR with time varying coefficients identified by a combination of long-run and sign restrictions. Our estimated conditional wage Phillips curves show that most qualitative findings from the reduced form evidence are not driven by endogeneity problems or possible changes in the relative importance of shocks, though such factors may have overstated some of the actual changes.

Finally, we have shown that the reduced sensitivity of wage inflation to unemployment is also reflected in the estimated changes in a dynamic multiplier statistic, based on the estimated time-varying impulse responses to monetary policy and demand shocks.

¹⁵Blanchard and Galí (2009) point to that mechanism as an explanation of the smaller macroeconomics effects of oil price shocks in the 2000s relative to the 1970s.

REFERENCES

- Barnichon, Régis and Geert Mesters (2018): "Identifying the Phillips Curve from Shifts in Demand," unpublished manuscript.
- Blanchard, Olivier and Lawrence Katz (1999): "Wage Dynamics: Reconciling Theory and Evidence," *American Economic Review*, Vol. 89, No. 2, pp. 69-74
- Blanchard, Olivier and Jordi Galí (2009): "The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s?" (with O. Blanchard) in J. Galí and M. Gertler (eds.), *International Dimensions of Monetary Policy*, University of Chicago Press (Chicago, IL), 2009, 373-428
- Blanchard, Olivier (2016): "The Phillips Curve: Back to the '60s?" *American Economic Review* 106(5), 31-34.
- Clarida, Richard, Jordi Galí and Mark Gertler (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics* 115(1), 147-180.
- Ciccarelli, M. and C. Osbat (2017), "Low inflation in the euro area: Causes and Consequences", ECB Occasional Paper Series, No 181.
- Coibion, Olivier and Yuriy Gorodnichenko (2015), "Is the Phillips curve Alive and Well after All? Inflation Expectations and the Missing Disinflation", *American Economic Journal: Macroeconomics*, Vol. 7(1), pp. 197-232.
- Constancio, Vitor (2017): "Understanding and overcoming low inflation," remarks, Frankfurt, 21-22 September 2017.
- Daly, Mary C. and Bart Hobijn (2014): "Downward Nominal Wage Rigidities Bend the Phillips Curve," *Journal of Money Credit and Banking* 46(2), 51-93.
- Del Negro, Marco and Giorgio Primiceri (2013): "Time-Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum" Federal Reserve Bank of New York Staff Report No. 619.
- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin (2000): "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics* vol. 46, no. 2, 281-314.

Gali, Jordi (2011a): "The Return of the Wage Phillips Curve," *Journal of the European Economic Association*, vol. 9, issue 3, 436-461.

Gali, Jordi (2011b): *Unemployment Fluctuations and Stabilization Policies: A New Keynesian Perspective*, MIT Press (Cambridge, MA).

Gali, Jordi, Frank Smets and Raf Wouters (2012): "Unemployment in an Estimated New Keynesian Model," *NBER Macroeconomics Annual 2011*, 329-360.

Gali, Jordi and Luca Gambetti (2009): "On the Sources of the Great Moderation," *American Economic Journal: Macroeconomics* 1(1), 26-57.

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.

Smets, Frank, and Rafael Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, vol 97, no. 3, 586-606.

Stock, James H. and Mark W. Watson (2018): "Slack and Cyclically Sensitive Inflation," presented at the ECB Forum on Central Banking.

Uhlig, Harald (2005): "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics* 52 (2), 381-419

APPENDIX

Let $\boldsymbol{\theta}_t = \text{vec}(\mathbf{A}'_t)$ where $\mathbf{A}_t = [\mathbf{A}_{0,t}, \mathbf{A}_{1,t}, \dots, \mathbf{A}_{p,t}]$ and $\text{vec}(\cdot)$ is the column stacking operator. We assume $\boldsymbol{\theta}_t$ evolves over time according to the following equation:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t \quad (4)$$

where $\boldsymbol{\omega}_t$ is Gaussian white noise vector process with covariance matrix $\boldsymbol{\Omega}$.

Time variation of $\boldsymbol{\Sigma}_t$ is modeled as follows. Let $\boldsymbol{\Sigma}_t = \mathbf{F}_t \mathbf{D}_t \mathbf{F}'_t$, where \mathbf{F}_t is lower triangular, with ones on the main diagonal, and \mathbf{D}_t a diagonal matrix. The vector containing the diagonal elements of $\mathbf{D}_t^{1/2}$, denoted by $\boldsymbol{\sigma}_t$, is assumed to evolve according to the process

$$\log \boldsymbol{\sigma}_t = \log \boldsymbol{\sigma}_{t-1} + \boldsymbol{\zeta}_t. \quad (5)$$

Moreover let $\boldsymbol{\phi}_{i,t}$ denote the column vector with the non-zero elements of the $(i+1)$ -th row of \mathbf{F}_t^{-1} . We assume

$$\boldsymbol{\phi}_{i,t} = \boldsymbol{\phi}_{i,t-1} + \boldsymbol{\nu}_{i,t} \quad (6)$$

where $\boldsymbol{\zeta}_t$ and $\boldsymbol{\nu}_{i,t}$ are Gaussian white noise vector processes with zero mean and (constant) covariance matrices $\boldsymbol{\Xi}$ and $\boldsymbol{\Psi}_i$, respectively. We further assume that $\boldsymbol{\nu}_{i,t}$ is independent of $\boldsymbol{\nu}_{j,t}$, for all $j \neq i$, and that $\boldsymbol{\omega}_t$, $\boldsymbol{\varepsilon}_t$, $\boldsymbol{\zeta}_t$ and $\boldsymbol{\nu}_{i,t}$ (for all i) are mutually independent.

Priors Specification

We make the following assumptions about prior distributions :

$$\begin{aligned} \boldsymbol{\theta}_0 &\sim N(\hat{\boldsymbol{\theta}}, 4\hat{\mathbf{V}}_{\boldsymbol{\theta}}) \\ \log \boldsymbol{\sigma}_0 &\sim N(\log \hat{\boldsymbol{\sigma}}_0, \mathbf{I}_n) \\ \boldsymbol{\phi}_{i0} &\sim N(\hat{\boldsymbol{\phi}}_i, \hat{\mathbf{V}}_{\boldsymbol{\phi}_i}) \\ \boldsymbol{\Omega}^{-1} &\sim W(\underline{\boldsymbol{\Omega}}^{-1}, \underline{\rho}_1) \\ \boldsymbol{\Xi}^{-1} &\sim W(\underline{\boldsymbol{\Xi}}^{-1}, \underline{\rho}_2) \\ \boldsymbol{\Psi}_i^{-1} &\sim W(\underline{\boldsymbol{\Psi}}_i^{-1}, \underline{\rho}_{3i}) \end{aligned}$$

where $W(\mathbf{S}, d)$ denotes a Wishart distribution with scale matrix \mathbf{S} and degrees of freedom d and \mathbf{I}_n is a $n \times n$ identity matrix, with n the number of variables in the VAR.

We use a time invariant VAR for \mathbf{x}_t estimated using the first $\tau = 64$ observations to calibrate prior means and variances. $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{V}}_{\boldsymbol{\theta}}$ are set equal to the OLS estimates. Let $\hat{\boldsymbol{\Sigma}}$ be the covariance matrix of the residuals $\hat{\mathbf{u}}_t$ of the initial time-invariant VAR. We apply the decomposition $\hat{\boldsymbol{\Sigma}} = \hat{\mathbf{F}}\hat{\mathbf{D}}\hat{\mathbf{F}}'$ and set $\log \hat{\boldsymbol{\sigma}}_0$ equal to the log of the diagonal elements of $\hat{\mathbf{D}}^{1/2}$. $\hat{\boldsymbol{\phi}}_i$ is set equal to the OLS estimates of the coefficients of the regression of $\hat{\mathbf{u}}_{i+1,t}$, the $i + 1$ -th element of $\hat{\mathbf{u}}_t$, on $-\hat{\mathbf{u}}_{1,t}, \dots, -\hat{\mathbf{u}}_{i,t}$ and $\hat{\mathbf{V}}_{\boldsymbol{\phi}_i}$ equal to the estimated variances.

The scale matrices are parametrized as follows: $\underline{\boldsymbol{\Omega}} = \underline{\rho}_1(\lambda_1 \hat{\mathbf{V}}_{\boldsymbol{\theta}})$, $\underline{\boldsymbol{\Xi}} = \underline{\rho}_2(\lambda_2 \mathbf{I}_n)$ and $\underline{\boldsymbol{\Psi}}_i = \underline{\rho}_{3i}(\lambda_3 \hat{\mathbf{V}}_{\boldsymbol{\phi}_i})$. The degrees of freedom $\underline{\rho}_1$ and $\underline{\rho}_2$ are set equal to the number of rows $\underline{\boldsymbol{\Omega}}^{-1}$ and \mathbf{I}_n plus one respectively while $\underline{\rho}_{3i}$ is $i + 1$ for $i = 1, \dots, n - 1$. Finally $\lambda_1 = 0.0002$, $\lambda_2 = 0.01$ and $\lambda_3 = 0.01$.

Gibbs sampling algorithm

The Gibbs sampling algorithm is identical to that described in the online appendix of Gali and Gambetti (2015). We use all the data points available from 1964:Q2 to 2017:Q4. We draw 50000 realizations discarding the first 40000 and then taking 1 out of 10 draws, therefore collecting a total of 1000 realizations.

Table 1A. Empirical Wage Phillips Curves			
Earnings, GDP Deflator (Q)			
	π_{t-1}^p	u_t	R^2
Pre-Crisis			
<i>1964Q1-2007Q2</i>	0.67 (0.04)	-0.29 (0.07)	0.57
<i>1986Q1-2007Q2</i>	0.17 (0.07)	-0.58 (0.07)	0.45
<i>p value</i>	0.003	0.016	
Crisis and Recovery			
<i>2007Q3-2017Q4</i>	0.01 (0.12)	-0.11 (0.04)	0.10
<i>p value 1</i>	0.002	0.102	
<i>p value 2</i>	0.056	0.001	

Table 1B. Empirical Wage Phillips Curves			
Earnings, CPI (Q)			
	π_{t-1}^p	u_t	R^2
Pre-Crisis			
<i>1964Q1-2007Q2</i>	0.45 (0.04)	-0.16 (0.08)	0.42
<i>1986Q1-2007Q2</i>	0.09 (0.04)	-0.57 (0.07)	0.45
<i>p value</i>	0.001	0.021	
Crisis and Recovery			
<i>2007Q3-2017Q4</i>	0.05 (0.04)	-0.11 (0.05)	0.14
<i>p value 1</i>	0.001	0.56	
<i>p value 2</i>	0.243	0.001	

Note:

Table 1C. Empirical Wage Phillips Curves
Earnings, GDP Deflator (Q), Lagged Unemployment

	π_{t-1}^p	u_{t-1}	R^2
Pre-Crisis			
<i>1964Q1-2007Q2</i>	0.65 (0.04)	-0.22 (0.07)	0.55
<i>1986Q1-2007Q2</i>	0.16 (0.07)	-0.56 (0.07)	0.43
<i>p value</i>	0.014	0.367	
Crisis and Recovery			
<i>2007Q3-2017Q4</i>	0.01 (0.11)	-0.15 (0.05)	0.18
<i>p value 1</i>	0.002	0.34	
<i>p value 2</i>	0.020	0.001	

Table 1D. Empirical Wage Phillips Curves			
Earnings, GDP Deflator (YOY)			
	π_{t-1}^p	u_t	R^2
Pre-Crisis			
<i>1964Q1-2007Q2</i>	0.80 (0.04)	-0.49 (0.06)	0.63
<i>1986Q1-2007Q2</i>	0.25 (0.08)	-0.61 (0.07)	0.47
<i>p value</i>	0.045	0.028	
Crisis and Recovery			
<i>2007Q3-2017Q4</i>	0.03 (0.19)	-0.11 (0.05)	0.10
<i>p value 1</i>	0.015	0.001	
<i>p value 2</i>	0.045	0.001	

Note:

Table 3A. Conditional Wage Phillips Curves			
Earnings, GDP Deflator (Q)			
	π_{t-1}^p	u_t	R^2
Pre-Crisis			
<i>1964Q1-2007Q2</i>	0.32 (0.04)	-0.55 (0.07)	0.39
<i>1986Q1-2007Q2</i>	0.21 (0.05)	-0.74 (0.06)	0.64
<i>p value</i>	0.226	0.019	
Crisis and Recovery			
<i>2007Q3-2017Q4</i>	0.07 (0.09)	-0.29 (0.05)	0.10
<i>p value 1</i>	0.24	0.04	
<i>p value 2</i>	0.36	0.001	

Table 3B. Conditional Wage Phillips Curves			
Earnings, GDP Deflator (Q), Lagged Unemployment			
	π_{t-1}^p	u_{t-1}	R^2
Pre-Crisis			
<i>1964Q1-2007Q2</i>	0.31 (0.05)	-0.40 (0.00)	0.31
<i>1986Q1-2007Q2</i>	0.20 (0.06)	-0.70 (0.07)	0.44
<i>p value</i>	0.199	0.009	
Crisis and Recovery			
<i>2007Q3-2017Q4</i>	0.07 (0.08)	-0.32 (0.05)	0.56
<i>p value 1</i>	0.28	0.46	
<i>p value 2</i>	0.34	0.001	

Table 3C. Conditional Wage Phillips Curves			
Earnings, GDP Deflator (YOY)			
	π_{t-1}^p	u_t	R^2
Pre-Crisis			
<i>1964Q1-2007Q2</i>	0.43 (0.05)	-0.65 (0.07)	0.47
<i>1986Q1-2007Q2</i>	0.33 (0.06)	-0.76 (0.06)	0.68
<i>p value</i>	0.420	0.026	
Crisis and Recovery			
<i>2007Q3-2017Q4</i>	0.32 (0.12)	-0.23 (0.05)	0.57
<i>p value 1</i>	0.665	0.001	
<i>p value 2</i>	0.96	0.001	

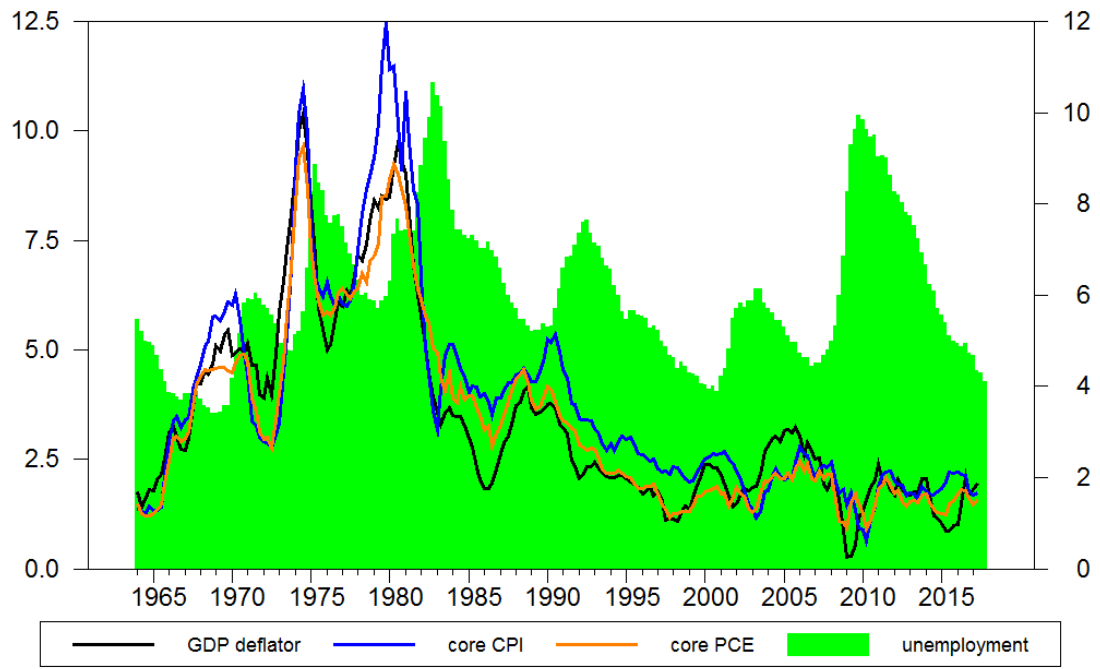


Figure 1. Unemployment and Price Inflation

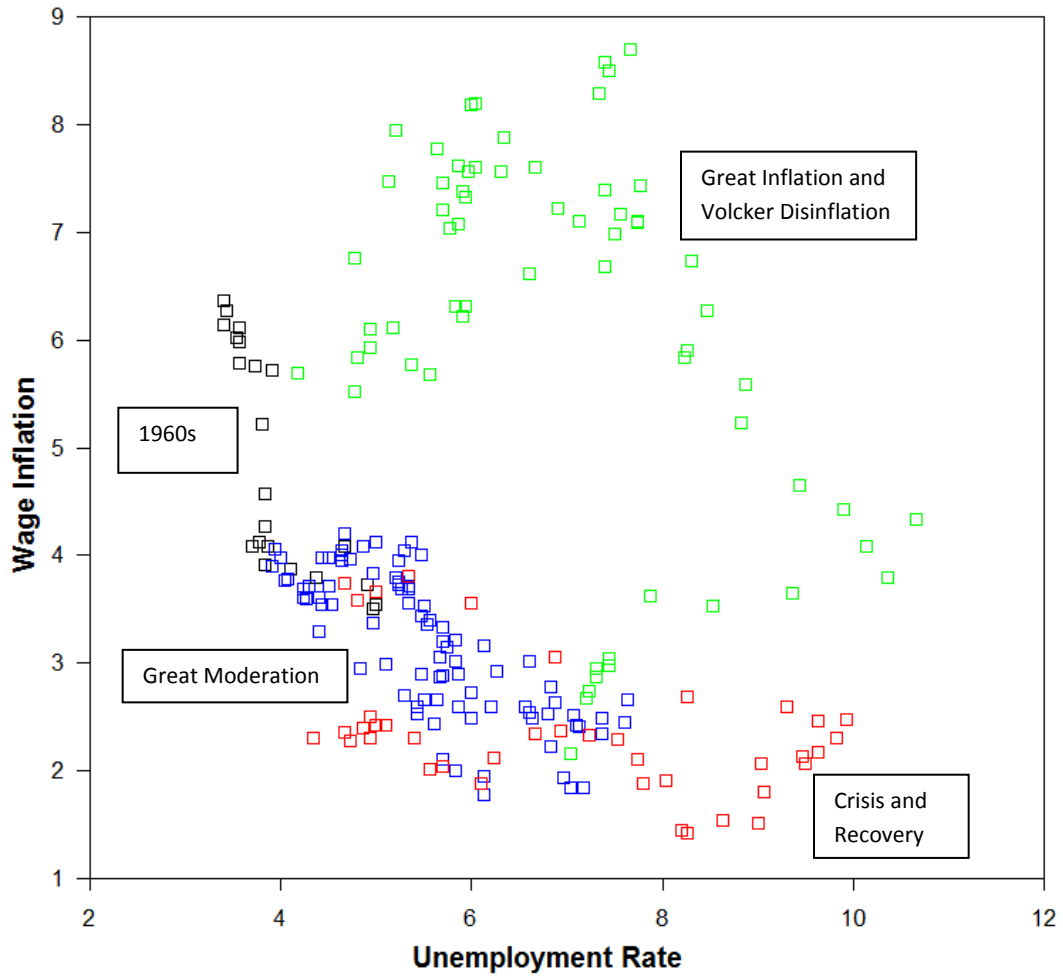


Figure 2. The U.S. Wage Phillips curve

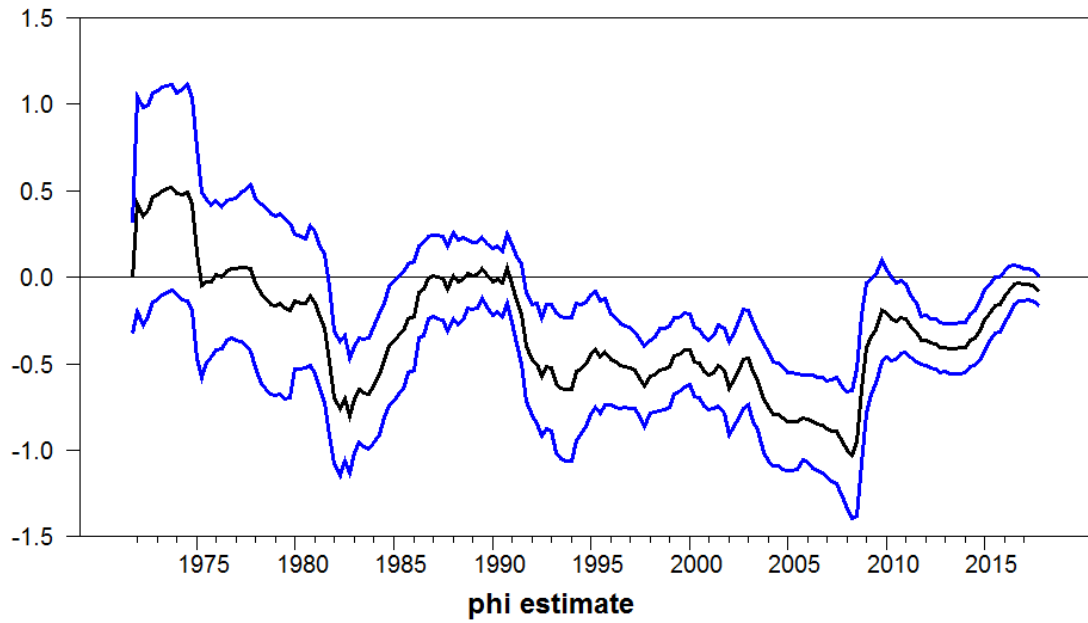


Figure 3A. Unemployment Coefficient (unconditional estimate)

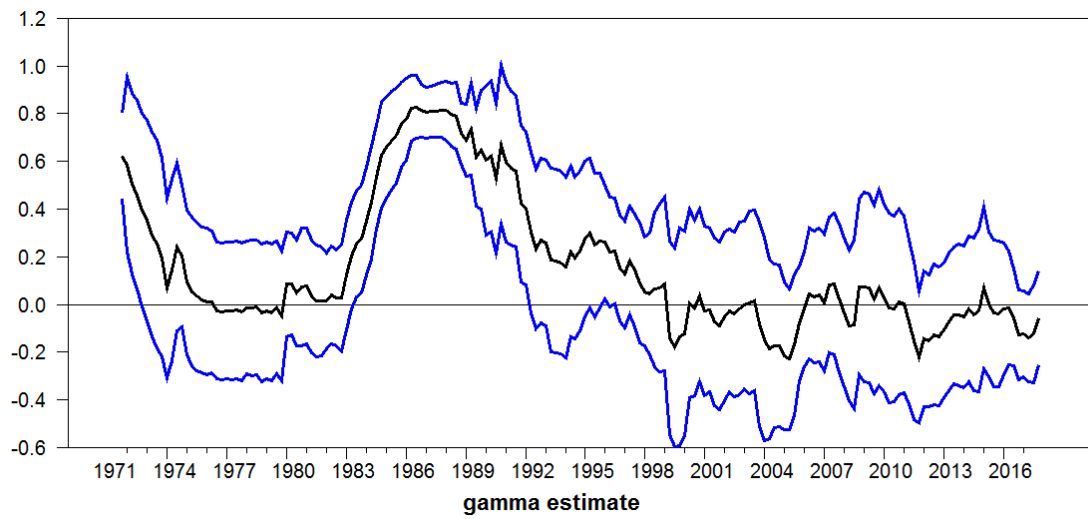


Figure 3B. Inflation Coefficient (unconditional estimate)

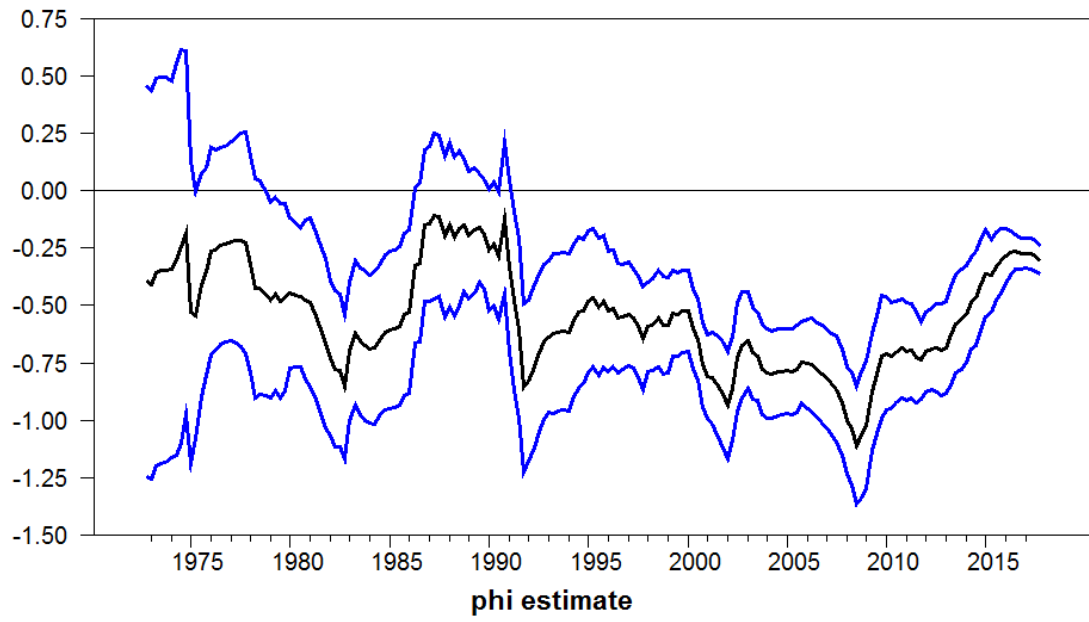


Figure 4A. Unemployment Coefficient (conditional estimate)

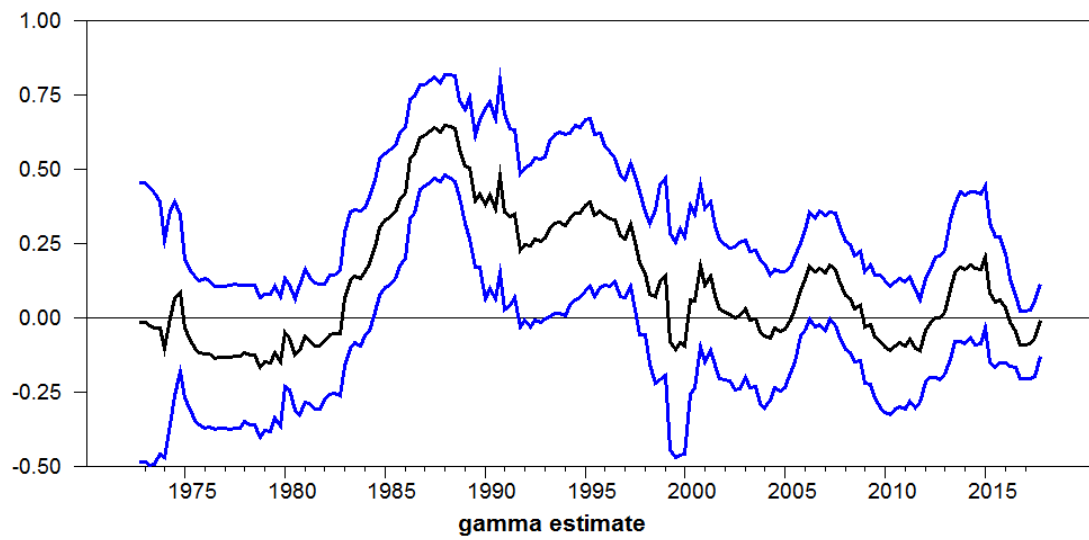


Figure 4B. Inflation Coefficient (conditional estimate)

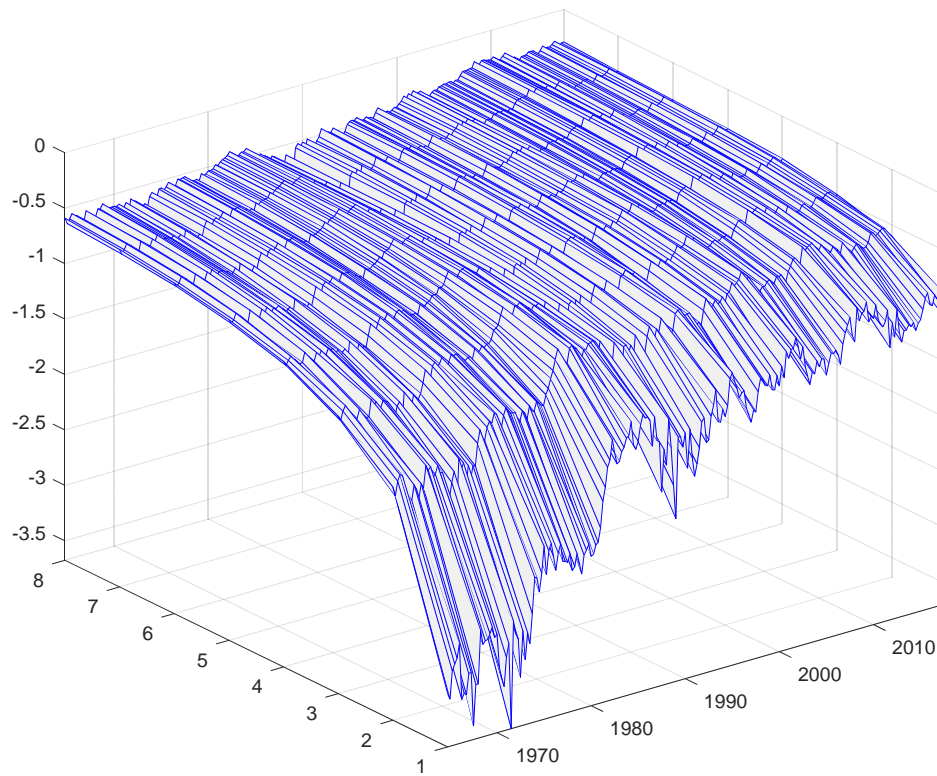


Figure 5A. Dynamic Multiplier: Monetary Policy Shocks

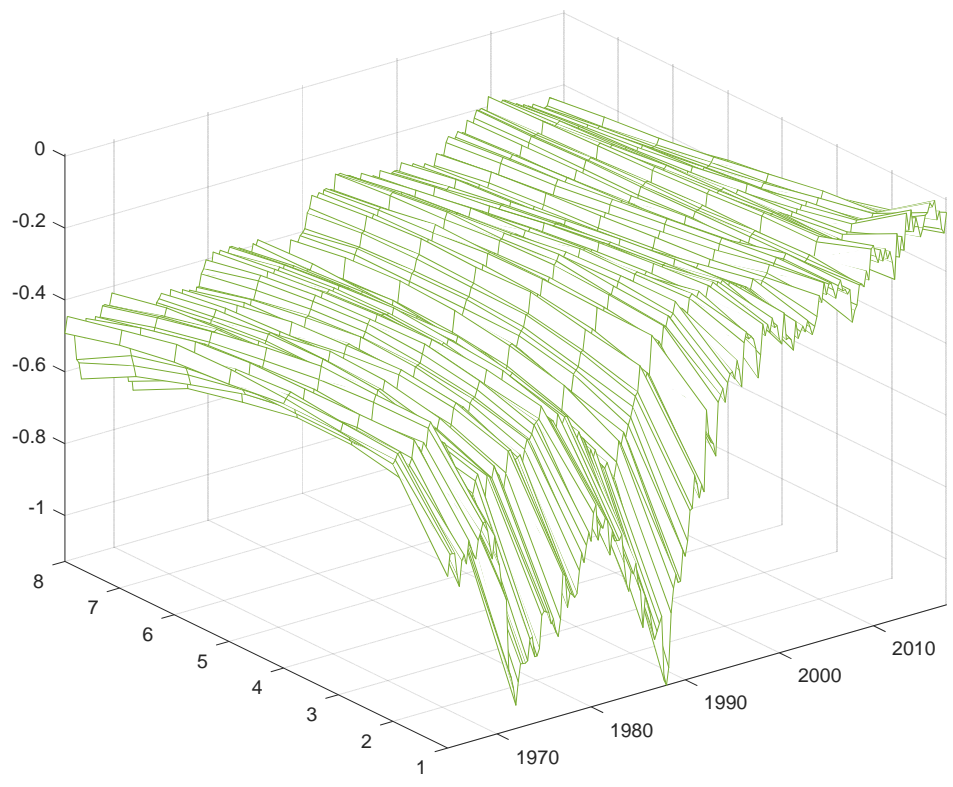


Figure 5B. Dynamic Multiplier: Demand Shocks

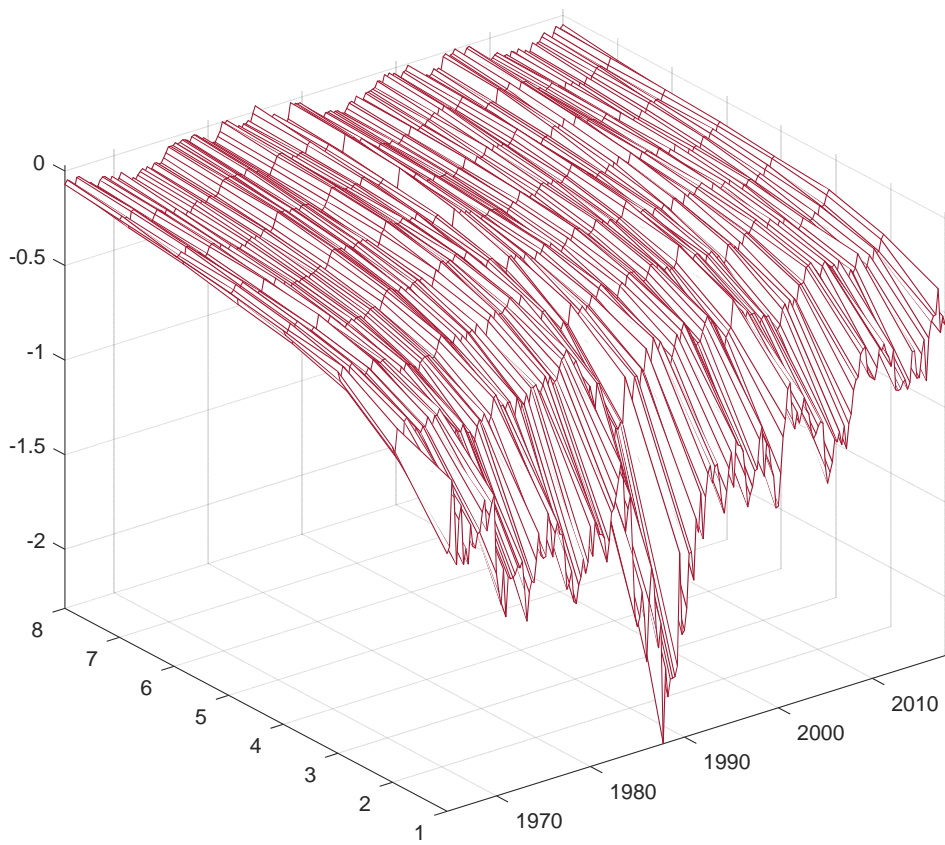


Figure 5C. Dynamic Multiplier: Price Markup shocks