Has the Phillips Curve Flattened and Why? *

Atsushi Inoue[1], Barbara Rossi[2], and Yiru Wang[3]

[1]*Vanderbilt University*†
[2]*ICREA-Universitat Pompeu Fabra, Barcelona School of Economics and CREI*‡
[3]*University of Pittsburgh*§

September 12, 2022

Abstract

A recent debate centers on the instability of the slope of the Phillips curve over time. Most of the empirical evidence on the dynamic evolution of the Phillips curve, however, is based on either structural or semi-structural models with time-varying parameters or instrumental variable estimates in ad-hoc sub-samples. We contribute to the literature by offering insights from a flexible time-varying instrumental variable approach. Even after controlling for endogeneity, we find evidence of a weakening of the structural slope of the Phillips curve starting around 1980. We find that the weakening of the cyclical relationship between unemployment and inflation is due to a flattening of the Phillips curve over time rather than to monetary policy. We also offer new insights on the Phillips curve during the recent pandemic: We find that the flattening has reverted and the Phillips curve is coming back.

Keywords: Inflation, Unemployment, Phillips curve, Local Projections, Instabilities, Time-varying Parameters, IV estimation.
JEL codes: C13, C32, E32, Q35

1 Introduction

Inflation and unemployment seem to have become disconnected during recent years. The correlation between inflation and real activity at business cycle frequencies has decreased in the 1990s (e.g. Atkeson and Ohanian, 2001, Stock and Watson, 2007, 2008, 2020) and in particular during the years of the expansion that followed the recent financial crisis of 2007-2009 (the so-called missing deflation – see Hall, 2011, Ball and Mazumder, 2011, Coibion and Gorodnichenko, 2015, Bobeica and Jarociński, 2019).

The decrease in the cyclical correlation between inflation and unemployment has been attributed by some to a flattening of the slope of the Phillips curve. If the Phillips curve indeed became flat, it would imply that more extreme policy measures would be necessary to maintain

*We thank seminar participants at the Belgian Macroeconomics Workshop.
†Address: Department of Economics, Vanderbilt University. Email: ainoue@vanderbilt.edu.
‡Address: c/Ramon Trias Fargas 25/27, 08005 Barcelona, Spain. Email: barbara.rossi@upf.edu.
§Address: Department of Economics, University of Pittsburgh. Email: yiru.wang@pitt.edu.

inflation at its target value. Thus, the question of whether the Phillips curve flattened is of high empirical relevance.[1]

One of the main challenges in the estimation of the Phillips curve is the presence of endogeneity problems, as inflation and unemployment are jointly determined in equilibrium. There are two main approaches to handling endogeneity: estimating the Phillips curve as part of a structural macroeconomic model (either Structural VARs or DSGEs) or focusing on the Phillips curve relationship by relying on instrumental variables methods. Both have advantages and disadvantages. On the one hand, it is well-known how to estimate DSGEs and Structural VARs in the presence of instabilities; however, DSGEs and Structural VARs are full-information estimation procedures, and potential mis-specification in any other part of the model might potentially contaminate the Phillips curve estimate. This could be a potentially serious problem during the recent financial crisis, where structural models have to confront serious mis-specification challenges. On the other hand, limited-information approaches, such as instrumental variables (e.g. Galí and Gertler, 1999, or Galí, Gertler and Lopez-Salido, 2005), are less affected by potential mis-specification but methodologies to address time-variation in instrumental variable models were lacking in the literature.

The contribution of this paper is to directly estimate the time-varying structural Phillips curve via limited-information methods. Our approach relies on the novel methodology proposed by Inoue et al. (2022) to estimate local projections and instrumental variables models with time-varying parameters (TVP-LP-IV). We also propose novel TVP-LP-IV estimates robust to weak instruments.

We find that the slope of the Phillips curve weakened since the early 1980s but it started reverting back in the most recent pandemic period. We also find that the decrease in the correlation between unemployment and inflation cannot be attributed to monetary policy; rather, to the decrease in the slope of the Phillips curve.

In contrast to our work, most of the existing literature relies either on reduced-form time-varying parameter approaches (Ball and Mazumder, 2019) or semi-structural time-varying parameter approaches (Galí and Gambetti, 2018); structural models estimation in given sub-samples (Del Negro et al., 2020); or instrumental variable estimation in given sub-samples (Barnichon and Mesters, 2020, 2021).

More in detail, a first strand of the literature uses time-varying parameter methods in reduced-form or semi-structural models. Reduced-form approaches attempt at studying the correlation between inflation and unemployment without resolving the endogeneity problem. For example, Stock and Watson (2008) survey the literature on the evaluation of inflation forecasts in the United States and suggest that Phillips curve forecasts are better than competing multivariate forecasts, although their performance is episodic relative to a univariate benchmark, again pointing to instabilities. Ball and Mazumder (2019) argue that expected inflation was backward-looking until the late 1990s, but then became strongly anchored at the central bank's target value, which would explain why inflation did not decrease in the high unemployment period around Great Recession. Galí and Gambetti (2018) adopt a semi-structural approach by estimating a time-varying parameter Vector Autoregression (VAR) model to identify economic shocks, then using such shocks to purge the Phillips curve variables and achieve identification of the Phillips curve parameters. In a similar spirit, Bergholt et al. (2022) estimate structural

[1]The flattening of the Phillips curve is not the only potential explanation for the disconnect between inflation and real variables during the Great Recession; other explanations which we entertain in this paper include the possibility that inflation or economic slack are mis-measured and that monetary policy is better at stabilizing inflation, thus flattening aggregate demand.

shocks using sign restrictions in constant-parameter VARs; then, they investigate changes in the Phillips curve over time using inflation and unemployment data purged by the relevant shocks in either sub-samples or rolling windows. Differently from Galí and Gambetti (2018) and Bergholt et al. (2022), we directly estimate the Phillips curve using instrumental variable methods that do not require identifying all the structural shocks in the economy.

A second strand of the literature relies on structural models. For example, Primiceri (2006) estimates a time-varying parameter structural model for the US economy. Del Negro et al. (2020) investigate whether the flattening of the Phillips curve is the explanation behind the disconnect between inflation and unemployment by focusing on (Structural) Time-varying parameter VARs and DSGE models. They account for the potential time-variation in the relationship between inflation and real activity by separately estimating their Structural VAR and DSGE models in two sub-samples, before and after 1989. The break date is determined by an a-priori choice, as a compromise between choosing a date where the economy became more stable (i.e. the Great Moderation, that started in 1984) and the stability of inflation itself, which seems to date back to the mid-1990s. Differently from their work, we rely directly on estimating the structural Phillips curve via limited-information methods, which are more robust to mis-specification, and let the instability in inflation dynamics to freely emerge within our time-varying instrumental variable (TVP-IV) estimator.

A third strand of the literature focuses on instrumental variables or external information. McLeay and Tenreyro (2019) argue that the fact that inflation follows a seemingly exogenous statistical process, unrelated to the output gap, does not mean that the Phillips curve has disappeared. They show that, in a theoretical model, by increasing inflation when output is below potential, monetary policy can generate a negative correlation between inflation and the output gap, blurring the identification of the Phillips curve. They find evidence against the disappearance of the Phillips curve using regional data. The identification problem pointed out by McLeay and Tenreyro (2019), however, can be addressed by using instrumental variables, like we do. Barnichon and Mesters (2020, 2021) estimate the Phillips curve and the Phillips multiplier using narrative monetary policy shocks as instruments to address the endogeneity problem. To take into account time-variation, they split the sample at a known break date. Overall, while these papers perform sub-sample analysis (typically using 1990 as the break point estimate), none of these papers allow for general patterns of time variation, which is instead the main contribution of our paper.

The remainder of the paper is organized as follows. The next section presents the methodology. Section 3 discusses the empirical evidence on the evolution of both the Phillips relation as well as the slope of the structural Phillips curve over time. Section 4 investigates whether the decrease in the correlation between unemployment and inflation is due to monetary policy or to a decrease in the slope of the Phillips curve, and Section 5 discusses the most recent evidence on the Phillips curve, including the recent pandemic. Section 6 concludes.

2 The Phillips Curve Model and the Methodology

Our benchmark Phillips curve is the classic version by Galí and Gertler (1999):

$$\pi_t = c + \gamma_f E_t\left(\pi_{t+1}\right) + \gamma_b \pi_{t-1} + \lambda x_t + u_t,$$

where π_t denotes inflation, x_t denotes the measure of real marginal cost, $E_t\left(.\right)$ denotes conditional expectations at time t and u_t is an unobserved shock. This specification is the same as Galí and Gertler (1999, GG herafter) and Galí, Gertler, and López-Salido (2005, GGLS herafter), that is a hybrid New Keynesian Phillips Curve (NKPC) with lagged inflation and the unemployment gap as the forcing variable. We will estimate the NKPC using instrumental variables under the rational expectations assumption.

Our main focus is estimating the slope of the Phillips curve, namely λ, which is the object of a lively debate. On the one hand, several researchers found that the slope of the Phillips curve has flattened or even that the Phillips curve "died", in the sense that the slope is close to zero – see Coibion and Gorodnichenko (2015), Blanchard (2016), Ball and Mazumder (2019), and Stock and Watson (2020), among others. On the other hand, Barnichon and Mesters (2021) and Bergholt, Furlanetto and Vaccaro-Grange (2022), among others, argue that it is not. For example, using a Phillips-multiplier approach, Barnichon and Mesters (2021) argue that the inflation-unemployment trade-off went from being very large before 1990 to being small, but still significant, after 1990, and that the decline in the trade-off is mostly due to the anchoring of inflation expectations. Tenreyro and Twaites (2016) argue that the disconnect between inflation and real activity may not only be due to a flat Phillips curve but also to a flat aggregate demand, such as one where monetary policy strongly responds to inflation. For example, if the central bank achieves a perfect inflation stability, the researcher would observe inflation to be uncorrelated with real activity even if the Phillips curve slope is not zero. Del Negro et al. (2020) find that the slope of the Phillips curve substantially weakened over time, and that is the main reason for the disconnect between inflation and unemployment. Bergholt, Furlanetto and Vaccaro-Grange (2022) find that the Phillips curve is "dead" only unconditionally: once it is purged for supply shocks, the Phillips curve is alive and well, and may even have steepened since the financial crisis.[2]

As is well-known in the literature, there are several econometric challenges in estimating the structural Phillips curve. A first challenge is that the forcing variable x_t may be correlated with the structural error term u_t, thus resulting in an endogeneity problem. An additional challenge is that the expected inflation term $E_t(\pi_{t+1})$ is not only endogenous but also unobservable. To address these issues, we consider an Instrumental Variable (IV) approach to identification, as in Galí and Gertler (1999), Galí, Gertler, López-Salido (2005) and Barnichon and Mesters (2020), among others. Suppose Z_t is a vector of valid instruments such that

$$E[Z_t(\pi_t - \gamma_b \pi_{t-1} - \gamma_f \pi_{t+1} - \lambda x_t)] = 0.$$

Under the rational expectations assumption, we can include as instruments in Z_t any predetermined variables, such as lags of the endogenous variables.

Our contribution is to estimate the Phillips curve slope allowing the parameters to be time-varying using the TVP-IV approach developed in Inoue et al. (2022). We will estimate the model:

$$\pi_t = c_t + \gamma_{f,t}\pi_{t+1} + \gamma_{b,t}\pi_{t-1} + \lambda_t x_t + u_t,$$

where the parameters evolve slowly over time according to a random walk with a small variance, which is proportional to a nuisance parameter. Modeling the time variation in the parameters in a random walk fashion has a long tradition in the literature – see e.g. Cogley and Sargent (2005) and Primiceri (2006).

Let us highlight that we not only allow the coefficients of the Phillips curve to be time-varying, but also the volatility of the cost-push shock, u_t.

[2] Bergholt et al. (2022) estimate the Phillips curve using OLS after purging the variables by supply shocks.

We estimate the TVP-IV regression as follows:

$$\begin{bmatrix} x_t \\ \pi_{t+1} \\ \pi_t \end{bmatrix} = \begin{bmatrix} \beta'_{x,z,t} & c_{x,t} & \beta_{x,\pi_b,t} \\ \beta'_{\pi_f,z,t} & c_{\pi_f,t} & \beta_{\pi_f,\pi_b,t} \\ \lambda\beta'_{x,z,t} + \gamma_{f,t}\beta'_{\pi_f,z,t} & \lambda_t c_{x,t} + \gamma_{f,t} c_{\pi_f,t} + c_{\pi,t} & \lambda_t \beta_{x,\pi_b,t} + \gamma_{f,t}\beta_{\pi_f,\pi_b,t} + \gamma_{b,t} \end{bmatrix} \begin{bmatrix} z_t \\ 1 \\ \pi_{t-1} \end{bmatrix} + \begin{bmatrix} v_{x,t} \\ v_{\pi_f,t} \\ v_{\pi,t} \end{bmatrix}, \quad (1)$$

where $(x_t, \pi_{t+1})'$ are endougenous variables, and $Z_t = (z'_t, \pi_{t-1})'$, where z_t are the excluded instruments. Our main parameter of interest is λ_t. The parameter path is estimated according to a minimum weighted average risk criterion, as in Inoue et al. (2022), to which we refer the readers for more details.

3 Has the Phillips Curve Flattened Over Time?

In this section we discuss our main empirical evidence on the three main concepts surrounding estimation of the Phillips curve in the literature: the Phillips relation, the slope of the Phillips curve and the Phillips multiplier.

The Phillips Relation

We start by empirically investigating the reduced-form relationship between inflation and the labor share over time, following Stock and Watson (2020). We focus on the estimated slope ($\beta_{1,t}$) in the following Phillips relation:

$$E_t \Delta_4 \pi_t^4 = \beta_{0,t} + \beta_{1,t} x_t^4, \quad (2)$$

where x_t^4 is the change in the average value of variable "x" between times t and t-3 and $\Delta_4 = (1 - L^4)$, L denotes the lag operator such that $Lx_t = x_{t-1}$. There are several candidate choices for both inflation and real marginal cost measures – see e.g. the literature review in Mavroeidis, Plagborg-Møller and Stock (2014). In our analysis, inflation (π_t) is measured by personal consumption expenditure price index (PCE excluding food and energy, PCExFE) and x_t is a measure of slack.

Figure 1 reports the time-varying estimate of $\beta_{1,t}$ for various measures of slack for the US. Each panel in the figure corresponds to a different measure of slack, inspired by Stock and Watson (2020, Table 1). We consider: the unemployment gap, as measured by the Congressional Budget Office (CBO) in panel (a); the GDP gap, also from the CBO, in panel (b); the unemployment gap filtered using a two-sided filter[3] in panel (c); a measure of the short-term unemployment gap[4] in panel (d); the employment-population ratio (again obtained via a two-sided filter) in panel (e); the employment-population ratio focusing on population of age between 25 and 54 year-old (again obtained via a two-sided filter) in panel (f); the capacity utilization rate in panel (g); the unemployment rate measured as a real-time slack in panel

[3]The two-sided filter used in this section to obtain the gap measure is the same as Stock and Watson (2020) and it is a band-pass Butterworth filter of degree 6, with lower and upper cutoffs corresponding to periods of 32 and 6 quarters, respectively.

[4]The short-term unemployment gap is obtained from the short-term unemployment rate (those unemployed 26 weeks or less as a fraction of the labour force), i.e. the measure of slack in Ball and Mazumder (2014), using the two-sided filter described in the previous footnote.

(h); and the short-term unemployment rate in panel (i).[5] The estimate based on our TVP-LP estimator is reported by the dashed line (in red).

We compare our estimate with Stock and Watson's (2020), who estimate the Phillips relation in three sub-samples; the latter is depicted by the solid (black) line, together with their 90 percent confidence band (black dotted lines).

INSERT FIGURE 1 HERE

As Figure 1 shows, the slope of the Phillips relation substantially flattened over time, and this emerges clearly in the data no matter whether we estimate the relation in sub-samples or using our time-varying estimator. Thus, the Phillips relation has disappeared in the data in the most recent period.

But does it mean that the Phillips curve has disappeared? Not necessarily, as the Phillips relation measures the correlation between inflation and unemployment, while the Phillips curve measures the trade-off between inflation and unemployment due to supply shocks.

The Slope of the Phillips Curve

In what follows, we will directly estimate the structural Phillips curve using the TVP-IV-based approach that flexibly allows the parameters to change over time while, at the same time, avoiding the endogeneity problem. Figure 2 plots the estimates of λ_t, $\gamma_{f,t}$, and $\gamma_{b,t}$ using the TVP-IV framework and contrasts it with the estimates from a constant parameter model a' la Galí and Gertler (1999) and Galí, Gertler and López-Salido (2005). The main sample ranges from 1970Q1 to 2008Q1,

We focus on a model specification where x_t is the unemployment gap (estimated by the Congressional Budget Office, CBO) and expected inflation is the three-quarter-ahead forecast of mean PGDP inflation from the Survey of Professional Forecasters. The set of instruments includes two lags of unemployment gap (from the CBO) and two lags of the output gap (estimated in real-time using a one-sided quadratically detrending procedure).[6]

Our instruments are both valid and strong. In fact, Hansen's J-statistic equals 1.955, with a p-value of 0.3763, indicating that the instruments are valid. Lewis and Mertens's (2022) weak IV test statistic equals 16.0254, and it is greater than the 90% critical value 14.0533, indicating that the instruments are strong. The Ganics et al. (2021) weak-instrument robust confidence interval for the strength of identification also points to strong instruments (the min eigen value is 1.6634, with a confidence interval equal to $(1.1505, 5.7933)$, which excludes zero).

As previously anticipated, the TVP-IV method can flexibly and robustly estimate the Phillips curve coefficient over time. In fact, it has both the advantage of letting the parameters change over time in a flexible way, including the variance of the error term, as well as being robust to endogeneity, as it relies on instrumental variable estimation.

INSERT FIGURE 2 HERE

[5] All slack measures are standardized, and they have the same mean and standard deviation as the unemployment gap from the CBO. They have also been transformed in order to be positively correlated with the CBO output gap. The Appendix provides more details on the data.

[6] The set of instruments is inspired by Galí and Gertler (1999) and Galí, Gertler and López-Salido (2005); however, we have excluded lagged inflation as an instrument because our analysis suggests that it is a weak instrument. Figure 8 in the Appendix shows that our main results are robust to using the same set of instrumental variables as Galí and Gertler (1999) and Galí, Gertler and López-Salido (2005).

Our results confirm a flattening of the slope of the Phillips curve (λ_t) in the last two decades. The slope has decreased, in absolute value, from around -0.12 in the early 1970s to around -0.04 in the most recent sample. In particular, notice how it trended downward in the 1990s, becoming effectively indistinguishable from zero.

Note also that the importance of the forward-looking component ($\gamma_{f,t}$) has slightly decreased from around 0.45 in the early 1970s to around 0.35 in the most recent sample. This suggests that agents have become less forward-looking over time. At the same time, the importance of the backward-looking component in the Phillips curve ($\gamma_{b,t}$) has remained constant over time.

Previous papers, notably Galí and Gambetti (2018), have estimated Phillips curve coefficients allowing the parameters to be time-varying. Their approach to addressing the endogeneity issues affecting the estimation of the Phillips curve relies on purging the OLS estimates using shocks identified via time-varying parameter Structural VARs and, hence, is semi-structural. In particular, their Structural VAR identification identifies several macroeconomic shocks, using both sign and long-run restrictions. Our approach instead does not require to separately identify shocks, as it relies on instrumental variables, in a way that directly parallels the pioneering work of Galí and Gertler (1999).

The Phillips Multiplier

In a recent paper, Barnichon and Mesters (2021) propose the "Phillips multiplier" as an alternative measure of the inflation-unemployment trade-off faced by policymakers, different from the slope of the Phillips curve. Relative to Barnichon and Mesters (2021), our approach (discussed in the previous sub-section) can directly and flexibly estimate the time-varying trade-off between inflation and unemployment in the classical specification of the Phillips curve, while at the same time having the advantage of allowing for a time-varying coefficient on expected inflation.[7] Furthermore, we can use the same instruments and specification as in the seminal contribution by Galí and Gertler (1999), which makes our approach more directly comparable to theirs.

However, the Phillips multiplier is also a measure of the trade-off between inflation and unemployment that we can consider in our framework. Therefore, we also estimate the Phillips multiplier using our time-varying method. By being estimated using instrumental variable methods, the Phillips multiplier also avoids the typical endogeneity problems afflicting the estimation of the Phillips curve. We focus on the same model specification as in Barnichon and Mesters (2021), and use monetary policy shocks as instruments, as monetary policy shocks are uncorrelated with supply shocks. Therefore, the Phillips multiplier will estimate the effect of an increase in unemployment on inflation conditional on the presence of monetary policy shocks.

The top two panels in Figure 3 report the impulse responses of inflation and unemployment to a monetary policy shock, whereas the bottom panel depicts the Phillips multiplier. The dashed (red) lines depict our time-varying estimates (each line corresponds to an impulse response estimated at a given point in time), while the continuous (black) line reports the full-sample estimate. The top two panels in the picture show that, conditionally on a contractionary monetary policy shock, inflation decreases and unemployment increases. The time-varying estimates show that the quantitative extent to which inflation and unemployment respond to

[7] Under some assumptions, in particular a constant forward-looking behavior of inflation, Barnichon and Mesters (2021) relate the Phillips multiplier to the slope of the Phillips curve. In our approach, we instead let the forward-looking component be freely time-varying.

a monetary policy shock changes significantly over time. As a result, the estimated Phillips multiplier also varies over time, as shown in the bottom panel of Figure 3.

INSERT FIGURE 3 HERE

In order to shed more light on the nature of the time-variation in the Phillips multiplier, Figure 4 depicts the Phillips multiplier over time for a selected horizon $h = 12$. Figure 4 (a) depicts the Phillips multiplier over time before 1990, using Romer and Romer (2004) monetary policy shocks as instruments. Figure 4 (b) depicts the Phillips multiplier over time after 1990, using the recent high-frequency identification (HFI) monetary surprises as instruments. For each sub-figure, the TVP-IV estimates are reported by dash-dot red lines, while the continuous black lines report the sub-sample estimates of the Phillips multiplier in Barnichon and Mesters (2021). The figure shows that the Phillips multiplier also decreased substantially over time, and the decrease dates back to the 1970s. After 1990, the TVP-IV multiplier estimates are close to 0.

INSERT FIGURE 4 HERE

Comparing our results in Figure 2 to Figures 3 and 4, the flattening of the Phillips curve seems a robust result in our data, no matter whether we consider the Phillips (cor)relation, the Phillips multiplier or the slope of the structural Phillips curve.

4 Why Has the Cyclical Correlation Between Inflation and Output Decreased?

Several researchers have made compelling arguments that the reason for the decrease in the cyclical correlation between inflation and unemployment is related to monetary policy actions. According to this explanation, a more responsive monetary policy to inflation and economic conditions would tighten monetary policy more when it perceives inflation to be increasing, in order to keep the latter under control: this causes unemployment to rise, resulting in a positive correlation between inflation and unemployment that biases the slope coefficient of the Phillips curve toward zero. See Haldane and Quah (1999); Roberts (2006); Williams (2006); Mishkin (2007); Carlstrom, Fuerst, and Paustian (2009); and, more recently, McLeay and Tenreyro (2018).

As is well-known, the correlation between inflation and unemployment is the same as the slope of the Phillips curve only in the presence of no endogeneity bias and no measurement error. Thus, the endogeneity problem can be solved using valid and relevant instruments.[8] In the presence of an endogeneity bias due to monetary policy actions, IV estimates will still be consistent provided the instruments satisfy the required statistical conditions – that is, the chosen instruments should be both valid and relevant.

The IV approach we discussed in the previous section, that uses lagged macroeconomic variables as instruments, suggests that the slope of the Phillips curve decreased over time in a manner similar to the decrease in the correlation between inflation and unemployment. Our results therefore suggest that the decrease in the correlation between inflation and unemployment is due to a decrease in the slope of the Phillips curve and not to other factors, among which monetary policy.

[8]Some researchers have attempted to solve the endogeneity problem by using regional data (see McLeay and Tenreyro, 2018, and Hazell et al., 2022, among others). Here we will instead maintain the same framework as in the classical literature on the Phillips curve debate, which focuses on macroeconomic data; however, we will shed light on the issue by using both valid and strong instruments, as well as weak identification robust confidence intervals.

As mentioned, our analysis in the previous section is based on an instrumental variable method where instruments are both valid and relevant. However, the theoretical validity of the instruments also requires that the residuals of the Phillips curve are not serially correlated, otherwise the IV procedure may not correctly estimate the slope of the Phillips curve. Lurking inside the residual of the Phillips curve there are both cost push shocks as well as measurement error, and both might be correlated with the lagged instruments via a correlation with their own lags. In our analysis, the residuals of the Phillips curve show some evidence of serial correlation.

To be robust to this potential problem, we consider aggregate demand shocks as instruments, as in Barnichon and Mesters (2020). In particular, we consider monetary policy shocks, which are potentially valid instruments, as both the measurement error and unobserved supply shocks (such as labor supply, price or technology shock) would be uncorrelated with such shocks. Although we use the same model specification and the same identification strategy as Barnichon and Mesters (2020), we will use it in our framework to analyze whether the culprit behind the decrease in the correlation between inflation and output is monetary policy, while their approach does not allow for time-varying parameters.

We estimate the following Phillips curve specification:[9]

$$\pi_t = c_t + \lambda_t x_t + \gamma_{f,t}\pi_{t+1} + \gamma_{b,t}\pi_{t-1} + u_t,$$

where the set of instrument is the Almond parameterization of twenty lags of the Romer and Romer (2004) monetary policy shocks and the forcing variable x_t is the unemployment gap. We estimate the equation in the sample starting in 1974:Q1 and ending in 2007:Q4, due to data limitations in the availability of the monetary policy shock data.[10]

Hansen's J-statistic for over-identification is 0.022, with a p-value of 0.8813, indicating that the instruments are valid in the our sample. Ganics et al. (2021) weak instrument-robust procedure implies weak instruments, as the min eigenvalue is 0.0001, with a confidence interval equal to $(0, 0.5533)$. Lewis and Mertens's (2022) weak IV test statistic equals 4.1991, also indicating that the instruments are weak.[11] Therefore, we develop a novel TVP-IV methodology to obtain estimates and confidence bands that are robust to weak instruments, described in what follows.

We estimate the reduced form parameters in the following eq. (3):[12]

$$\begin{bmatrix} x_t \\ \pi_{t+1} \\ \pi_t \end{bmatrix} = \begin{bmatrix} \beta'_{x,z,t} & c_{x,t} & \beta_{x,\pi_b,t} \\ \beta'_{\pi_f,z,t} & c_{\pi_f,t} & \beta_{\pi_f,\pi_b,t} \\ \beta'_{\pi,z,t} & \tilde{c}_{\pi,t} & \beta_{\pi,\pi_b,t} \end{bmatrix} \begin{bmatrix} z_t \\ 1 \\ \pi_{t-1} \end{bmatrix} + \begin{bmatrix} v_{x,t} \\ v_{\pi_f,t} \\ v_{\pi,t} \end{bmatrix}, \tag{3}$$

where $\beta'_{\pi,z,t} = \lambda\beta'_{x,z,t} + \gamma_{f,t}\beta'_{\pi_f,z,t}$, $\tilde{c}_{\pi,t} = \lambda_t c_{x,t} + \gamma_{f,t} c_{\pi_f,t} + c_{\pi,t}$, and $\beta_{\pi,\pi_b,t} = \lambda_t\beta_{x,\pi_b,t} + \gamma_{f,t}\beta_{\pi_f,\pi_b,t} + \gamma_{b,t}$. In particular, the reduced form parameters, denoted by θ_t^{rf}, are $\beta'_{\pi,z,t}$, $\tilde{c}_{\pi,t}$ and $\beta_{\pi,\pi_b,t}$, rather than λ_t, γ_f, and γ_b. In this step, we obtain the point estimates of the reduced form parameters $\hat{\theta}_t^{rf}$ as well as their joint distribution, denoted by $F_{\hat{\theta}_t^{rf}}(\cdot)$, for $t = 1, 2, ...$ The estimates $\hat{\lambda}_t, \hat{\gamma}_f, \hat{\gamma}_b$ can be recovered from the reduced-from parameter estimate $\hat{\theta}_t^{rf}$. Then,

[9] This is the same model specification as Barnichon and Mesters (2020).

[10] The monetary policy shock series that we use ends in 2007. During the zero lower bound period there are fewer monetary policy shocks anyway, which might invalidate the strength of the instrument.

[11] HAC-robust variance estimates are implemented with Barnichon and Mesters's (2020) choice, which is 5 lags.

[12] That is, instead of estimating eq. (1).

we randomly draw $\mathcal{M}$ times from the joint distribution $F_{\hat{\theta}_t^{rf}}(\cdot)$ and repeatly recover $\hat{\lambda}_t^i, \hat{\gamma}_f^i, \hat{\gamma}_b^i$, $i = 1, 2, ..., \mathcal{M}$. The confidence bands and the median estimate can be obtained by the corresponding quantiles. This estimation procedure is valid since the weak instrument issue won't affect the estimation of the reduced form parameters θ_t^{rf}.

Figure 5 shows the time-varying estimate of the slope of the Phillips curve; dotted lines report the 90 percent confidence bands robust to weak instruments. The figure confirms our result that the slope of the Phillips curve substantially flattened in the 1980s and 1990s.[13]

Finally, we revisit the model specification of GGLS, using exactly the same set of instruments they used. In our data, their set of instruments is weak; therefore we report estimates and confidence bands using our TVP-IV method robust to weak instruments. Panel (a) in Figure 6 confirms, again, a decrease in the slope of the Phillips curve over time.[14]

In conclusion, using a variety of different specifications and, in particular, a specification robust to the presence of measurement errors and serial correlation in the residuals of the Phillips curve, and using weak instrument robust techniques, we find robust evidence that the decrease in the correlation between inflation and unemployment is due to a flattening of the Phillips curve, rather than to monetary policy.

5 What's Up with the Phillips Curve in the Recent Pandemics?

In this section we focus on the Phillips curve during the recent financial crisis and, especially, the recent pandemics. Both have contributed to a substantially unstable macroeconomic environment, which the TVP-IV approach in this paper can handle in a robust and flexible way.

Figure 7 plots the time-varying estimates of the Phillips curve parameters using data up to the end of 2021 in the main GGLS specification. The top panel in the picture shows the slope of the Phillips curve (λ_t) together with 90 percent confidence bands; the bands are constructed using our method robust to weak instruments. It is clear from the figure that, after hovering close to zero (in absolute value) until the mid-2000s, the slope has started to increase again. Thus, contrary to the literature that attributes the missing disinflation in the recent financial crisis to the weakening of the Phillips curve, we find evidence that the Phillips curve is becoming again alive and well.

What else happened during the financial and pandemic crises? The bottom panel in Figure 7 shows a steady increase in the degree of forward-looking behavior in inflation. The upward trend, that started during the great moderation, has recently become even stronger, and the most recent estimate is around 0.6. On the other hand, the degree of backward-looking behavior in inflation has weakened substantially. The downward trend, which started since the 1970s, has brought the parameter to 0.2, and statistically insignificantly different from zero.

INSERT FIGURE 7 HERE

Overall, our findings suggest that, in setting prices, agents pay more attention to the future and less to the past. The fact that past inflation has lost importance in the agents' price-setting behavior may explain the decrease in the overall serial correlation in inflation and its lack of predictability over time (Stock and Watson, 2007).

[13] Barnichon and Mesters (2020) also study the specification considering the output gap as the forcing variable. We report our result of this specification in Figure 9 in the Appendix, which also confirms that the slope of the Phillips curve substantially flattened in the 1980s and 1990s.

[14] HAC-robust variance estimates are implemented with Lazarus, Lewis, Stock, Watson's (2018) recommendation, which implies 16 lags.

6 Conclusion

We contribute to the debate surrounding the instability of the relation between unemployment and inflation over time by offering insights from a flexible time-varying instrumental variable approach.

We find that the weakening of the cyclical correlation between inflation and unemployment is due to a flattening in the slope of the Phillips curve, rather than to monetary policy. The slope of the structural Phillips curve has decreased over time since the 1980s. In the most recent period since the Great Recession and during the recent pandemic, the slope of the Phillips curve has increased again.

Our results are based on an approach that has the advantage of avoiding endogeneity while, at the same time, being robust to changes in the economic environment. In addition, by virtue of the approach taken in this paper, our conclusions do not require making auxiliary assumptions on the rest of the economy nor estimating a fully specified model, and hence are more robust to mis-specification than existing, full-information approaches. We demonstrate the robustness of our results to various specifications that feature both strong instruments as well as weak-identification robust ones.

Figures and Tables

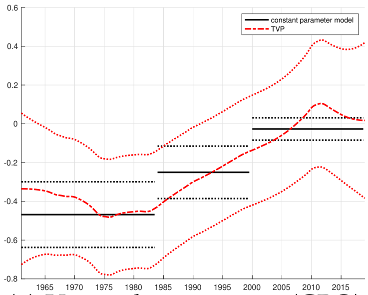

(a) Unemployment gap (CBO)

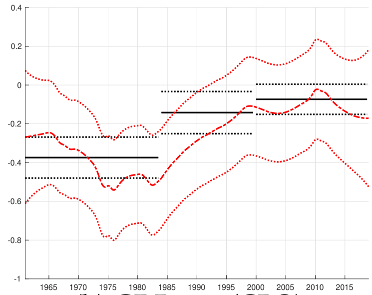

(b) GDP gap (CBO)

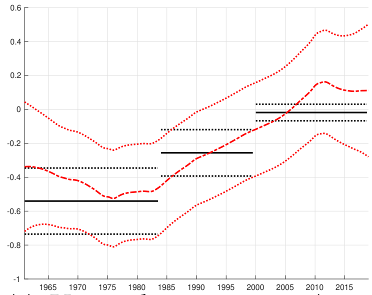

(c) Unemployment gap (two-sided filtered)

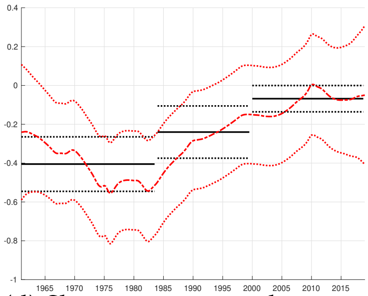

(d) Short-term unemployment gap (two-sided filtered)

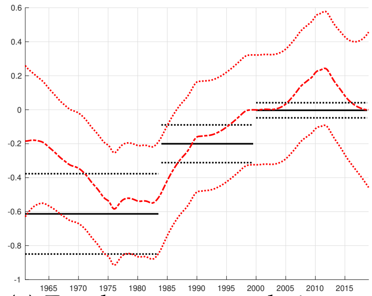

(e) Employment-population ratio (two-sided filtered)

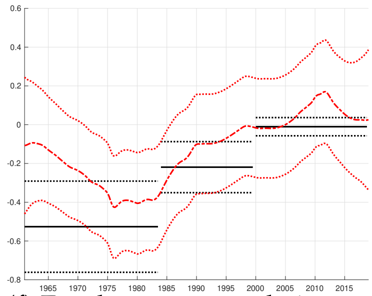

(f) Employment-population ratio ages 25-54 (two-sided filtered)

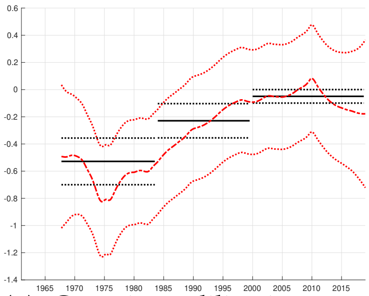

(g) Capacity utlilization rate (two-sided filtered)

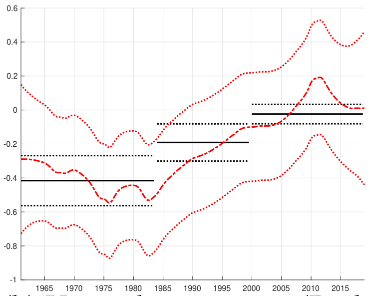

(h) Unemployment rate (Real-time slack)

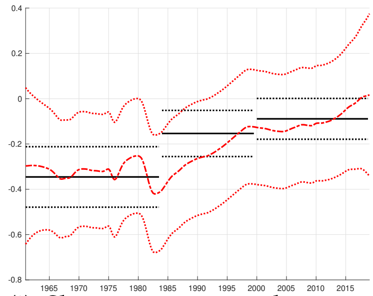

(i) Short-term unemployment rate (Real-time slack)

Figure 1: **The time-varying Phillips relation**. The figure shows the estimated slope (β_1) in the Phillips relation: $E_t\Delta_4\pi_t^4 = \beta_0 + \beta_1 x_t^4$, where x_t^4 is the change in the average value of variable "x" between times t and t-3 and $\Delta_4 = (1 - L^4)$. Inflation is measured by PCE-xFE and x_t are the various measures of slack for the US (See Stock and Watson's (2020) Table 1). The period is 1961-2019.

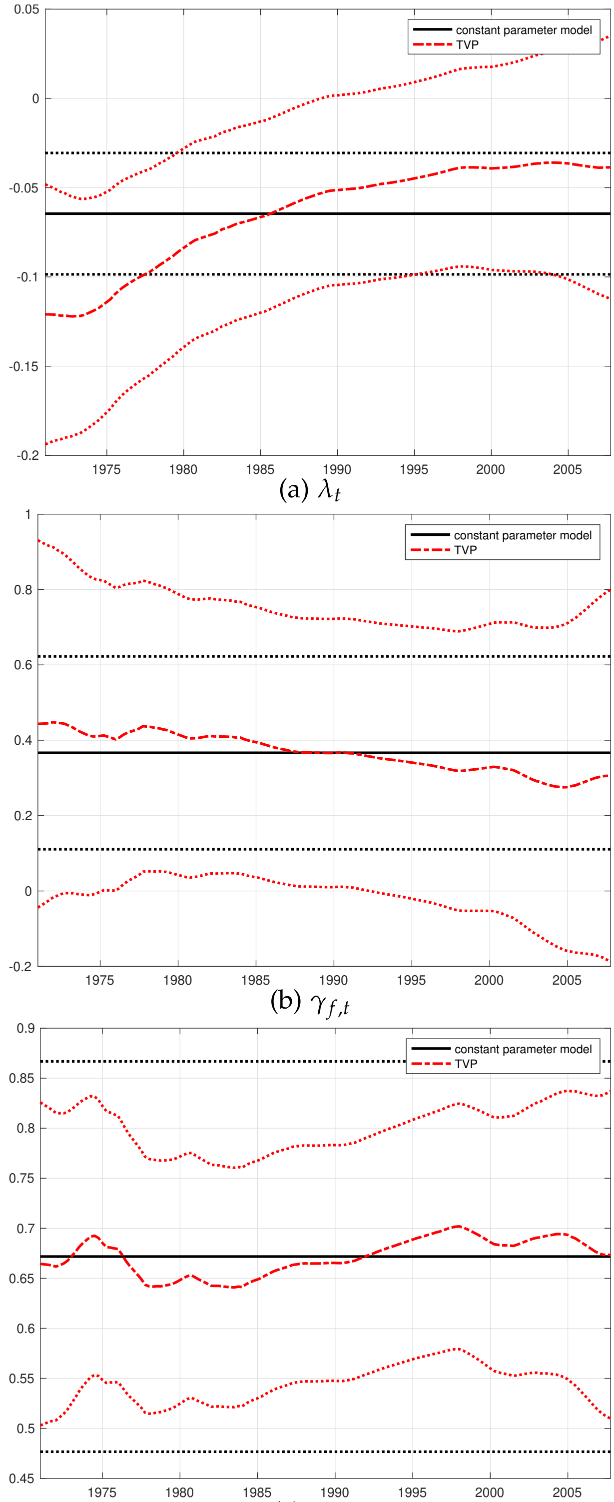

(a) λ_t

(b) $\gamma_{f,t}$

(c) $\gamma_{b,t}$

Figure 2: **The time-varying Phillips curve**. The figure shows the coefficients of the structural Phillips curve estimated using the TVP-IV method (dashed lines) versus the full-sample constant estimate, together with 90% confidence bands. The slope of the Phillips curve is reported in panel (a). The sample is 1970Q1-2008Q1. HAC-robust variance estimates are implemented with 4 lags, following Mavroeidis, Plagborg-Møller and Stock (2014).

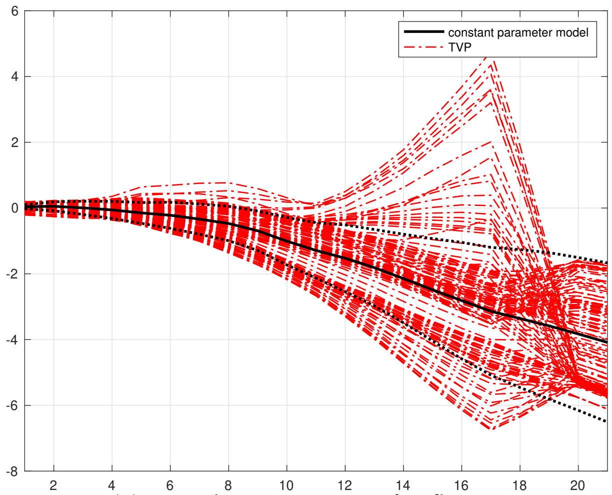

(a) impulse responses of inflation

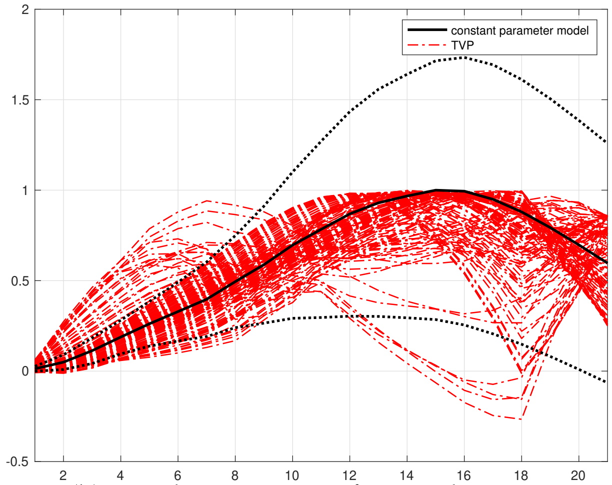

(b) impulse responses of unemployment

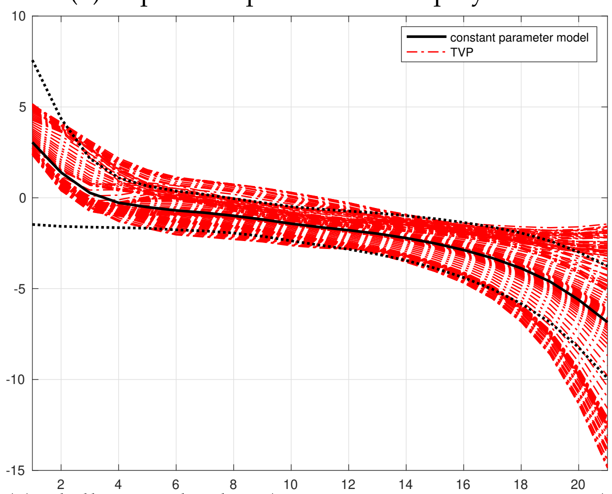

(c) Phillips multiplier (one-step TVP-IV estimation)

Figure 3: **The time-varying Phillips multiplier**. The figure shows the US Phillips multiplier in the sample 1969q1-2007q4. The black continuous line is Barnichon and Mesters' (2021, Fig. 1) full-sample multiplier and the red dashed lines are the TVP-IV multipliers.

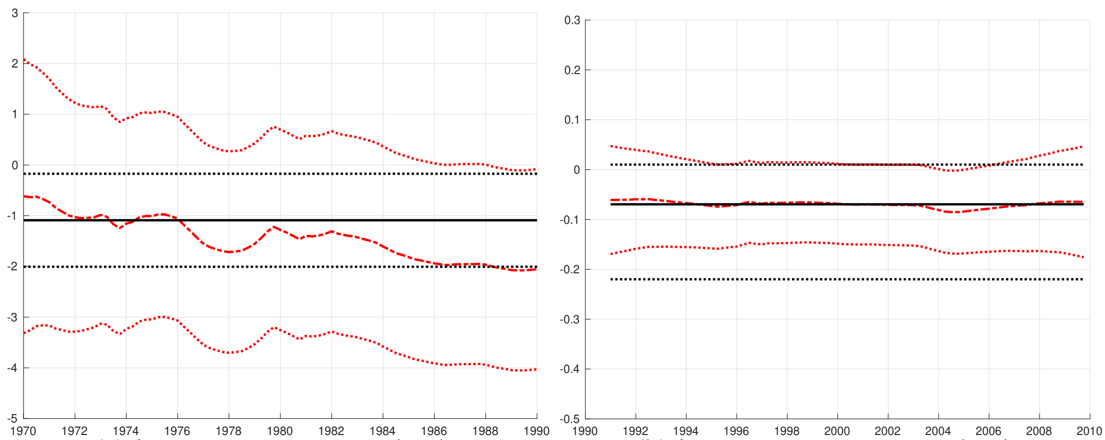

(a) $h = 12$, pre-1990, RR shocks (b) $h = 12$, post-1990, HFI shocks

Figure 4: **The time-varying Phillips multiplier**. The figure shows the TVP-LP estimated US Phillips multiplier over time for selected horizons (dashed red line), using Romer and Romer's shocks as instruments for pre-1990 sub-samples and HFI shocks as instruments for post-1990 sub-samples as to compare with Barnichon and Mesters's (2021) constant multiplier estimates, reported by the black solid line.

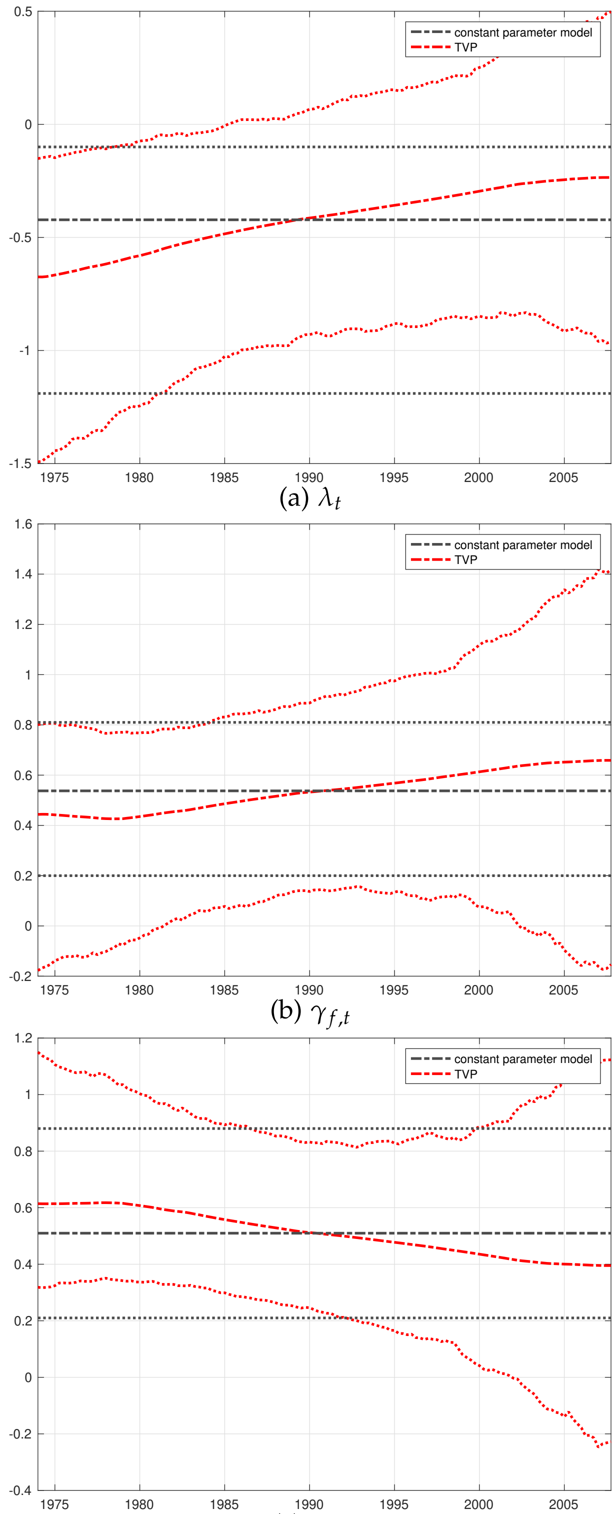

(a) λ_t

(b) $\gamma_{f,t}$

(c) $\gamma_{b,t}$

Figure 5: **The time-varying Phillips curve using monetary policy shocks as instruments**. The figure reports the estimated TVP-IV coefficients as well as their full-sample counterparts, together with 90% weak-instrument robust confidence bands. The sample is 1974Q1:2007Q4. The specification is as follows: we use the Almond parameterization of the 20 lags of Romer and Romer monetary policy shocks, and the unemployment gap obtained via the Hodrick-Prescott filter. The number of lags considered in the Newey and West's (1987) estimator follows Barnichon and Mesters (2020), which implies 5 lags. The bands are smoothed using a seven quarter centered moving average.

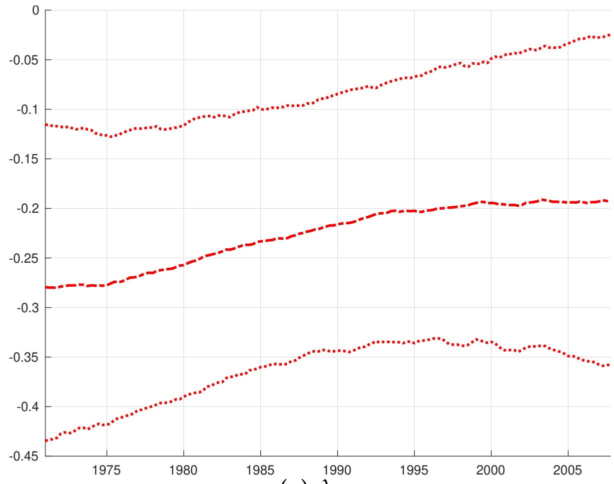

(a) λ_t

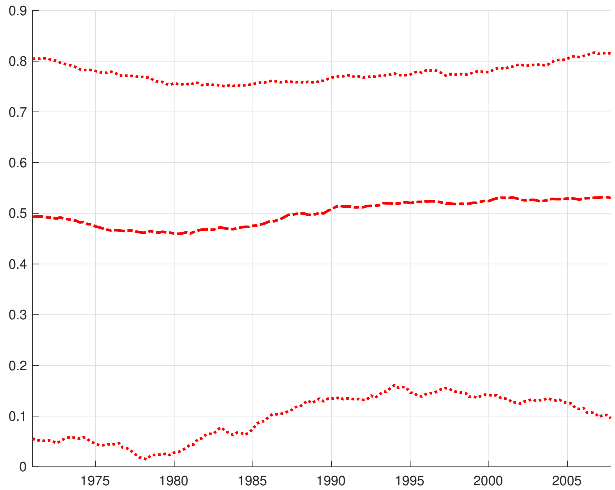

(b) $\gamma_{f,t}$

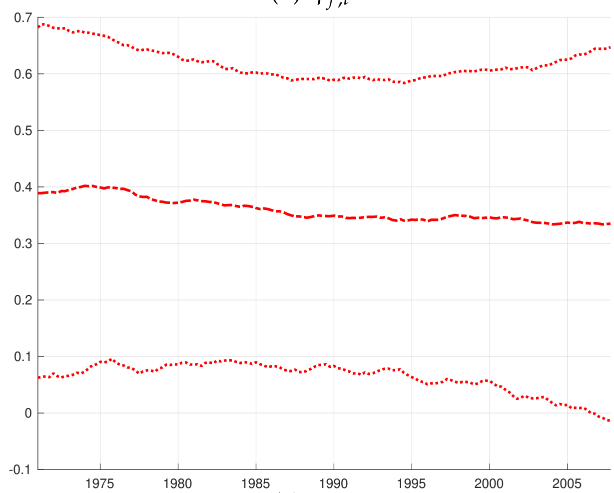

(c) $\gamma_{b,t}$

Figure 6: **The time-varying Phillips curve: Robustness to weak instruments.** The figure shows the estimated TVP-IV coefficients together with 90% weak-instrument robust confidence bands, 1970Q1:2021Q4. The specification is the same as GGLS. We estimate the hybrid NKPC with one lag of inflation and the unemployment gap (cbo) as forcing variable using the instrument set: four lags of inflation and two lags of the unemployment gap (CBO), wage inflation, and output gap (CBO). HAC-robust variance estimates are implemented with Lazarus, Lewis, Stock, Watson's (2018) recommendation, which implies to 16 lags. The bands are smoothed using a seven quarter centered moving average.

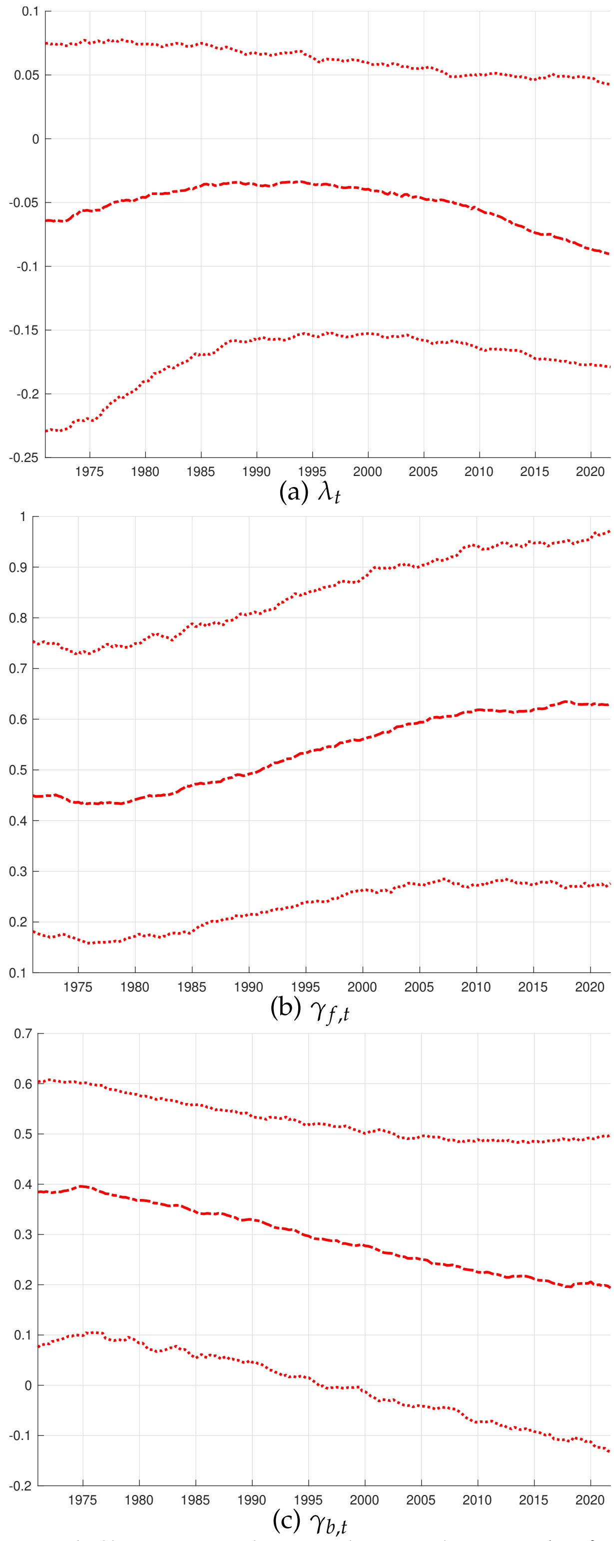

Figure 7: **The time-varying Phillips curve during the pandemic.** The figure shows the estimated TVP-IV coefficients together with 90% weak-instrument robust confidence bands, 1970Q1:2021Q4. The specification is the same as GGLS. We estimate the hybrid NKPC with one lag of inflation and the unemployment gap (cbo) as forcing variable using the instrument set: four lags of inflation and two lags of the unemployment gap (CBO), wage inflation, and output gap (CBO). HAC-robust variance estimates are implemented with Lazarus, Lewis, Stock, Watson's (2018) recommendation, which implies to 19 lags. The bands are smoothed using a seven quarter centered moving average.

7 Appendix. Data Description

This Appendix describes the data used in this paper. The data are quarterly and the span of the data is determined by data availability.

Inflation is measured as the "Implicit GDP deflator" (mnemonics "GDPDEF"). The data is transformed as follows: 100 times the log difference of the GDP deflator. For labor share, we use the "Business Sector: Labor Share for All Employed Persons" (mnemonics "PRS84006173"). The data is transformed as: 100*ln(PRS84006173/100). The instruments used follow Galí, Gertler, and López-Salido (2001), including four lags of inflation and two lags of the labor share, wage inflation, and output gap. For wage inflation, we use the "Business Sector: Hourly Compensation for All Employed Persons" (mnemonics "HCOMPBS"). The data is transformed as: 100 times the log difference of the HCOMPBS. The output gap is an economic measure of the difference between the actual output of an economy and its potential output. For the output, we use "Real Gross Domestic Product" (mnemonics "GDPC96") and "Population Level" (mnemonics "CNP16OV"). The data is transformed as: 100*ln(GDPC96/CNP16OV). For the potential output, we use "Real Potential Gross Domestic Product" (mnemonics "GDPPOT"). The data is transformed as: 100*ln(GDPPOT/CNP16OV). All are available from the Federal Reserve Bank of St. Louis's FRED database. "GDPDEF" is available from 1947Q1 - 2022Q1. "PRS84006173" is available from 1947Q1 - 2022Q1. "HCOMPBS" is available from 1947Q1 - 2022Q1. "GDPC96" is available from 1947Q1 - 2017Q2. (This series has been discontinued. It was a duplicate of "GDPC1", which will continue to be updated.) "CNP16OV" is available from 1948Q1 - 2022Q1. "GDPPOT" is available from 1949Q1 - 2031Q4.

The government spending shocks is Ramey's (2011) military news variable. The update series (up to 2015) that we use is from Ramey and Zubairy (2018).

The excess bond premium shock is from Gilchrist and Zakrajsek (2012). The updated data we use is available at:
https://www.federalreserve.gov/econres/notes/feds-notes/updating-the-recession-risk-and-the-excess-bond-premium-20161006.htm and available at:
https://www.federalreserve.gov/econres/notes/feds-notes/ebp_csv.csv

The monetary policy shock is from Romer and Romer (2004) from 1969Q1 to 2007Q4 and the updated data we use is available from Wieland and Yang (2020) at:
https://www.openicpsr.org/openicpsr/project/135741/version/V1/view.[15]

[15] The Romer and Romer (2004) monetary policy shock series is available from 1969Q1 to 2007Q4. The excess bond premium series (EBP) is available from 1973Q1 to 2022. The government spending shock series is available from 1989Q1 to 2015Q4. Both the monetary policy and the EBP shock series are originally available at the monthly frequency, and we aggregated them at the quarterly frequency by summing the monthly values.

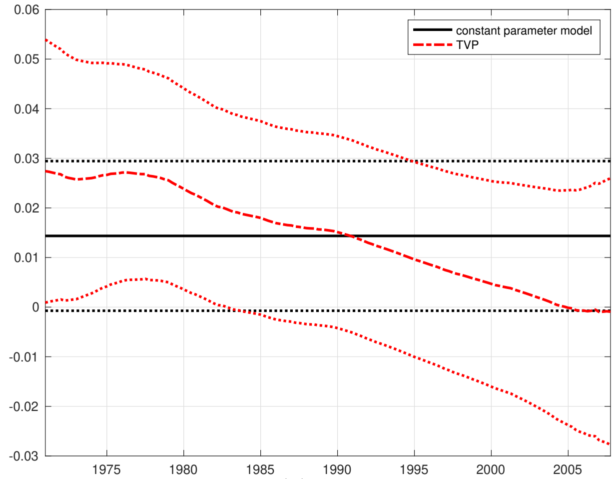

(a) λ_t

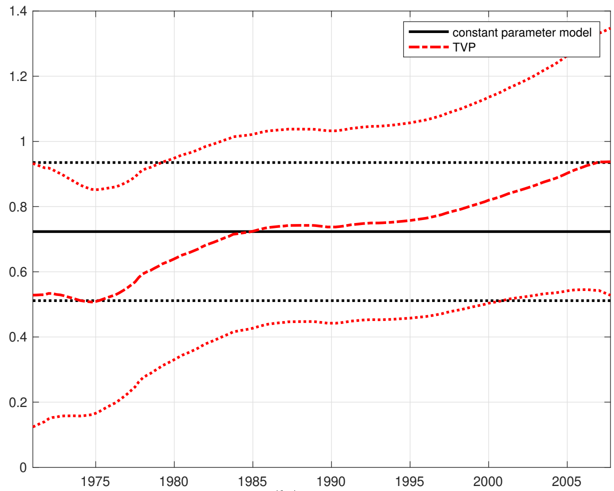

(b) $\gamma_{f,t}$

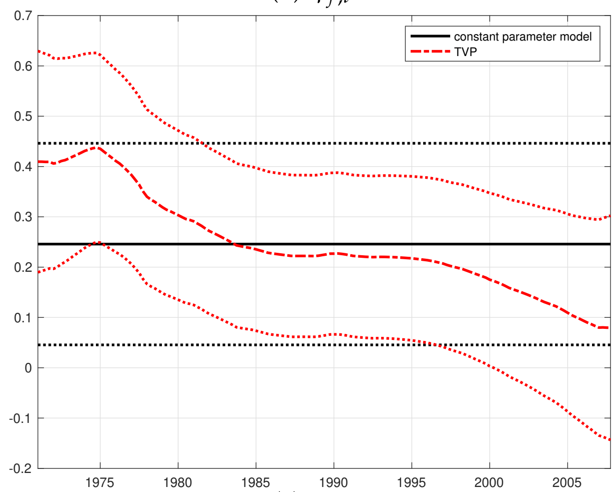

(c) $\gamma_{b,t}$

Figure 8: **The time-varying Phillips curve using labor share as the forcing variable.** The figure shows the estimated TVP-IV coefficients as well as the full-sample estimates together with 90% confidence bands. The sample is 1970Q1:2008Q1. The specification is the same as GGLS. We estimate the hybrid NKPC with one lag of inflation and the labor share as forcing variable. The GGLS instrument set includes four lags of inflation and two lags of the labor share, wage inflation, and quadratically-detrended output. HAC-robust variance estimates are implemented with 4 lags, following Mavroeidis, Plagborg-Møller and Stock (2014).

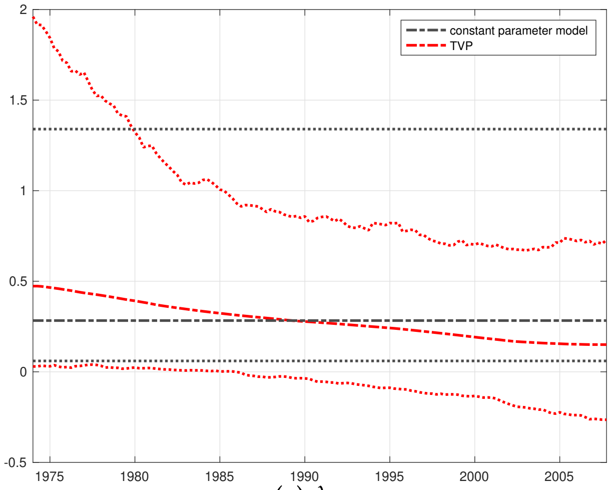

(a) λ_t

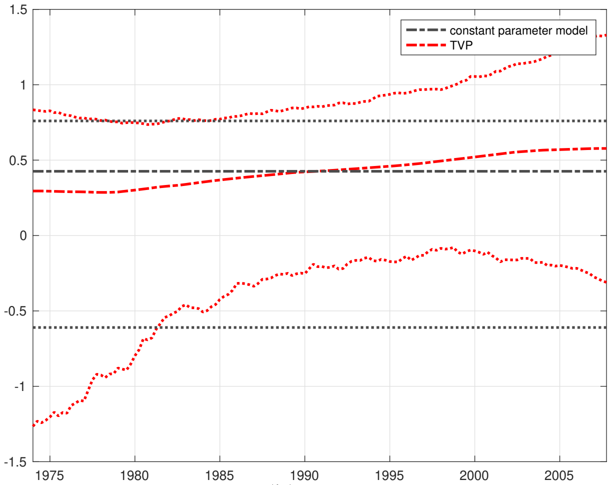

(b) $\gamma_{f,t}$

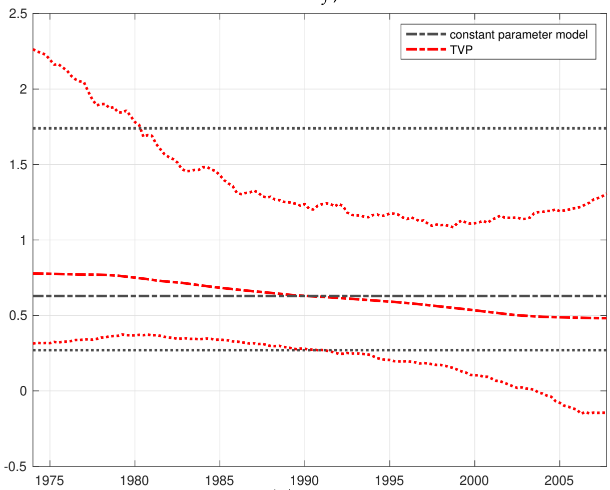

(c) $\gamma_{b,t}$

Figure 9: **The time-varying Phillips curve using monetary policy shocks as instruments and output gap as the forcing variable.** The figure shows the TVP-IV estimated coefficients and the full-sample estimates, together with the 90% confidence bands robust to weak instruments. The sample is 1974Q1:2007Q4. The specification uses as instruments the Almond parameterization of 20 lags of the Romer and Romer monetary policy shocks, and the output gap obtained via the Hodrick-Prescott filter as the forcing variable. HAC-robust variance estimates are implemented with 5 lags, following Barnichon and Mesters (2020). The bands are smoothed using a 7-quarter moving average.

8 References

Atkeson, A. and L. E. Ohanian (2001). "Are Phillips Curves Useful for Forecasting Inflation?" Federal Reserve of Minneapolis Quarterly Review 25(1), 2–11.

Ball, L. and S. Mazumder (2011). "Inflation Dynamics and the Great Recession," Brookings Papers on Economic Activity 42, 337–405.

Ball, L. and S. Mazumder (2019). "A Phillips Curve with Anchored Expectations and Short-Term Unemployment," Journal of Money, Credit and Banking 51, 111–137.

Barnichon, R. and G. Mesters (2020). "Identifying Modern Macro Equations with Old Shocks," The Quarterly Journal of Economics 135(4), 2255-2298.

Barnichon, R. and G. Mesters (2021). "The Phillips Multiplier," Journal of Monetary Economics 117, 689-705.

Baxter, M. and R.G. King (1999), "Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series," Review of Economics and Statistics 81(4), 575-593.

Bergholt, D., F. Furlanetto, and E. Vaccaro-Grange (2022). "The Death and Resurrection of the US Phillips Curve," mimeo.

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

Bobeica, E. and M. Jarociński (2019). "Missing Disinflation and Missing Inflation: A VAR Perspective," International Journal of Central Banking 15(1), 199-232.

Cogley, T. and T.J. Sargent (2005), "Drifts and volatilities: monetary policies and outcomes in the post WWII US", Review of Economic Dynamics 8(2), 262-302.

Coibion, O. and Y. Gorodnichenko (2015). "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," American Economic Journal: Macroeconomics 7, 197–232.

Del Negro, M., M. Lenza, G. Primiceri and A. Tambalotti (2020). "What's Up with the Phillips Curve?", Brooking Papers on Economic Activity, 301-357.

Galí, J. and L. Gambetti (2018). "Has the U.S. Wage Phillips Curve Flattened? A Semi-Structural Exploration," mimeo, CREI.

Galí, J. and M. Gertler (1999). "Inflation dynamics: A structural econometric analysis," Journal of Monetary Economics 44, 195–222.

Galí, J. M. Gertler and D. López-Salido (2005). "Robustness of the estimates of the hybrid New Keynesian Phillips curve," Journal of Monetary Economics 52(6), 1107-1118.

Ganics, G., I. Atsushi and B. Rossi (2021). "Confidence Intervals for Bias and Size Distortion in IV and Local Projections-IV Models," Journal of Business & Economic Statistics, 39(1), pp.307-324.

Gilchrist, S. and E. Zakrajsek (2012). "Credit Spreads and Business Cycle Fluctuations," American Economic Review 102, 1692–1720.

Hall, R. E. (2011). "The long slump," American Economic Review, 101(2), 431-69.

Hazell, J., J. Herreno, E. Nakamura, and J, Steinsson (2022). "The Slope of the Phillips Curve: Evidence from U.S. States," Quarterly Journal of Economics, 137(3), 1299-1344.

Inoue, A., B. Rossi and Y. Wang (2022). "Local Projections in Unstable Environments: How Effective is Fiscal Policy?" CEPR Discussion Paper No. DP17134.

Lazarus, E., Lewis, D.J., Stock, J.H. and Watson, M.W., 2018. HAR Inference: Recommendations for Practice. Journal of Business & Economic Statistics, 36(4), pp.541-559.

Lewis, D.J. and Mertens, K., (2022). A Robust Test for Weak Instruments with Multiple Endogenous Regressors. Available at SSRN 4144103.

Mavroeidis, S., M. Plagborg-Moller, and J. H. Stock (2014). “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve,” Journal of Economic Literature 52, 124–188.

McLeay, M. and S. Tenreyro (2019). “Optimal Inflation and the Identification of the Phillips Curve,” in NBER Macroeconomics Annual 2019, volume 34.

Newey, W.K. and West, K.D. (1994). ”Automatic Lag Selection in Covariance Matrix Estimation,” The Review of Economic Studies, 61(4), pp.631-653.

Newey, W.K. and West, K.D. (1987). “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” Econometrica 55, no. 3 (1987): 703–8.

Primiceri, G. E. (2006). “Why Inflation Rose and Fell: Policy-Makers’ Beliefs and U.S. Postwar Stabilization Policy,” Quarterly Journal of Economics 121, 867–901.

Ramey, V.A., 2011. ”Identifying government spending shocks: It’s all in the timing,” The Quarterly Journal of Economics, 126(1), pp.1-50.

Ramey, V.A. and Zubairy, S., 2018. ”Government spending multipliers in good times and in bad: evidence from US historical data,” Journal of political economy, 126(2), pp.850-901.

Romer, C.D. and D.H. Romer (2004). ”A New Measure of Monetary Shocks: Derivation and Implications,” American Economic Review 94(4), 1055–84.

Stock, J.H. and M.W. Watson (2007). “Why Has U.S. Inflation Become Harder to Forecast?” Journal of Money, Credit and Banking 39, 3–33.

Stock, J.H. and M.W. Watson (2008). “Phillips Curve Inflation Forecasts,” NBER Working Papers 14322.

Stock, J.H., Watson, M.W. (2020). ”Slack and cyclically sensitive inflation,” Journal of Money, Credit and Banking, 52(S2), 393-428.

Tenreyro, S. and G. Thwaites (2016). ”Pushing on a String: US Monetary Policy Is Less Powerful in Recessions,” American Economic Journal: Macroeconomics 8(4), 43-74.

Wieland, J.F. and M. Yang (2020), “Financial Dampening,” Journal of Money, Credit and Banking 52(1), 79–113.