DID MONETARY POLICY KILL THE PHILLIPS CURVE? SOME SIMPLE ARITHMETICS*

D. BERGHOLT†, F. FURLANETTO‡ AND E. VACCARO-GRANGE§

May 2022

Abstract: Price inflation in the U.S. economy has been remarkably stable over the past 25 years, in spite of large fluctuations in real economic activity. This observation has led some to believe that the Phillips curve has flattened. We argue that this viewpoint may be premature unless one accounts explicitly for all supply-side variation in data. In fact, we show that it is crucial to control for *an entire array* of supply shocks (and not only cost-push shocks) when evaluating alternative explanations for the puzzling behavior of inflation. Equipped with a combination of New Keynesian theory and SVAR models, we decompose the unconditional variation in data into the components driven by demand and supply, respectively. This allows us to conduct a simple yet novel accounting exercise, which reveals that the Phillips curve remains relatively stable once supply shocks are properly controlled for. The demand curve, in contrast, has flattened substantially. Our results are fully consistent with an explanation based on a more aggressive monetary policy response to inflation.

Keywords: *Inflation, the Phillips curve, monetary policy, structural VAR models.*
JEL Classification: *C3, E3, E5.*

1 INTRODUCTION

Since the mid 1990s price inflation has remained remarkably stable in the United States, even in presence of large cycles in economic activity, as shown in Figure 1. This tendency has been even more pronounced in the years preceding the COVID recession. During the pre Great Recession boom, inflation was stubbornly stable, barely above 2 percent. During the Great Recession, in face of the largest decline in real economic activity since

*This working paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. We would like to thank for useful comments our discussant Benjamin Wong and Knut Are Aastveit, Regis Barnichon, Martin Beraja, Fabio Canova, Jeff Campbell, Michael Ehrmann, Luca Fosso, Luca Gambetti, Domenico Giannone, Nicolo Maffei Faccioli, Elmar Mertens, Emi Nakamura, Giorgio Primiceri, Federico Ravenna, Robert Rich, Jon Steinsson, Carl Walsh, Chris Wolf as well as seminar participants at Norges Bank, Federal Reserve Bank of Cleveland and conference in Sydney, CEBRA virtual conference 2020, virtual SNDE conference 2022, New Developments in Business Cycle Analysis 2021 conference hosted virtually by Deutsche Bundesbank, EABCN conference on Challenges in Empirical Macroeconomics since 2020 in Eltville, Barcelona Summer Forum 2022 and CEF conference 2022 in Dallas. This work was supported by the French National Research Agency Grant ANR-17-EURE-0020, and by the Excellence Initiative of Aix-Marseille University - A*MIDEX.

†Norges Bank. Email: drago.bergholt@norges-bank.no
‡Norges Bank and BI Norwegian Business School. Email: francesco.furlanetto@norges-bank.no
§New York University Abu Dhabi. Email: etienne.vaccaro-grange@nyu.edu

the Great Depression, inflation (measured with the GDP deflator) declined by only one percent. In the aftermath of the Great Recession, real economic activity recovered (albeit slowly), unemployment reached a 50-year low slightly under 4 percent but inflation remained consistently below 2 percent.

Why has inflation been so stable? At least three resolutions have been proposed to solve the puzzle. The first (and perhaps most widely accepted) explanation points towards a decline in the slope of the Phillips curve, as documented by Ball and Mazumder (2011), Blanchard (2016), Stock and Watson (2019) and Del Negro, Lenza, Primiceri, and Tambalotti (2020) among others. Explanations for such a flattening include global factors, such as increased import competition, rising market concentration and changes in the network structure of the US production sector (cf. Forbes (2019), Obstfeld (2020), Heise, Karahan, Şahin, et al. (2020), Rubbo (2020) and Ascari and Fosso (2021) among others) or the heterogeneity of financial conditions across firms (cf. Gilchrist, Schoenle, Sim, and Zakrajšek (2017)). In such a scenario, demand shocks have large effects on real economic activity but barely affect inflation. A second explanation, possibly complementary to the first, highlights the role of monetary policy that may have become more aggressive over time with respect to achieving inflation stability (McLeay and Tenreyro (2020)). According to this view, the Phillips curve is alive and stable but demand shocks leave no footprint and generate limited fluctuations in inflation (but also in measures of slack in the economy). A third possibility is that the correlation between inflation and real economic activity declines because the relative importance of demand and supply shocks has changed (Galí and Gambetti (2009), Gordon (2013) and Hobijn (2020)). One could imagine that supply shocks became more important over time or more concentrated in specific periods, as it was the case for oil shocks in the aftermath of the Great Recession (Coibion and Gorodnichenko (2015)). Under that view, inflation may co-move less with real economic activity even if both the Phillips curve and the behavior of monetary policy are perfectly stable. We refer to these three broad explanations for the stability of inflation as the *slope hypothesis*, the *policy hypothesis* and the *shock hypothesis*.

In this paper, we run a horserace between the three main explanations of the inflation puzzle. In order to achieve our goal, we estimate a simple bivariate Structural Vector Autoregression (SVAR) model identified with sign restrictions. We use data on the GDP deflator and on the measure of the output gap computed by the Congressional Budget Office (CBO) to identify demand and supply shocks. Such a simple model is sufficient to disentangle the three explanations if combined with insights derived from the simple three equation New Keynesian model summarized in the textbook by Galí (2008). The only modification from the textbook model is that we define the output gap in deviation from trend (or steady state), and not in deviation from the flexible prices equilibrium, to better match its empirical proxy computed by the CBO. Our main result is that the policy hypothesis is the winner of our horserace: according to our results, inflation has not moved much over recent years simply because the Federal Reserve has prevented inflation (or deflation) to materialize.

We argue that identifying *both* demand and supply shocks and inspecting changes in their transmission mechanisms is crucial to obtain our evidence in favor of the *policy hypothesis*. In fact, standard New Keynesian theory provides clear testable implications to disentangle the three stories based on the conditional correlations between inflation and output gap in response to demand and supply shocks. Let us consider first the slope

Figure 1: The evolution of inflation and the output gap

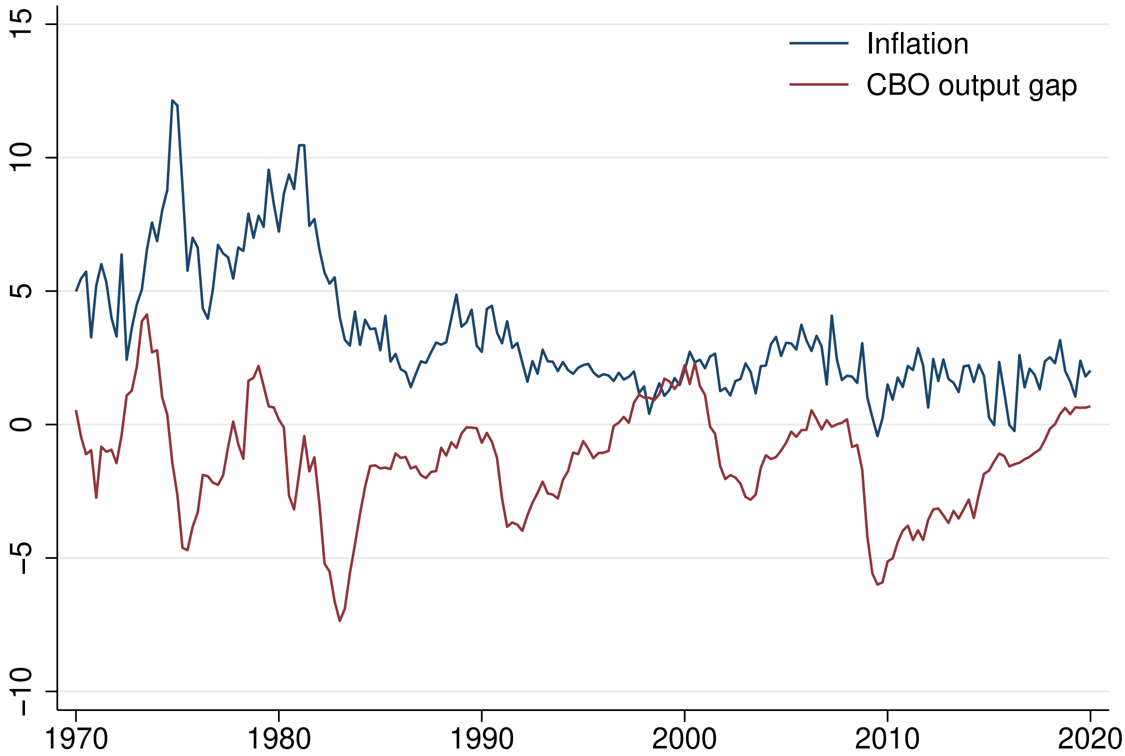

story and imagine a decline in the slope of the New Keynesian Phillips curve. The New Keynesian model implies that the conditional correlation between inflation and the output gap should decrease in response to a demand shock. Given the lower slope of the supply curve, a demand shock will have larger effects on real economic activity and lower effects on inflation. In contrast, the conditional correlation between inflation and output gap is unaffected in response to a supply shock. In fact, the impact of a supply shock is determined by the slope of the aggregate demand curve which is unchanged under the slope hypothesis. The opposite pattern arises if the monetary policy authority starts responding more aggressively to inflation. The conditional correlation between inflation and output gap does not change in response to demand shocks, although the volatility of the two variables decreases and reaches the Divine Coincidence limit in the case of strict inflation target. In this case, however, the conditional correlation is expected to decline, i.e. to become less negative, in response to supply shocks: a supply shock will have limited effect on inflation (because of the endogenous monetary policy response which leads to a flattening of the IS curve) but large effects on the output gap, as long as the output gap is measured in deviation from trend. Finally, if the volatility of shocks changes none of the two conditional correlations is expected to change since the propagation of shocks is unaffected. Only the unconditional correlation will change. All in all, the slope story implies a change only in the conditional correlation to demand shocks, the policy story implies a change only in the conditional correlation to supply shocks while the shocks story implies no change in both correlations.

In addition, the basic New Keynesian model has some testable implications also in terms of volatilities. As already mentioned, aggregate demand shocks have a larger effect on output if the supply curve is flatter. In contrast, aggregate demand shocks have a smaller effect on output if the demand curve is flatter (and no effect at all under the limiting case of strict inflation targeting). Therefore, aggregate demand shocks should explain a larger share of output volatility under the slope story and a smaller share of output volatility under the policy story. Note, however, that changes in the relative importance of shocks for output could also be rationalized by the shocks story. In that sense, only the implications for the slopes (conditional correlations) are mutually exclusive and jointly exaustive.

We estimate our simple bivariate SVAR model over two adjacent sample periods (1968:Q4-1994:Q4 and 1995:Q1-2019:Q4) in order to inspect changes in the propagation of shocks and compare them with predictions from theory. We decompose fluctuations in the data into two components: one driven by demand shocks and one driven by supply shocks. Such a historical decomposition allows us to compute the correlation between inflation and the output gap conditional on the two shocks. These conditional correlations correspond to the slopes of the regression lines fitting the cloud of data points generated by the two shocks. We find three results that can be directly compared with the theoretical predictions of the New Keynesian model. First, the conditional correlation between inflation and the output gap in response to demand shocks hardly changes at all between the two samples. Second, the conditional correlation between inflation and the output gap in response to supply shocks is substantially reduced in the second sample. Put differently, supply shocks have larger effects on real economic activity and more limited effects on inflation. Third, demand shocks are less important in the variance decomposition for the output gap in the second part of the sample. All these three results are consistent with the policy hypothesis and survive a battery of sensitivity checks. In particular, our results are confirmed when we include alternative measures of inflation or real economic activity in the SVAR and when we consider larger specifications of the SVAR including data on inflation expectations or the interest rate.

We are definitely not the first highlighting the crucial role of supply shocks for inflation dynamics (cf. Hobijn (2020) and Gordon (2013) among others). Our contribution is to show that changes in the transmission of supply shocks can be as informative as changes in the transmission of demand shocks to explain the inflation puzzle. In addition, we stress that it is important to control for *all* supply shocks while the literature often controls only for a very specific kind of supply shock, the so-called price mark-up shock. We show that the residual in the Phillips curve equation is a combination of all supply shocks if the output gap is measured in deviation from a smooth trend, as it is the case for the CBO measure of the output gap (cf. Coibion, Gorodnichenko, and Ulate (2018)). The residual is equivalent to a mark-up shock only when the output gap is defined as the welfare relevant concept in the New Keynesian literature, which is very different from the trend concept used in empirical work.[1]

Our paper contributes to the empirical literature studying the drivers of the connection between inflation and real economic activity. Traditionally, most papers discuss approaches and challenges to the estimation of the New Keynesian Phillips curve in a single

[1]In the New Keynesian model potential output does not necessarily respond smoothly to shocks as does the CBO estimate of potential output (as shown by Coibion et al. (2018)). Technology and labor supply shocks, for example, move potential output more than actual output in the standard New Keynesian model (see Bilbiie and Melitz (2020) for an extension with entry and exit which does not feature that property).

equation framework (cf. Galí and Gertler (1999), Sbordone (2002) and Kleibergen and Mavroeidis (2009) among many others). However, Mavroeidis, Plagborg-Møller, and Stock (2014) highlight how estimates of Phillips curve parameters are subject to a weak instrument problem and conclude that new datasets and new identification approaches are needed to reach an empirical consensus.

In terms of new data sets, Imbs, Jondeau, and Pelgrin (2011) estimate Phillips curves at the sectoral level using French data and then derive implications for monetary policy. More recently, several papers rely on regional data. Hazell, Herreno, Nakamura, and Steinsson (2022) estimate the slope of the Phillips curve in the cross section of U.S. states using newly constructed state-level price indexes for non-tradeable goods. They find only a modest decline in the slope of the Phillips curve since the 1980s. In addition, they use a multi-region model to infer the aggregate slope and conclude that there is no missing disinflation or missing inflation over the most recent business cycles. The same result is found by Fitzgerald, Jones, Kulish, Nicolini, et al. (2020) using city-level and state level data for the US while Beraja, Hurst, and Ospina (2019) combine regional and aggregate data to investigate the connection between wages and unemployment with a special focus on the slope of the wage Phillips curve.

In terms of new identification strategies, the literature has developed to find better instruments to estimate the Phillips curve. Barnichon and Mesters (2020) use as instruments independently identified structural shocks rather than predetermined variables. They estimate the Phillips curve using monetary shocks as instruments and find that conventional methods substantially underestimate the slope of the Phillips curve. A few recent papers have taken a multivariate approach using SVAR models to isolate the variation in the data due to demand shocks. A prominent example is Del Negro et al. (2020) who show that inflation barely reacts in the post 1990s sample in response to shocks to the excess bond premium, i.e. shocks that propagate through the economy like a typical demand shock, while it used to significantly respond before. Relatedly, Ascari and Fosso (2021) estimate a SVAR with common trends and find evidence of a lower response of inflation to business cycle shocks in recent years because of limited pass-through from wages to prices. A benefit of our set-up with respect to the above mentioned papers, is that we can evaluate all the three main explanations for the inflation puzzle in a unified framework based on the mutually exclusive implications derived from the theoretical model.[2] The paper closest to us is perhaps Galí and Gambetti (2019) who use a SVAR to purge the data from the variation induced by wage mark-up shocks to estimate the slope of the wage Phillips curve. Based on our argument on the measurement of the output gap, we argue that it is important to control for the variation induced by *all* supply shocks.

The rest of the paper is organized as follows: Section 2 discusses the structural relationship between output and inflation using a textbook New Keynesian model. Section 3 describes our methodological approach, Section 4 documents the main empirical results. Section 5 provides robustness tests, while Section 6 concludes.

[2] Historically, the VAR approach has been used mainly to study the long-run trade-off between inflation and unemployment (cf. King and Watson (1994), Cecchetti and Rich (2001), Benati (2015), Barnichon and Mesters (2021) and Ascari, Bonomolo, and Haque (2022)).

2 Theoretical discussion

We start with a log-linearized, textbook New Keynesian model (see Woodford (2003) and Galí (2008) for further details) and then briefly look at some extensions. The model is summarized below:

$$y_t = \mathbb{E}_t y_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - u_t\right) \tag{1}$$

$$y_t = a_t + n_t \tag{2}$$

$$w_t = \psi_t + \sigma y_t + \varphi n_t \tag{3}$$

$$mc_t = w_t - a_t \tag{4}$$

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \lambda mc_t + z_t \tag{5}$$

Conditional on monetary policy and exogenous disturbances, these five equations characterize the dynamics of five endogenous variables, all defined in log deviations from their respective steady state (or equivalently, from trend) values: the output gap y_t, hours worked n_t, the real wage w_t, real marginal costs mc_t, and price inflation π_t. The variables a_t, ψ_t, and z_t are interpreted as a productivity shock, a labor supply shock, and a cost-push shock respectively. u_t is a demand shock or discount factor shock. All parameters have the usual interpretation, including $\lambda = \frac{(1-\theta)(1-\beta\theta)}{\theta}$, with θ being the Calvo probability in any given period of not being able to adjust the price. The model is closed with a specification of monetary policy. As a baseline, we assume that the nominal interest rate i_t is determined by a simple Taylor rule:

$$i_t = \phi_\pi \pi_t + \phi_y y_t + m_t \tag{6}$$

m_t is interpreted as a monetary policy shock. One can simplify the model by inserting for equations (2)-(4) into (5) and arrive at the canonical New Keynesian Phillips curve:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t + s_t, \tag{7}$$

where $s_t = z_t + \lambda\psi_t - \lambda\left(1+\varphi\right) a_t$ collects the three supply shocks. Our object of interest, the slope of the Phillips curve, is given by $\kappa = \lambda\left(\sigma + \varphi\right)$. A flattening of the structural Phillips curve amounts to a decline in κ.

Important for our purpose, the model equations are derived as approximations of the *log-linear deviations from steady state or trend*, not from the counterfactual equilibrium under flexible prices. While the flex-price equilibrium gap of output is important for welfare and appears frequently in textbooks, we believe that output in deviation from some trend is more in line with the operational definitions used by statistical agencies. The measure of potential output published by the CBO, for example, is a slow-moving variable rather than an erratic series driven by short-term fluctuations. It is exactly the latter one would expect, had one defined potential as the flex-price outcome. Our choice to focus on trend gaps has important implications for the Phillips curve, as all three supply shocks enter it directly.[3] Thus, with this output gap definition the divine coincidence (Blancard and Gali, 2007) seizes to hold, and we need to control for all supply shocks

[3]Only the cost-push shock enters the Phillips curve if we consider output in deviation from its counterfactual when prices are flexible. The other supply shocks are part of flex-price output in this case.

when estimating equation (7). Finally, note that the two demand shocks u_t and m_t do not enter the Phillips curve. This illustrates that demand shocks may serve as valid and relevant instruments for y_t.

In order to discuss the challenges associated with estimation of κ, and to highlight our proposed identification strategy, we find it instructive to work with the model's solution. To this end we collect the two demand shocks in $d_t = m_t - u_t$ and impose the heroic assumption that d_t and s_t are independently and identically distributed with variances σ_d^2 and σ_s^2, respectively.[4] Analytical solutions for output and inflation follow:

$$y_t = \frac{1}{\sigma + \phi_y + \kappa\phi_\pi}\left(d_t - \phi_\pi s_t\right)$$

$$\pi_t = \frac{1}{\sigma + \phi_y + \kappa\phi_\pi}\left[\kappa d_t + \left(\sigma + \phi_y\right) s_t\right]$$

Suppose that we estimate, by OLS, the simple regression equation

$$\pi_t = \gamma y_t + \varepsilon_t$$

using the *unconditional* data generated from this stylized model. Given the analytical solutions, the estimator given by $\gamma^{OLS} \equiv \frac{cov(\pi_t, y_t)}{var(y_t)}$ can be written in closed form as

$$\gamma^{OLS} = \frac{\kappa\left(\sigma_m^2 + \sigma_u^2\right) - \phi_\pi\left(\sigma + \phi_y\right)\sigma_z^2 - \frac{\phi_\pi\kappa^2(\sigma+\phi_y)}{(\sigma+\varphi)^2}\sigma_\psi^2 - \frac{\phi_\pi\kappa^2(\sigma+\phi_y)(1+\varphi)^2}{(\sigma+\varphi)^2}\sigma_a^2}{\sigma_m^2 + \sigma_u^2 + \phi_\pi^2\sigma_z^2 + \left(\frac{\phi_\pi\kappa}{\sigma+\varphi}\right)^2\sigma_\psi^2 + \left(\frac{\phi_\pi\kappa(1+\varphi)}{\sigma+\varphi}\right)^2\sigma_a^2}$$

$$= \frac{\kappa\sigma_d^2 - \phi_\pi\left(\sigma + \phi_y\right)\sigma_s^2}{\sigma_d^2 + \phi_\pi^2\sigma_s^2},$$

where $\sigma_d^2 = \sigma_m^2 + \sigma_u^2$ represents the total variance of demand side shocks, and $\sigma_s^2 = \sigma_z^2 + \left(\frac{\kappa}{\sigma+\varphi}\right)^2\sigma_\psi^2 + \left(\frac{\kappa(1+\varphi)}{\sigma+\varphi}\right)^2\sigma_a^2$ represents the total variance of supply side shocks. Note that σ_s^2 is a function of κ, an observation that we will exploit later.

What can γ^{OLS} tell us about κ? A couple of implications can immediately be drawn from this simple model: first, the estimator based on *unconditional* data, given by

$$\gamma_u = \frac{\kappa - \phi_\pi\left(\sigma + \phi_y\right)\frac{\sigma_s^2}{\sigma_d^2}}{1 + \phi_\pi^2\frac{\sigma_s^2}{\sigma_d^2}},$$

is almost surely biased downwards relative to κ. The bias stems from two supply driven sources of variation: (i) the variance in y_t given by $\phi_\pi^2\sigma_s^2$, and (ii) the negative covariance between π_t and y_t, given by $-\phi_\pi\left(\sigma + \phi_y\right)\sigma_s^2$. Second, the bias might evolve over time. Suppose that one decides to estimate the Phillips curve over different sub-samples of unconditional data. Then, one may erroneously conclude that the Phillips curve has flattened even when the true slope κ has remained unchanged. For example, one may

[4] This i.i.d. assumption greatly simplifies the notation. It is, however, relaxed in the appendix where we instead consider the more common assumption that shocks follow separate AR(1) processes. None of our conclusions are altered in this case.

reach this conclusion if the stance of monetary policy has changed towards stricter inflation targeting ($\phi_\pi \uparrow$), or if supply shocks have become more volatile relative to demand shocks ($(\sigma_s^2/\sigma_d^2) \uparrow$). In fact, it is clear from the expression above that unconditional data deliver an inherent identification problem: there is indeed no way to distinguish a change in κ from changes in policy or shocks, based on estimates of γ_u.

2.1 CONDITIONAL SLOPE REGRESSIONS

It is only when $\sigma_s^2 = 0$, i.e. when *all* variation in the data due to supply shocks is filtered out, that the potentially time varying bias is neutralized. This brings us another lesson: suppose that we are able–somehow–to purge the data for all variation due to supply side shocks. The estimator *conditional on demand* shocks follows:

$$\gamma_d = \kappa$$

This estimator is unbiased: if we somehow are able to trace changes in γ_d, then it seems reasonable to attribute those to changes in the structural Phillips curve slope κ. The stance of policy, or the composition of shocks, play no role for γ_d in our framework. Next, suppose that we are able–somehow–to purge the data for all variation due to demand side shocks. The estimator *conditional on supply* shocks follows:

$$\gamma_s = -\frac{\sigma + \phi_y}{\phi_\pi}$$

This estimator is purely a function of the two policy parameters and σ, all of which are important for the responsiveness of *demand* to prices. Stricter inflation targeting, in particular, shifts γ_s up towards zero. Changes in the Phillips curve slope or in the composition of shocks play no role for γ_s in our framework. The fact that both γ_d and γ_s are independent of shock volatilities allows us to ignore the role of κ for σ_s^2. Moreover, it allows us to formalize and set apart a third explanation, the shocks story: a rise in the relative importance of supply shocks has no effect on our conditional regression slopes, and it is the only explanation among those considered that has this particular property.

Figure 2 illustrates the implications from our discussion so far, and serves to highlight a number of testable predictions pursued in this paper. In the figure we report a battery of scatter plots with data simulated from the theoretical model. Scatter plots are augmented with simple regression slopes.[5] Each column in the figure represents a separate explanation–the slope story is demonstrated in the first column, the policy story in the second, and the shocks story in the third.[6] Each row represents a particular dataset–unconditional data augmented with γ_u are shown in the first row, data conditional only on demand shocks and augmented with γ_d in the second, and data conditional on supply shocks and augmented with γ_s in the third. Note that, in each column, the scatter plots in

[5]The model's parameters are set to standard values: the baseline calibration includes $\beta = 0.99$, $\sigma = 1$, $\varphi = 2$, $\phi_\pi = 1.5$, $\phi_y = 0.125$, and $\theta = 0.75$. The standard deviations of shocks are set to $\sigma_u = \sigma_d = 1$, $\sigma_a = \sigma_\psi = 0.2$, and $\sigma_z = 0.05$ respectively. We let each of the shocks follow an AR(1) with an autoregressive coefficient equal to 0.75.

[6]The slope story is illustrated with a rise in the Calvo parameter from 0.75 in sample one (blue) to 0.875 in sample two (red). The policy story is shown as a rise in ϕ_π from 1.5 to 2.5 across the samples, while the shocks story is a decline in the volatility of demand shocks, from 1 to 0.65.

Figure 2: Three alternative explanations – simulated data

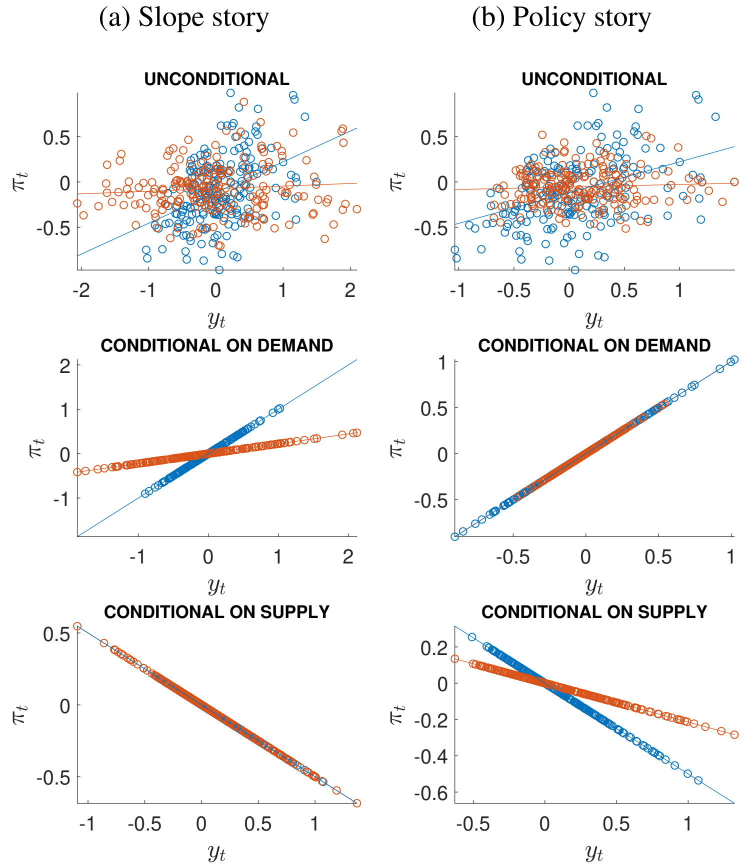

the second and third rows sum to the unconditional scatter plot in the first row. Consider first the slope story. In this case the following applies: if the observed weakening of the statistical relationship between output and inflation (first row) is predominantly a consequence of a flatter Phillips curve slope κ, then we should see a decline in γ_d (second row), combined with a relatively stable γ_s (third row). If, instead, the weakened relationship is predominately driven by stricter inflation targeting, then we should find a relative stable γ_d over time, combined with a flattening (towards less negative values) of γ_s. Finally, if the main explanation is that supply side shocks have become more volatile relative to demand side shocks, then we should find relatively stable estimates of both γ_d and γ_s across sub-samples. The three alternative explanations are both mutually exclusive and jointly exhaustive in our framework. That is, given a weakened relationship between output and inflation, we *must see* it in either one of the two slopes or in a change in the shock composition. Of course, in practice it is likely that a linear combination of the three stories can best describe data. We quantify the relative importance of each story in the empirical section.

2.2 Conditional variances

While the slope regressions introduced above provide us with a useful apparatus to disentangle the competing explanations, we also acknowledge their limitations. Perhaps most importantly, a flattening of γ_s is strictly speaking just a flattening in the IS equation in (1). As such, it may come about from changes in financial markets, or other... However, to further asses changes in the slope of demand we proceed by decomposing the variance of output. This variance decomposition leaves us with the following expression:

$$VD\left(y|d\right) = \frac{\sigma_d^2}{\sigma_d^2 + \phi_\pi^2 \sigma_z^2 + \left(\frac{\phi_\pi \kappa}{\sigma+\varphi}\right)^2 \sigma_\psi^2 + \left(\frac{\phi_\pi \kappa(1+\varphi)}{\sigma+\varphi}\right)^2 \sigma_a^2}$$

$VD\left(y|d\right)$ is the share of the total variance in output that is attributed to the two demand shocks.[7] It follows that we can exploit this variance decomposition to further disentangle the competing explanations: if the slope story ($\kappa \downarrow$) is true, then we should observe a rise in $VD\left(y|d\right)$, i.a. a more important role for demand shocks over time. The policy story ($\phi_\pi \uparrow$), in contrast, implies a decline in $VD\left(y|d\right)$. These implications for conditional variances are also visible in Figure 2. When comparing sub-samples, we see that the demand shocks become relatively more important for output when the slope story applies, while the opposite is true for the policy story. A decline in the role of demand for output is naturally also the results given by the shocks story. However, the shocks story would imply a constant estimate of γ_s, as emphasized earlier. In the appendix, we show that our results hold also in a richer model with several bells and whistles added to the system. Moreover, we show that the results do not depend on our abstraction from expectations.

3 Empirical approach

The empirical approach we pursue in this paper is essentially a two-step procedure: in the first step, referred to as the filtering step, we decompose the data into two components driven by demand and supply shocks, respectively. This is done with SVAR models estimated with Bayesian techniques over the adjacent samples 1968:Q4-1994:Q4 and 1995:Q1-2019:Q4. In the second step, referred to as the regression step, we perform our joint test on slopes and variance decompositions to evaluate the merits of the three proposed explanations for the inflation puzzle. We thus run regressions on the filtered data in order to make inference about the conditional relationship between output and inflation in both samples. In addition, we inspect variance decompositions across samples.

The structural representation of our simple VAR model can be written as follows:

$$B^{-1} Y_t = \sum_{k=1}^{p} Z_k Y_{t-k} + \varepsilon_t$$

Y is a vector of size n, $B \in \mathcal{M}_{n\times n}(\mathbb{R})$, and $Z_k \in \mathcal{M}_{n\times n}(\mathbb{R})$ are the matrices of structural parameters, and $\varepsilon_t \in \mathbb{R}^n$ is the vector of structural shocks with $\varepsilon_t \sim \mathcal{MVN}(0, I_n)$. We

[7] A further decomposition into the two different demand shocks would lead to the same expression for each of the shocks, except that σ_d^2 in the nominator would be replaced with either σ_m^2 or σ_u^2, respectively.

Table 1: Sign restrictions - VAR models

(A) Baseline SVAR	Demand $\uparrow$	Supply $\downarrow$	
Inflation	+	+	
Output gap	+	-	
(B) SVAR with policy shocks	Demand $\uparrow$	Supply $\downarrow$	Monetary Policy $\uparrow$
Inflation	+	+	+
Output gap	+	-	+
Interest (or shadow) rate	+	$\star$	-
(C) SVAR with inflation expectations	Demand $\uparrow$	Supply $\downarrow$	Residual
Inflation	+	+	+
Output gap	+	-	$\star$
Inflation expectations	+	+	-

Note: Restrictions are imposed only on impact. $\star$ means that no restriction is imposed.

assume that all the roots of the model's characteristic polynomial lie outside the unit disk, so that the VAR model is stationary. Our baseline VAR contains two variables:

$$Y_t = (\pi_t, y_t)'$$

where π_t is inflation measured using the GDP deflator and y_t is the output gap computed by the CBO.

We estimate the SVAR model with four lags and a constant on quarterly data. We use Bayesian methods with standard natural conjugate (Normal-Wishart) priors. Moreover, we specify flat priors for the reduced form parameters in order to remain agnostic about the data generating processes. We also impose sign restrictions on impact using the QR decomposition algorithm proposed by Arias and Waggoner (2018) to identify the structural shocks. This algorithm enables to draw from a conjugate uniform-normal-inverse-Wishart posterior distribution over the orthogonal reduced form parameterization and then to transform the draws into the structural parameterization. This procedure is continued until we have obtained $10,000$ draws that satisfy the imposed sign restrictions.

Sign restrictions are specified in Table 1. Panel A summarizes our identification scheme, which disentangles demand shocks from supply shocks based on the co-movement between inflation and the output gap. In some robustness tests we extend the baseline VAR by adding interest rates. Panel B reports how we identify an additional demand disturbance–a monetary policy shock–in those cases. Finally, in some specifications we include inflation expectations. Then we follow the identification scheme shown in Panel C. The last shock is treated as a residual shock which, in contrast to conventional demand and supply shocks, implies negative co-movement between inflation and inflation expectations.

We use the SVAR model as a filtering device to isolate the variation in historical data due to supply and demand shocks, respectively. Thus, we are essentially interested in historical decompositions. Given that our model is set identified, each of the $10,000$ accepted draws will be associated with a different historical decomposition (and also with a different variance decomposition). In order to summarize this information into *one* decomposition of the unconditional data, we proceed by choosing, for each point in time,

Figure 3: Empirical scatter plots

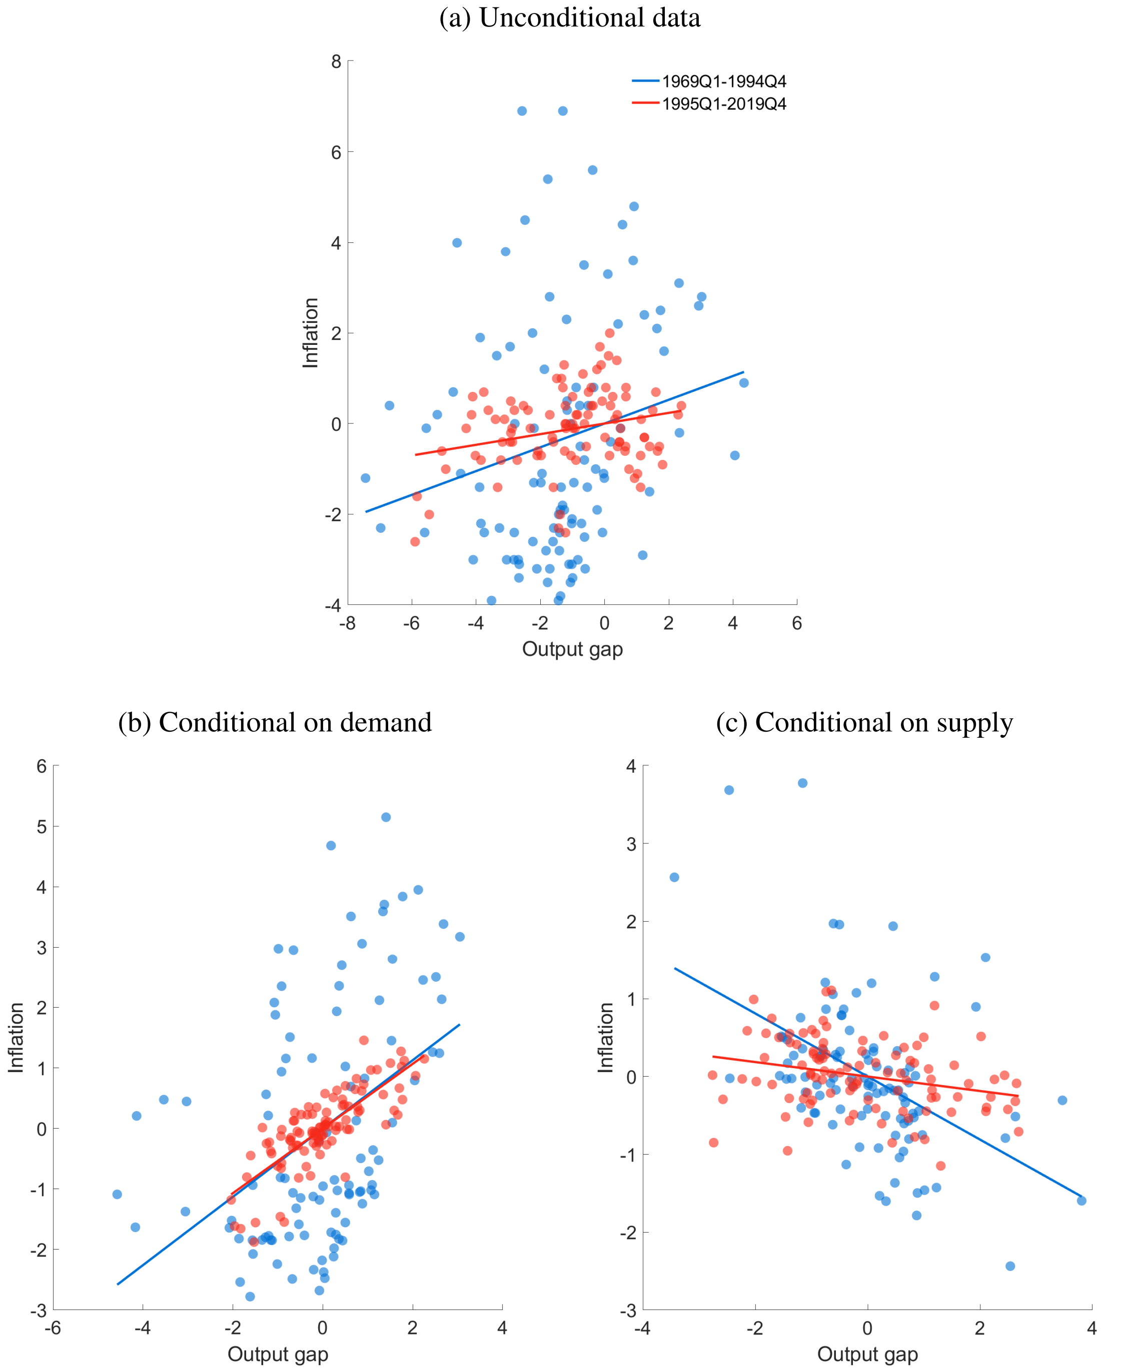

Notes: Unconditional data vs. conditional data obtained from the estimated SVAR model. Corresponding slope estimates are provided in Table 2.

the posterior median decomposition. The final outcome of this exercise is one historical decomposition, with each observation representing the median across the $10,000$ alternative models. In a similar fashion, we obtain one a summary measure for the variance decomposition.

4 RESULTS

This section presents estimates for conditional correlations and conditional variances, i.e. the empirical counterparts of the statistics discussed earlier in the textbook three-equation New Keynesian model.

4.1 FROM UNCONDITIONAL TO CONDITIONAL CORRELATIONS

In order to estimate the unconditional and conditional correlations between inflation and the CBO output gap, we consider the following simple regression equation:

$$\pi_t = c_1 \left(1 + D_t \delta_c\right) + \gamma_1 \left(1 + D_t \delta_\gamma\right) y_t + u_t \tag{8}$$

As in previous sections we denote inflation by π_t, and the output gap by y_t. The dummy variable D_t, which is equal to one in the second sample (and zero otherwise), allows us to separately estimate output gap coefficients across samples. In particular, we denote the slope coefficient in the first sample by γ_1, and the slope coefficient in the second sample by $\gamma_2 = \gamma_1 + \delta_\gamma$. A weakened relationship between output and inflation is captured by a negative value of δ_γ.

In order to assess the competing explanations for a flatter statistical slope, we estimate equation 8 both on unconditional data and on conditional data generated by the SVAR model in the filtering step. Put differently, we contrast the unconditional estimates $\gamma_{u,1}$ and $\gamma_{u,2}$ with the estimates based on conditional data across the two samples. That is, we obtain $\gamma_{d,1}$ and $\gamma_{d,2}$ from data purged for supply shocks, and $\gamma_{s,1}$ and $\gamma_{s,2}$ from the data purged for demand shocks. This leaves us with a set of estimated slope coefficients which, when evaluated jointly, allows us to test the different explanations in a common framework.

As a reference, we start with a discussion of the unconditional data which are presented in the scatter plot in Figure 3a. The horizontal axis measures observations of CBO's output gap, the vertical axis measures inflation in the GDP deflator (plotted in deviation from its mean $\pi_t - c_1 \left(1 + D_t \delta_c\right)$). The scatter plot is augmented with the regression lines which represent the sample specific, unconditional slope coefficients $\gamma_{u,1}$ and $\gamma_{u,2}$.

Panel A in Figure 3 illustrates a significant decline in the slope coefficient estimated on unconditional data. Quantitatively, the estimated slope is $\gamma_{u,1} = 0.26$ during the sample period $1969Q1$-$1994Q4$, but only $\gamma_{u,2} = 0.12$ during the period $1995Q1$-$2019Q4$. Thus, we observe a decline of more than 50% in the unconditional slope in the second sample. As stated earlier, this decline is in line with a large, yet growing literature emphasizing the weakened statistical relationship between inflation and measures of economic activity. At first glance, one may (as several researchers and commentators have done) reach the conclusion that the weakened relationship comes about from a flatter, structural Phillips curve. Note also, that the unconditional variance of both inflation and output is significantly smaller in the latter sample, consistent with the observation that it spans the lion's share of the so-called "Great Moderation".

Panel B in Figure 3 plots the relationship between output and inflation when we condition on the empirical variation attributed solely to identified demand shocks. A number of important observations emerge: first, the estimated slope conditional on demand is

Figure 4: Variance decomposition of CBO's output gap

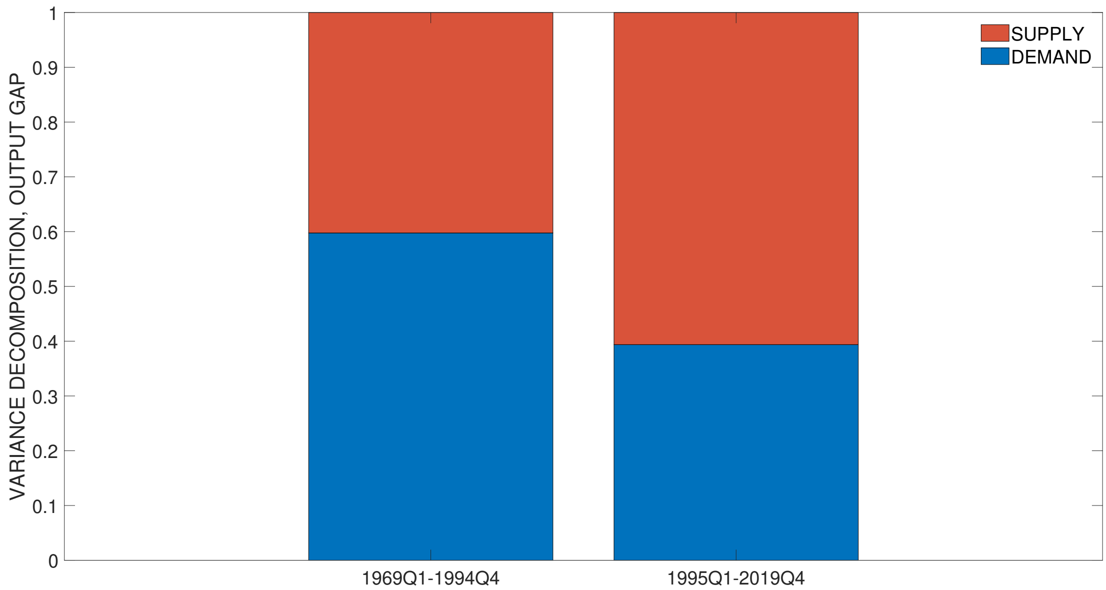

substantially higher than its unconditional counterpart. This is reassuring given the likely downward bias in unconditional slopes, as emphasized earlier. Second, the slope conditional on demand features only a minor decline across samples, from $\gamma_{d,1} = 0.56$ to $\gamma_{d,2} = 0.53$. That is, the empirical relationship between output and inflation, which has weakened substantially in unconditional data, remains relatively constant once we purge out supply side variation in data. In this sense, we do not find statistical evidence of a flattening of the Phillips curve. Rather, it seems to be alive and well. Third, the reduced volatility of both inflation and output carries over when we zoom in on demand driven variation in data.

Finally, Panel C in Figure 3 shows the results when we condition only on identified supply shocks. Naturally, once we consider supply side variation only the relationship between output and inflation turns negative. However, we find large differences in the conditional slopes across samples. In fact, the slope increases from $\gamma_{s,1} = -0.41$ to $\gamma_{s,2} = -0.09$. That is, the demand curve goes from clearly negative to relatively flat. All in all, we conclude that there seems to be a significant flattening of the slope conditional on supply, and this flattening is substantially larger than what we find conditional on demand shocks. To summarize, the evidence reported here points to a flattening of the demand curve rather than the supply curve.

4.2 Conditional variances

We present the variance decomposition in Figure 4 over the two samples. These variance decompositions are based on the scatterplots presented in Figure 3. Supply shocks, in particular, explain 40% of the output gap in the first sample, but 61% in the second sample. An increasing role of supply shocks for the output gap is consistent with the view that inflation targeting has gained focus for the central bank, and at the same time speaks against a flattening of the structural Phillips curve. Ceteris paribus we would actually expect to see a more dominant role of demand shocks in the second sample, had the slope of the Phillips curve declined.

4.3 DISCUSSION

Our results on conditional correlations and conditional variances are consistent with a flattening of the demand curve whose slope is given by $\gamma_s = -\frac{\sigma+\phi_y}{\phi_\pi}$ in the textbook three-equation New Keynesian model. The most natural interpretation for such a flattening is related to a more aggressive response of the monetary policy authority against inflation (i.e. an increase in the coefficient ϕ_π relatively to the coefficient ϕ_y), consistently with the estimates presented in the seminal paper by Clarida, Galí, and Gertler (2000). Incidentally, our results support their explanation of the Great Moderation based on ”good policy” rather than the alternative based on ”good luck”. Using the taxonomy presented in Section 2, the ”good luck” interpretation builds on a lower volatility in the shocks' processes, thus implying unchanged slopes in response to demand and supply shocks. We find that the slope has indeed changed in response to supply shocks, suggesting that the “good luck” story, at best, does not provide the full picture.

While the change is monetary policy behavior seems to us the most natural interpretation for our results, it is important to note that an alternative interpretation based on a higher impact of interest rates on aggregate demand (lower σ) could also explain a flattening in the demand curve. Such a higher impact could reflect higher opportunities to smooth consumption induced by a larger asset market participation. However, two additional implications of our simple New Keynesian model are in contrast with this alternative interpretation. First, a decline in σ should translate also in a decline in κ and our SVAR does not support this implication. Second, a decline in σ should leave the variance decomposition unaffected across samples, in contrast with the evidence presented in Figure 3. This latter observation follows if we substitute $\kappa = \lambda\,(\sigma + \varphi)$ into the variance decomposition of output, given analytically in subsection 2.2. Nonetheless, a potential role for a decline in σ should be investigated more in depth in future research. While it is well known that the degree of asset market participation may have a direct impact on the slope of the demand curve (cf. Galí, López-Salido, and Vallés (2007), Bilbiie (2008) and Bilbiie and Straub (2013), the literature has not, as far as we know, considered changes in asset market participation as a relevant explanation for the inflation puzzle.

Finally, while the mutually exclusive and jointly exhaustive nature of our simple arithmetics is satisfied in our textbook three-equation New Keynesian model, it may not in more complex models with imperfect information. For example, the slope of the Phillips curve may depend endogenously on the stance of monetary policy (cf. Afrouzi and Yang (2021) and L'Huillier, Phelan, and Zame (2022)). The simplest way to introduce dependency of κ on ϕ_π in our framework would be to introduce a working capital channel, so that the interest rate affects marginal costs directly. However, this would lead to a steeper Phillips curve slope when ϕ_π increases, an implication that seems less relevant in this context. Nonetheless, it seems clear from our estimates that the conditional correlation in response to supply shocks has decreased substantially, a point that has is, to the best of our knowledge, new in the literature.

Table 2: Robustness exercises

	$\hat{\gamma}_u$		$\hat{\gamma}_d$		$\hat{\gamma}_s$		$VD(y\|s)$	
	S1	S2	S1	S2	S1	S2	S1	S2
(a)	0.26	0.12	0.56	0.53	-0.41	-0.09	0.40	0.61
(b)	0.23	0.21	0.43	0.79	-0.40	0.12	0.28	0.78
(c)	0.14	0.11	0.39	0.50	-0.21	-0.04	0.47	0.71
(d)	0.39	0.14	0.81	0.63	-0.15	-0.04	0.51	0.70
(e)	0.26	0.12	0.40	0.38	-0.62	-0.07	0.22	0.47
			0.77*	0.92*				
(f)	0.26	0.12	0.39	0.35	-0.63	-0.11	0.22	0.45
			0.76*	0.78*				
(g)	-0.04	0.19	0.25	0.50	-0.36	-0.04	0.43	0.50
(h)	0.26	-0.01	0.57	0.64	-0.41	-0.14	0.41	0.69
(i)	0.26	0.12	0.56	0.39	-0.80	0.06	0.19	0.45
(j)	0.25	0.12	0.50	0.66	-0.46	-0.05	0.19	0.54
(k)	0.44	0.19	0.81	1.02	-0.88	0.04	0.19	0.60

*Conditional on identified monetary policy shocks.

5 ROBUSTNESS

In this section we evaluate the robustness of our main results with respect to alternative measures of inflation and real economic activity included in the bivariate model. We consider also different sample periods and alternative specifications of the SVAR including measures of the interest rate (and thus identifying a second demand shock in the form of a monetary policy shock) or measures of inflation expectations. Results are presented in Table 2. We have considered the following ten specifications in addition to the baseline model presented in the previous section:

(a) Baseline as a reference

(b) Cyclically sensitive inflation Stock and Watson (2019) as a measure of inflation

(c) Unemployment rate as a measure of real economic activity

(d) Unemployment gap from u^* (trend) as a measure of real economic activity

(e) 3-variable VAR with the Federal Funds Rate

(f) 3-variable VAR with the shadow rate as computed by Wu and Xia (2016)

(g) Sample split in 1998Q4 as in Jorgensen and Lansing (2019)

(h) Second sample ends after 2008Q4

(i) Baseline augmented with SPF expectations

(j) Baseline augmented with Michigan expectations

(k) CPI inflation and Michigan expectations

5.1 ALTERNATIVE VARIABLES

In specification (b) we estimate the bivariate SVAR using the "cyclically sensitive" (CSI) measure of inflation computed by Stock and Watson (2019). CSI inflation places low weights on tradable goods and on the least well-measured sectors and high weights on nontradeable goods and services and on relatively well-measured sectors. In keeping with Stock and Watson (2019), we find that CSI maintains a relatively strong correlation with measures of real economic activity in the second sample. In such a context of stability, one could perhaps expect to find stable slopes also in the conditional relationships. In contrast, we confirm the flattening in response to supply shocks and a larger role of supply shocks in the second part of the sample, as in the baseline model. The stability in the unconditional slope and the flattening in response to supply shocks can coexist only in presence of a steepening in response to demand shocks. In fact, this is what we find: the estimate for $\hat{\gamma}_d$ increases from 0.43 to 0.79, thus pointing to a steepening of the Phillips curve.

In specifications (c) and (d) we use the unemployment rate and the CBO measure of the unemployment gap respectively as indicators of real economic activity. Inflation is negatively correlated with unemployment unconditionally and unemployment enters with a negative sign in empirical specifications of the Phillips curve. Therefore, in order to facilitate the comparison with our baseline model, we switch the sign of all slopes' coefficients reported in Table 2. In both specifications, we find a flattening in the unconditional data, a clear flattening in response to supply shocks and a larger role for supply shocks in the variance decomposition in the second sample. The specification using the unemployment rate finds a steepening in response to demand shocks while some flattening (from 0.81 to 0.63) is found when using the unemployment gap. All in all, these specifications are still consistent with the policy story being the main driver of the inflation puzzle.

5.2 EXTENSION WITH MONETARY POLICY SHOCKS

In specifications (e) and (f) we include the interest rate and the shadow interest rate as computed by Wu and Xia (2016) respectively into the VAR and also identify a monetary policy shock driving a negative co-movement between the interest rate and the output gap. The identification assumptions used in this trivariate SVAR are detailed in panel (b) on Table 1. One major benefit is that the monetary policy shock acts as a second demand shifter. In a way, we are decomposing the baseline demand shock into two components. Interestingly, there is no sign of a decline in $\hat{\gamma}_d$ conditional on monetary policy shocks (if anything, we find some steepening). Neither is there a decline when we look at the purified, non-monetary demand shock. Instead, we have a clear flattening of the supply slope, from -0.62 to -0.07 in specfication (e) and from -0.63 to -0.11 in specification (f). Finally, the role of supply shocks in the variance decomposition of output doubles in both specifications. We would likely argue that the model with shadow rates is preferable, given that the nominal interest rate was stuck at the lower bound period for several years in the second sample.

5.3 ALTERNATIVE SAMPLE PERIOD

In specification (g) we consider a sample splitting in 1998:Q4. We choose this date because, as shown by Jorgensen and Lansing (2019), it delivers a negative a correlation between the level of inflation and the output gap in the first sample and a positive correlation in the second sample. Uisng our data, the correlation changes sign from -0.04 to 0.19. According to theory, a more aggressive response of monetary policy against inflation can reduce the magnitude of the unconditional correlation between inflation and the output gap but cannot explain on its own the change in sign from negative to positive. This specification of the SVAR confirms a strong flattening in the slope of the demand curve (from -0.36 to -0,04), as in our baseline model. At the same time, we estimate a steepening of the supply curve from 0,25 to 0,5. Therefore, the SVAR combines a steepening in supply with a flattening in demand to explain the shift in sign in the unconditional correlation. Notably, supply shocks explain a larger share of output fluctuations in the second sample, in keeping with our baseline model.

In specification (h) we modify the length of the second sample by stopping the estimation when the zero lower bound starts binding at the end of 2008 (thus estimating the model over the period 1995:Q1-2008:Q4). This exercise is important because New Keynesian theory tells us that the propagation of shocks can change substantially in presence of a binding zero lower bound (although unconventional monetary policies seem to limit the changes in propagation in practice, according to the empirical evidence provided by Debortoli, Galí, and Gambetti (2020)). Therefore, we want to check that our results are not driven by a period that has been very peculiar for macroeconomic policy. Notably, our results are confirmed (if not reinforced) in specification (h): we find a tiny steepening in response to demand shocks, a clear flattening in response to supply shocks and an increase role for supply shocks in the output gap variance decomposition.

5.4 Including inflation expectations

In specifications (j), (k), and (l) we extend the SVAR to include data on inflation expectations. This exercise can be seen as equivalent to Coibion and Gorodnichenko (2015) with the crucial difference that estimates are conducted also on conditional data as filtered by our SVAR model. We include a third shock in the system moving inflation and inflation expectations in opposite directions, as detailed by the identification assumptions listed in panel (c) on Table 1, to match the number of observables with the number of identified shocks. This third shock is left without an economic interpretation and plays a minor role in the model since inflation and inflation expectations are positively correlated in the data. In addition, inflation expectations do not respond much to shocks, in particular in low-inflation environments (cf. Coibion, Gorodnichenko, Kumar, and Pedemonte (2020) and Coibion, Gorodnichenko, Knotek II, and Schoenle (2020)).

In our first experiment (specification (i)), we use inflation expectations data from the Survey of Professional Forecasters (SPF). We confirm a strong decline in the conditional correlation in response to supply shocks (from -0.8 to 0.06, thus much larger than in the baseline model) and a more important role for supply shocks in driving the output gap in the variance decomposition. However, we also estimate a reduction in the conditional correlation in response to demand shocks from 0.56 to 0.39. Since this result is consistent with a flattening of the Phillips curve, the reader could think that our main result is (at least in part) weakened by the inclusion of data on inflation expectations in the SVAR. However, such a conclusion is premature. In fact, the expectations relevant for pricing decisions in the context of the Phillips curve are firm inflation expectations as discussed in detail in Coibion and Gorodnichenko (2015). Unfortunately, there is no quantitative measure of firm inflation expectations available in the United States for a sufficiently long sample. Notably, however, Coibion and Gorodnichenko (2015) argue that household inflation expectations, as measured by the Michigan Survey of Consumers, are a better proxy for firm expectations than SPF expectations and provide supporting empirical evidence based on survey data from New Zealand. Therefore, in specification (j) we include data on inflation expectations from the Michigan survey in our baseline model. All results are now stronger than in our baseline specification. In addition, we now find an increase in $\hat{\gamma}_d$ from 0.5 to 0.66, thus pointing to a steepening of the Phillips curve. One may criticize this experiment because expectations in the Michigan Survey are about CPI inflation and not about the GDP deflator. To address this concern, we use data on CPI inflation (together with the output gap and the Michigan Survey measure of inflation expectations) in specification (k). The results are again stronger than in the baseline model. The estimate for $\hat{\gamma}_s$ (from -0.88 to 0.04) indicates a clear flattening of the demand curve while $\hat{\gamma}_s$ signals a steepening of the Phillips curve. Consequently, supply shocks become once again the main drivers of the output gap. Conditional on the Coibion and Gorodnichenko (2015) arguments and evidence in favor of the use of the Michigan Survey as a better proxy for firms expectations, we conclude that our results are reinforced when including inflation expectations in the SVAR. One potential explanation is related to the fact that households adjust their inflation forecasts more strongly in response to oil price changes than professional forecasters, thus favoring a more accurate identification of supply shocks.

5.5 Inspecting the posterior distribution

Following common practice, all of the results so far have been based on data conditional on the pointwise posterior median of the SVAR model. As a final robustness exercise, we instead evaluate the conditional regression slopes across a wide range of the posterior distribution. Recall that the Bayesian approach taken in this paper provides us, for each of the two samples under consideration, $10,000$ different posterior models. All of these posterior models are consistent with the sign restrictions presented earlier, yet all of them give rise to a unique decomposition of the raw unconditional data. To further investigate this, we use the $10,000$ sample-specific, conditional datasets to compute $10,000$ sample-specific estimates of $\hat{\gamma}_d$ and $\hat{\gamma}_s$. This is done both for the baseline specification and for all of the robustness specifications reported in Table 2. The results of this exercise are summarized in Table 3, where we report the posterior mean together with 68% bands for the credible set (in brackets). Unconditional slopes are also provided in order to facilitate comparison of the results.

A couple of observations stand out: first, regarding $\hat{\gamma}_d$, we confirm the general picture established earlier. Slope coefficients conditional on demand shocks do not tend to decrease, at least not to a major extent. The 68% credible set for $\hat{\gamma}_d$ in sample 2 spans its posterior mean in sample 1 in all specifications except (g) and (h). Specification (g) uses the same sample split as in Jorgensen and Lansing (2019), implying in fact a steepening of the slope–both unconditionally and conditional on demand. Specification (h), instead, is where we discard observations after $2008Q4$. This is the only significant exception pointing to a flattening of the Phillips curve. Second, regarding $\hat{\gamma}_s$, we find a major flattening

Table 3: Robustness to posterior credibility sets

	$\hat{\gamma}_u$		$\hat{\gamma}_d$		$\hat{\gamma}_s$	
	S1	S2	S1	S2	S1	S2
(a)	0.26	0.12	0.61 $[0.29, 0.89]$	0.51 $[0.15, 0.83]$	-0.59 $[-1.32, 0.13]$	-0.14 $[-0.34, 0.09]$
(b)	0.23	0.21	0.41 $[0.20, 0.62]$	0.52 $[0.24, 0.81]$	-0.54 $[-1.26, 0.08]$	0.12 $[0.04, 0.23]$
(c)	0.14	0.11	0.18 $[0.88, -0.51]$	0.43 $[0.69, 0.13]$	-0.35 $[0.42, -1.12]$	-0.07 $[0.12, -0.21]$
(d)	0.39	0.14	0.59 $[1.28, -0.20]$	0.54 $[0.87, 0.18]$	-0.31 $[0.54, -1.11]$	-0.09 $[0.15, -0.26]$
(e)	0.26	0.12	0.34 $[0.12, 0.56]$	0.34 $[0.12, 0.52]$	-0.42 $[-1.01, 0.14]$	-0.10 $[-0.27, 0.09]$
(f)	0.26	0.12	0.33 $[0.11, 0.56]$	0.34 $[0.11, 0.54]$	-0.41 $[-0.99, 0.14]$	-0.12 $[-0.32, 0.09]$
(g)	-0.04	0.19	0.12 $[-0.16, 0.40]$	0.54 $[0.24, 0.76]$	-0.51 $[-1.08, 0.06]$	-0.07 $[-0.29, 0.17]$
(h)	0.26	-0.01	0.62 $[0.29, 0.91]$	0.33 $[0.02, 0.62]$	-0.59 $[-1.29, 0.12]$	-0.11 $[-0.29, 0.10]$
(i)	0.26	0.12	0.49 $[0.21, 0.77]$	0.40 $[0.11, 0.66]$	-0.56 $[-1.17, 0.01]$	0.02 $[-0.14, 0.19]$
(j)	0.25	0.12	0.43 $[0.24, 0.61]$	0.49 $[0.15, 0.79]$	-0.14 $[-0.60, 0.34]$	-0.07 $[-0.24, 0.10]$
(k)	0.44	0.19	0.90 $[0.41, 1.35]$	0.73 $[0.19, 1.25]$	-0.28 $[-0.86, 0.27]$	0.03 $[-0.13, 0.22]$

(i.e. less negative) of the slope conditional on supply in all specifications when considering the posterior mean across the $10,000$ models. Moreover, the 68% credible set for $\hat{\gamma}_s$ in sample 2 does not span the posterior mean in sample 1 in any of the specifications except (j), where we add inflation expectations from the Michigan survey to the analysis. Finally, it seems that the biggest change in the posterior distribution of $\hat{\gamma}_s$ across samples is located in the left tail of the distribution. The 16% lower bound in the baseline model, for example, increases from -1.32 to -0.34, substantially more than the 84% upper bound. In any case, we conclude that the results presented in Table 3 supports the picture drawn earlier: evidence of a flatter supply curve is weak in our data, while there seems to be a major flattening of the demand curve.

6 Conclusion

In this paper we have reconsidered the puzzling stability of inflation in spite or large fluctuations in real economic activity over the last couple of decades. Using a combination of New Keynesian theory and estimated SVAR models, we argue that controlling for the effects of all supply shocks (and not only for cost-push shocks) is of paramount importance to evaluate alternative explanations of the inflation puzzle. While we reconfirm that the unconditional correlation between output gap and inflation has declined, we find that the correlation has been substantially stable when conditioning on demand shocks. In contrast, We find substantial support for an alternative explanation for the inflation puzzle based on a more aggressive response of monetary policy against inflation in the second part of the sample.

It is important to discuss our results in connection with a few crucial papers in the literature what have questioned the narrative of the Phillips curve flattening. The paper by McLeay and Tenreyro (2020) is a key reference to highlight how the Phillips curve can be alive and well even if inflation is unrelated to the output gap in the data. McLeay and Tenreyro (2020) show that under optimal monetary policy the residual variation in output and inflation is driven only by cost-push shocks: in that scenario, a negative correlation between inflation and the output gap should emerge, blurring the identification of the (positively sloped) Phillips curve. In contrast, our identification scheme assumes that monetary policy is represented by a Taylor rule that is unable to mimic perfectly optimal monetary policy. Our choice is motivated by the fact that, if optimal monetary policy (in the form of a targeting rule) was in place, we should observe a negative unconditional correlation between inflation and the output gap in the data: Panel A in Figure 3 shows that this is not the case. Barnichon and Mesters (2020) stress the importance of using demand shocks (monetary policy shocks in their case) as instruments to trace the slope of the Phillips curve. We follow their prescription using a more general demand shock, although we consider also monetary policy shocks in isolation in our sensitivity analysis. Importantly, we apply the Barnichon and Mesters (2020) recommendation also in using supply shocks to trace the slope of the demand curve. The focus on the *joint* identification of *both* shocks allows us to exploit fully our simple arithmetics and derive our main result. Hazell et al. (2022) and Jorgensen and Lansing (2019) find that the anchoring of inflation expectations is crucial to explain the inflation puzzle using regional and aggregate data respectively. Both papers find that the estimated slope coefficient is stable over time although its magnitude depends on whether the estimation is performed on regional or aggregate data. In our framework, the anchoring of inflation expectations is a by-product of a more aggressive monetary policy response against inflation and our focus on supply shocks is crucial to provide additional validation to the result obtained in these previous elegant papers. Finally, Coibion and Gorodnichenko (2015) document the importance of oil shocks between 2009 and 2011 to explain the rise in consumers' inflation expectations and the absence of disinflation during the Great Recession. Using a very different framework, we also stress the importance of accounting for supply shocks and their joint role in explaining the Great Recession and macroeconomic dynamics in general.

Our sample stops in 2019:Q4 just before the COVID recession. Since then, inflation has come back and the current debate is once again on how policy should be set in order to lower inflation, and not to increase inflation as in the previous decade. Therefore, one could be tempted to say that our analysis is outdated. We believe that this is not the case for two reasons. First, the COVID recession has re-confirmed the importance of supply-side factors for output and inflation dynamics. Disruptions in global supply chains, shocks to energy prices and labor supply factors (being the Great Resignation one example) are routinely mentioned as important drivers of price and wage inflation. Therefore, we would argue that the recent experience strengthen our argument that is important to control for all supply shocks in empirical analysis of the Phillips curve. Second, our simple arithmetics is potentially useful also in the current phase. Inflation is again well above the Fed's target and, if this tendency continues, we could have entered in a new regime, more similar to our first sample. A higher slope of the Phillips curve or a less aggressive monetary policy response of monetary policy are compatible with the resurgence of inflation (in addition of course, with variations in the volatilities of the underlying disturbances). While it is still too early to draw inference from data on what mechanism is at play, we believe that our simple arithmetics can be a useful input to organize the discussion.

References

Afrouzi, H. and C. Yang (2021). Dynamic rational inattention and the phillips curve. *Available at SSRN 3770462*.

Arias, Jonas E, R. R. J. F. and D. F. Waggoner (2018). Inference based on structural vector autoregressions identified with sign and zero restrictions: theory and applications. *Econometrica 82*(2), 685–720.

Ascari, G., P. Bonomolo, and Q. Haque (2022). The long-run phillips curve is... a curve.

Ascari, G. and L. Fosso (2021). The inflation rate disconnect puzzle: On the international component of trend inflation and the flattening of the phillips curve.

Ball, L. and S. Mazumder (2011). Inflation dynamics and the great recession. *Brookings Papers on Economic Activity*, 337–406.

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.

Benati, L. (2015). The long-run phillips curve: A structural var investigation. *Journal of Monetary Economics 76*, 15–28.

Beraja, M., E. Hurst, and J. Ospina (2019). The aggregate implications of regional business cycles. *Econometrica 87*(6), 1789–1833.

Bilbiie, F. O. (2008). Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of economic theory 140*(1), 162–196.

Bilbiie, F. O. and M. J. Melitz (2020). Aggregate-demand amplification of supply disruptions: The entry-exit multiplier. Technical report, National Bureau of Economic Research.

Bilbiie, F. O. and R. Straub (2013). Asset market participation, monetary policy rules, and the great inflation. *Review of Economics and Statistics 95*(2), 377–392.

Blanchard, O. (2016). The phillips curve: Back to the'60s? *American Economic Review 106*(5), 31–34.

Cecchetti, S. G. and R. W. Rich (2001). Structural estimates of the us sacrifice ratio. *Journal of Business & Economic Statistics 19*(4), 416–427.

Clarida, R., J. Galí, and M. Gertler (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *The Quarterly Journal of Economics 115*(1), 147–180.

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*(1), 197–232.

Coibion, O., Y. Gorodnichenko, E. S. Knotek II, and R. Schoenle (2020). Average inflation targeting and household expectations. Technical report, National Bureau of Economic Research.

Coibion, O., Y. Gorodnichenko, S. Kumar, and M. Pedemonte (2020). Inflation expectations as a policy tool? *Journal of International Economics 124*, 103297.

Coibion, O., Y. Gorodnichenko, and M. Ulate (2018). The Cyclical Sensitivity in Estimates of Potential Output. *Brookings Papers on Economic Activity 49*(2 (Fall)), 343–

441.

Debortoli, D., J. Galí, and L. Gambetti (2020). On the empirical (ir) relevance of the zero lower bound constraint. *NBER Macroeconomics Annual 34*(1), 141–170.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020). What's Up with the Phillips Curve? *Brookings Papers on Economic Activity*.

Fitzgerald, T. J., C. Jones, M. Kulish, J. P. Nicolini, et al. (2020). Is there a stable relationship between unemployment and future inflation?

Forbes, K. (2019). Inflation dynamics: Dead, dormant, or determined abroad? Technical report, National Bureau of Economic Research.

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

Galí, J. and L. Gambetti (2009). On the sources of the great moderation. *American Economic Journal: Macroeconomics 1*(1), 26–57.

Galí, J. and L. Gambetti (2019). Has the us wage phillips curve flattened? a semi-structural exploration. Technical report, National Bureau of Economic Research.

Galí, J. and M. Gertler (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics 44*(2), 195–222.

Galí, J., J. D. López-Salido, and J. Vallés (2007). Understanding the effects of government spending on consumption. *Journal of the european economic association 5*(1), 227–270.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek (2017). Inflation dynamics during the financial crisis. *American Economic Review 107*(3), 785–823.

Gordon, R. J. (2013). The phillips curve is alive and well: Inflation and the nairu during the slow recovery. Technical report, National Bureau of Economic Research.

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2022). The slope of the phillips curve: evidence from us states. *The Quarterly Journal of Economics 37*(3), 1299–1344.

Heise, S., F. Karahan, A. Şahin, et al. (2020). The missing inflation puzzle: The role of the wage-price pass-through.

Hobijn, B. (2020). The supply-side origins of us inflation. *Serie Banca Central, análisis y políticas económicas, no. 27*.

Imbs, J., E. Jondeau, and F. Pelgrin (2011). Sectoral phillips curves and the aggregate phillips curve. *Journal of Monetary Economics 58*(4), 328–344.

Jorgensen, P. L. and K. J. Lansing (2019). Anchored inflation expectations and the flatter phillips curve. Federal Reserve Bank of San Francisco.

King, R. G. and M. W. Watson (1994). The post-war us phillips curve: a revisionist econometric history. In *Carnegie-Rochester Conference Series on Public Policy*, Volume 41, pp. 157–219. Citeseer.

Kleibergen, F. and S. Mavroeidis (2009). Weak instrument robust tests in gmm and the new keynesian phillips curve. *Journal of Business & Economic Statistics 27*(3), 293–311.

L'Huillier, J. P., G. Phelan, and W. R. Zame (2022). Flattening of the phillips curve and the mandate of the central bank.

Mavroeidis, S., M. Plagborg-Møller, and J. H. Stock (2014). Empirical evidence on inflation expectations in the new keynesian phillips curve. *Journal of Economic Literature 52*(1), 124–88.

McLeay, M. and S. Tenreyro (2020). Optimal inflation and the identification of the phillips curve. *NBER Macroeconomics Annual 34*(1), 199–255.

Obstfeld, M. (2020). Global dimensions of us monetary policy. *International Journal of Central Banking 16*(1), 73.

Rubbo, E. (2020). Networks, phillips curves and monetary policy.

Sbordone, A. M. (2002). Prices and unit labor costs: a new test of price stickiness. *Journal of Monetary Economics 49*(2), 265–292.

Stock, J. H. and M. W. Watson (2019). Slack and cyclically sensitive inflation. Technical report, National Bureau of Economic Research.

Woodford, M. (2003). *Interest and prices: Foundations of a theory of monetary policy*. Princeton University Press.

Wu, J. C. and F. D. Xia (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking 48*(2-3), 253–291.

APPENDIX: THE MODEL WITH AUTOREGRESSIVE SHOCKS

Suppose that instead of being i.i.d. as in the main text, the structural shocks follow separate AR(1) processes:

$$\begin{aligned} d_t &= \rho_d d_{t-1} + \varepsilon_{d,t} \quad\quad \varepsilon_{d,t} \sim N\left(0, \sigma^2_{\varepsilon,d}\right) \\ s_t &= \rho_s s_{t-1} + \varepsilon_{s,t} \quad\quad \varepsilon_{s,t} \sim N\left(0, \sigma^2_{\varepsilon,s}\right) \end{aligned}$$

Closed form solutions for output and inflation follow:

$$\begin{aligned} y_t &= \frac{\sigma\left(1-\beta\rho_d\right)}{\Psi_d} d_t - \frac{\left(\phi_\pi - \rho_s\right)}{\Psi_s} s_t \\ \pi_t &= \frac{\sigma\kappa}{\Psi_d} d_t + \frac{\left[\sigma\left(1-\rho_s\right)+\phi_y\right]}{\Psi_s} s_t \end{aligned}$$

We have defined two auxiliary functions:

$$\begin{aligned} \Psi_d &= \left[\sigma\left(1-\rho_d\right)+\phi_y\right]\left(1-\beta\rho_d\right) + \left(\phi_\pi - \rho_d\right)\kappa > 0 \\ \Psi_s &= \left[\sigma\left(1-\rho_s\right)+\phi_y\right]\left(1-\beta\rho_s\right) + \left(\phi_\pi - \rho_s\right)\kappa > 0 \end{aligned}$$

OLS estimators under different assumptions follow readily:

(a) Unconditional data:

$$\begin{aligned} \kappa^{OLS} &= \frac{cov\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}, y_t\right)}{var\left(y_t\right)} \\ &= \frac{\left(\frac{\sigma(1-\beta\rho_d)}{\Psi_d}\right)^2 \kappa - \left(\frac{1}{\Psi_z}\right)^2 \left(\phi_\pi - \rho_z\right)\left[\sigma\left(1-\rho_z\right)+\phi_y\right]\left(1-\beta\rho_z\right)\frac{\sigma_z^2}{\sigma_d^2}}{\left(\frac{\sigma(1-\beta\rho_d)}{\Psi_d}\right)^2 + \left(\frac{\phi_\pi-\rho_z}{\Psi_z}\right)^2 \frac{\sigma_z^2}{\sigma_d^2}} \leq \kappa \end{aligned}$$

where $\sigma_d^2 = \frac{\sigma^2_{\varepsilon,d}}{1-\rho_d^2}$ and $\sigma_z^2 = \frac{\sigma^2_{\varepsilon,z}}{1-\rho_z^2}$.

(b) Purged of supply shocks, but ignoring expectations:

$$\kappa^{OLS} = \frac{cov\left(\pi_t, y_t\right)}{var\left(y_t\right)} = \frac{\kappa}{1-\beta\rho_d} \geq \kappa$$

(c) Purged of supply shocks and accounting for expectations:

$$\kappa^{OLS} = \frac{cov\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}, y_t\right)}{var\left(y_t\right)} = \kappa$$

(d) Purged of demand shocks:

$$\kappa^{OLS} = -\frac{\left[\sigma\left(1-\rho_z\right)+\phi_y\right]\left(1-\beta\rho_z\right)}{\phi_\pi - \rho_z} < 0$$