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

DRAGO BERGHOLT†, FRANCESCO FURLANETTO‡ AND ETIENNE VACCARO-GRANGE§

April 2024

Abstract: An apparent disconnect has taken place between inflation and economic activity in the US over the last 25 years, with price inflation remaining remarkably stable in spite of large fluctuations in the output gap and other measures of economic slack. This observation has led some to believe that the Phillips curve–a summary measure of aggregate supply–has flattened. We argue that this view may be premature and put forward a few, simple arithmetics which give rise to testable implications for demand and supply curve slopes. Equipped with New Keynesian theory and estimated SVAR models, we decompose the unconditional variation in US macro data into the components driven by demand and supply disturbances, and confront the inflation disconnect with our simple arithmetics. This exercise reveals a relatively stable supply curve slope once shocks to supply have been properly accounted for. The demand curve, instead, has flattened substantially in recent decades. Our results are at odds with a decline in the Phillips curve slope, but fully consistent with a shift towards a more firm monetary policy commitment to inflation stability.

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

---

*First draft: November 2019. This working paper should not be reported as representing the views of Norges Bank or the IMF. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank, the IMF, its Executive Board, or IMF management. We would like to thank for useful comments our discussants Benjamin Wong, Gregor von Schweinitz and Yoosoon Chang. We thank also Knut Are Aastveit, Guido Ascari, Larry Ball, Regis Barnichon, Martin Beraja, Fabio Canova, Jeff Campbell, Efrem Castelnuovo, Davide Debortoli, Thomas Drechsel, Michael Ehrmann, Luca Fosso, Jordi Galí, Luca Gambetti, Domenico Giannone, Refet Gürkaynak, Martin Holm, Rupal Kamdar, Nicolo Maffei Faccioli, Kevin Lansing, Cristina Manea, Elmar Mertens, Lorenzo Mori, Emi Nakamura, Giorgio Primiceri, Federico Ravenna, Robert Rich, Ørjan Robstad, Øistein Røisland, Leif Anders Thorsrud, Jon Steinsson, Anders Vredin, Carl Walsh, Chris Wolf as well as seminar participants at Norges Bank, Federal Reserve Bank of Cleveland, OMG Oslo seminar, Sveriges Riksbank, Central Bank of Ireland, Federal Reserve Bank of San Francisco, University of Le Mans, Sant'Anna School of Economics, 2019 CAMA-CAMP Workshop in Sydney, CEBRA virtual conference 2020, New Developments in Business Cycle Analysis 2021 conference hosted virtually by Deutsche Bundesbank, virtual SNDE conference 2022, EABCN conference on Challenges in Empirical Macroeconomics since 2020 in Eltville, Barcelona Summer Forum 2022, CEF conference 2022 in Dallas, EEA 2022 conference in Milano, SEM 2022 conference in Calgary, 2022 conference on The Return of Inflation at Norges Bank, Forskermøte 2022 in Stavanger, 23rd IWH-CIREQ-GW Macroeconometric Workshop in Halle, 2022 KIEP-CEPR conference on The Return of Inflation in Seoul, 2023 Padova Macro Talks, IAAE 2023 conference in Oslo and 2023 Sailing the Macro conference in Siracusa. 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. It also benefited from a Post-Doctoral Fellowship at New York University Abu Dhabi for one of its authors.

†Norges Bank. Email: drago.bergholt@norges-bank.no
‡Norges Bank. Email: francesco.furlanetto@norges-bank.no
§International Monetary Fund. Email: evaccaro-grange@imf.org

1 Introduction

Price inflation has remained remarkably stable in the United States since the mid 1990s. While a new regime may arise in the aftermath of COVID, the last few decades have delivered a major decline in inflation volatility, without a comparable decline in economic activity. This is shown in Figure 1. During the pre-Great Recession boom, inflation (measured with the GDP deflator) was stubbornly stable, barely above 2 percent. During the Great Recession, in face of the largest decline in real economic activity since the Great Depression, inflation declined by only around two percent. In the aftermath of the Great Recession, real economic activity recovered (albeit slowly), unemployment reached a 50-year low slightly below 4 percent but inflation remained consistently below 2 percent.

Why has inflation become so stable and why has it disconnected from real economic activity? These are the questions we aim to address in this paper. At least three different explanations are proposed in the literature: the first (and perhaps most widely accepted narrative) points to a decline in the slope of the Phillips curve, a structural equation describing how pressure in the economy translates into inflation. Examples include Ball and Mazumder (2011), Blanchard (2016), Stock and Watson (2020), Del Negro, Lenza, Primiceri, and Tambalotti (2020) and Inoue, Rossi, and Wang (2022) among others. Potential reasons for such a flattening include a more prominent role for global factors (such as increased import competition), rising market concentration, changes in the network structure of the US production sector (cf. Forbes (2019), Obstfeld (2020), Heise, Karahan, and Şahin (2022), Rubbo (2023) and Ascari and Fosso (2024) among others), or differences in financial conditions across firms (cf. Gilchrist, Schoenle, Sim, and Zakrajšek (2017)). A second explanation highlights the role of monetary policy: the Federal Reserve may have become more aggressive over time with respect to achieving its inflation stability mandate (McLeay and Tenreyro, 2020). Proponents of this view acknowledge that the Phillips curve could be alive and stable, and instead argue that stricter inflation targeting has led to smaller footprints associated with demand factors. This may have resulted in declining inflation volatility. The third explanation concerns a potential change in the composition of shocks over time (Galí and Gambetti (2009), Gordon (2013) and Hobijn (2020)). Supply shocks in particular may have become relatively more prominent or more concentrated in specific periods, as was the case for oil shocks in the aftermath of the Great Recession (Coibion and Gorodnichenko, 2015). If this is the case, then inflation may co-move less with real economic activity even if both the Phillips curve slope and the stance of monetary policy remain stable. We refer to the three explanations summarized above as the *slope hypothesis*, the *policy hypothesis*, and the *shock hypothesis*.

In this paper, we apply some simple arithmetics which allow us to evaluate the three main explanations for the observed disconnect between inflation and economic activity. Our arithmetics are rather uncontroversial and consistent, for example, with the canonical New Keynesian model described in Galí (2008). Since that model features nominal rigidities, it can be summarized in the output gap-inflation space by an upward-sloping supply curve (the Phillips curve) and a downward-sloping demand curve (the investment-savings curve). The supply curve is upward-sloping because demand creates resource scarcity and motivates firms to raise prices. The demand curve is downward-sloping because the central bank responds to inflation by hiking its policy rate, causing a decline in aggregate demand for goods and services. Now, consider an econometrician who lives in the

Figure 1: The evolution of inflation (GDP deflator) and the output gap

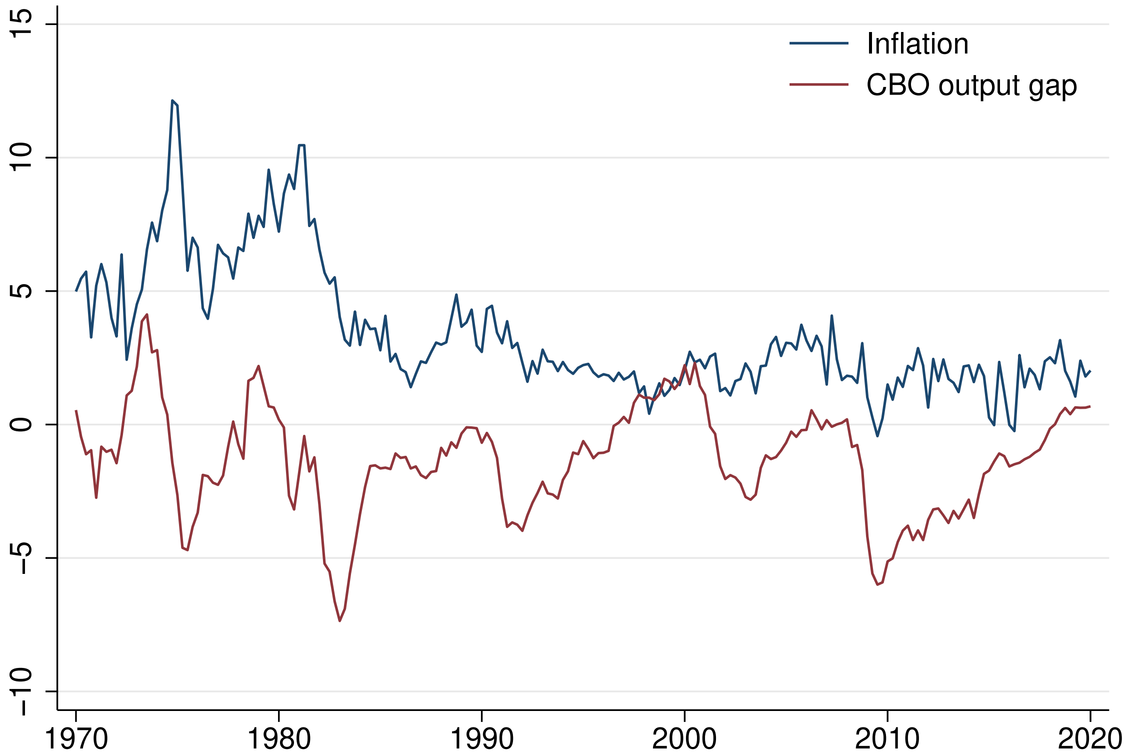

world just described and would like to regress inflation on the output gap in order to learn about an observed inflation disconnect. We inspect the properties of structural supply and demand, and formally derive the following testable implications that the econometrician may confront with data: first, the slope hypothesis–by proposing a flattening of aggregate supply while the curvature of demand remains fixed–is the only one that implies *a less positive regression slope* when data are *purged for supply-side shocks*. The policy hypothesis in contrast–by proposing a flattening of aggregate demand while the curvature of supply remains fixed–is unique in that it implies *a less negative regression slope* when data are *purged for demand-side shocks*. Intuitively, stricter central bank commitment to inflation stability translates into smaller inflation movements and a flatter demand curve. Finally, the shock hypothesis is the only narrative that implies *stable regression slopes* when we condition either on demand shocks or on supply shocks. This hypothesis advocates changes in the cocktail of realized shocks rather than in the responsiveness of supply and demand to a given disturbance. Combined, our simple arithmetics give rise to a set of testable hypotheses which enable us to disentangle the competing explanations. All that is needed is a way to purge demand from supply in the observable, unconditional data.

In order to confront the arithmetics just described with actual data, we fit a structural vector autoregression (SVAR) model to postwar US macro data on the GDP deflator and the output gap as computed by the Congressional Budget Office (CBO). Importantly, our observed output gap is measured in deviation from some notion of a slow-moving trend, not in deviation from a hypothetical equilibrium with flexible prices. While this is completely standard for output gap estimates provided by statistical agencies (cf. Coibion, Gorodnichenko, and Ulate (2018)), it implies that all supply-side shocks (not only cost-push or markup shocks) shift the supply curve and must be accounted for. Thus, our empirical model is relatively general and imposes only a *minimal set of sign restrictions* to disentangle demand shocks from supply shocks. With the two shocks identified, one can estimate the slopes of aggregate demand and supply (each corresponds to a regression line fitting a cloud of data points generated by one of the identified shocks), and check to what extent these two conditional slopes have changed over time. This is exactly what we do. The main result from this exercise is that *postwar US data call for a stable supply curve, combined with a substantially flatter demand curve* since the mid 1990s. These findings are clearly at odds with a decline in the Phillips curve slope, but fully consistent with stricter inflation targeting.

The empirical slope results just presented are complemented with an inspection of conditional variances: we derive analytical expressions for the shock decomposition and show that, according to theory, demand shocks should become increasingly important for the output gap if the slope hypothesis is true, while the policy hypothesis and the shock hypothesis both imply a relative rise in the importance of supply shocks. The policy hypothesis in particular leads to smaller shifts in the demand curve even if the volatility of fundamental demand shocks has remained unchanged. Thus, a variance decomposition of the output gap in data may serve as a useful cross-check for our main empirical results. Interestingly, when we compute the variance decomposition of the CBO output gap, we find a rise in the relative importance of supply-side shocks after the mid 1990s, lending additional support to the policy hypothesis.

Our baseline empirical results are robust to a large number of additional sensitivity checks: among others, we inspect alternative measures of inflation and real economic activity, consider alternative sample periods, as well as extensions of the SVAR model which include survey data on inflation expectations, or data on interest rates. A flatter demand slope, and a relative rise in supply-driven output gap volatility, emerge in all cases. A stable (or even steepening) supply slope is found in the vast majority of experiments. Thus, our simple arithmetics speaks clearly in favor of a policy explanation: the observed decline in inflation volatility, and the emerging disconnect between inflation and real economic activity, seem to arise mainly because the Federal Reserve has become more firmly committed to its inflation target over time. The slope and shock hypotheses, by contrast, are largely rejected when the simple arithmetics are applied to conditional volatility in the data.

How does this paper speak to existing literature? We are definitely not the first to highlight the potentially important role of supply shocks for inflation dynamics (cf. Hobijn (2020), Gordon (2013), Hasenzagl, Pellegrino, Reichlin, and Ricco (2022) among others). However, our contribution is to show how a joint assessment of potential changes in demand and supply slopes, which in turn requires that we identify both types of shocks, can be key to understanding the apparent disconnect between inflation and economic activity. In addition, we stress the importance of controlling for *all* supply shocks. The literature often controls only for cost-push shocks, a very specific type of supply shock.

Most importantly, our paper contributes to a vast empirical literature studying the structural relationship between inflation and economic activity. Traditionally, most papers discuss approaches and challenges to the estimation of the Phillips curve in a single 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 weak instrument problems. The authors conclude that new datasets and new identification approaches are needed to reach an empirical consensus.

In terms of new datasets, Imbs, Jondeau, and Pelgrin (2011) estimate Phillips curves at the sectoral level using French data and then derive implications for monetary policy. More recently Hazell, Herreno, Nakamura, and Steinsson (2022) estimate the slope of the Phillips curve in the cross-section of US states and find only a modest decline in the slope of the Phillips curve since the 1980s. They also construct an aggregate Phillips curve slope and conclude that there is no missing inflation or deflation in the most recent business cycles. Similar results are found by Fitzgerald, Jones, Kulish, and Nicolini (2020) who use US city and state level data. Finally, 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 on Phillips curve estimation has engaged in a search for better instruments for demand-driven economic activity. Barnichon and Mesters (2020), for example, find that conventional methods (including the use of predetermined variables as instruments) substantially underestimate the slope of the Phillips curve. They instead exploit identified monetary policy shocks. Taking a multivariate approach, Del Negro et al. (2020) show that post-1990s inflation barely reacts in response to a shock to the excess bond premium, i.e. a shock that behaves like a typical demand shifter. Relatedly, Ascari and Fosso (2024) estimate a SVAR with common trends and find evidence of more muted inflation responses to business cycle shocks in recent years. A flattening of the Phillips curve is also found by Inoue et al. (2022) in the context of an instrumental variables model with time-varying parameters while Wong (2013) finds only a limited flattening. While admittedly simple, our set-up allows us to evaluate all of the three explanations for the inflation puzzle in a *unified* framework, based on a set of simple identification restrictions derived from theory. To the best of our knowledge, this is the first paper to provide such a joint analysis.[1] The paper closest to us is perhaps Galí and Gambetti (2020), who use a SVAR model to purge the data for wage mark-up shocks when estimating the wage Phillips curve. We instead focus on the conventional price Phillips curve, and also stress the importance of controlling for all supply shocks. Finally, the use of conditional regressions on data filtered through a SVAR model is pursued also in Debortoli, Galí, and Gambetti (2020) to estimate a monetary policy rule.

The rest of the paper is organized as follows: Section 2 uses a textbook New Keynesian model to discuss and formally derive some simple arithmetics which govern the structural relationship between output and inflation. Section 3 describes our methodological approach. Section 4 documents the main empirical results while Section 5 provides a battery of robustness tests. Section 6 relates our findings to recent events and to selected contributions in the literature. Finally, Section 7 concludes.

[1] Historically the SVAR 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

To organize ideas, we start with a textbook New Keynesian model in log-linearized form (see Woodford (2003) and Galí (2008) for further details). It 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[2]: 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, respectively, as exogenous productivity, labor supply, and cost-push shock disturbances. u_t is a demand or discount factor disturbance. 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 captures exogenous deviations from the rule, so-called monetary policy disturbances. Importantly, stricter inflation targeting is captured by a rise in ϕ_π.

Two simplifications will be useful for our purpose: first, one can combine equations (2)-(4) with (5) in order to 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 we have introduced the parameter $\kappa = \lambda(\sigma + \varphi)$ and collected the three supply disturbances in $s_t = z_t + \lambda\psi_t - \lambda(1+\varphi)a_t$. Second, one can use equation (6) to substitute out the nominal interest rate i_t from (1). Collecting the two demand disturbances in $d_t = u_t - m_t$ we arrive at

$$y_t = \frac{1}{\sigma + \phi_y}\left(\sigma \mathbb{E}_t y_{t+1} - \phi_\pi \pi_t + \mathbb{E}_t \pi_{t+1} + d_t\right). \tag{8}$$

The two equations above summarize our main objects of interest: first, aggregate supply is given by the New Keynesian Phillips curve in (7) and naturally upward-sloping in the (y_t, π_t)-space. A flattening of the structural Phillips curve amounts to a decline in

[2] Variable deviations from a counterfactual flex-price equilibrium, in contrast, are important for welfare and appear frequently in textbooks. However, the notion of e.g. output in deviation from some trend is more in line with 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.

κ. Second, aggregate demand (accounting for the central bank's reaction function) is instead downward-sloping in the (y_t, π_t)-space and given by (8).[3] Stricter inflation targeting amounts to a rise in ϕ_π and a flatter (less negatively sloped) demand curve. Importantly, the weakened relationship between inflation and output must stem from a combination of changes in demand and supply curve slopes, as well as changes in how much the two curves shift around.

A few key challenges emerge for applied researchers hoping to understand the empirical disconnect between inflation and output: first, we observe only the evolving intersection between demand and supply, i.e. the realization of shifts in the two curves. Their slopes, as well as changes in these slopes over time, are not directly observable. This fundamental issue is essentially what we aim to address in the current paper. Second, since the model variables are expressed in deviation from steady state or trend rather than the flex-price counterfactual, all three supply shocks enter directly in equation (7).[4] Thus, simply controlling for cost-push shocks by adding e.g. energy prices to the empirical specification is not sufficient: we need to control for all supply shocks when estimating (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 if we want to recover the supply slope. Likewise, the components of s_t shift the supply curve and may, therefore, help us recover the demand curve slope.

In order to further discuss the identification challenges involved, and to highlight our proposed identification strategy, we find it instructive to work with the model's solution. To this end we 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. This i.i.d. assumption is relaxed in Appendix A where we instead consider the more common assumption that shocks follow separate AR(1) processes. None of our conclusions are altered in this case. Analytical solutions for output and inflation follow:

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

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

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

using data generated from this stylized model. The projection coefficient γ is best understood as a *Phillips correlation*, along the lines of Reis and Watson (2010) and Stock and Watson (2020). It may or may not be informative about the *structural Phillips curve slope* κ. Importantly, the analytical solutions above allow us to derive expressions for the $\gamma \equiv \frac{cov(\pi_t, y_t)}{var(y_t)}$ in closed form. What can these closed form expressions tell us about demand and supply curve slopes? That critically depends on our treatment of data.

[3] A similar picture emerges if we instead assume optimal monetary policy: suppose that the central bank minimizes $\frac{1}{2}\left[(\pi_t + m_t)^2 + \alpha y_t^2\right]$ subject to equations (1) and (7), and under full discretion. The weight α captures the relative importance of a stable output gap while m_t is interpreted as a monetary policy shock. The implied targeting rule, $\pi_t = -\frac{\alpha}{\kappa} y_t - m_t$, is once again a downward sloping demand curve in the (y_t, π_t)-space.

[4] If we consider output in deviation from its counterfactual when prices are flexible, then only the cost-push shock enters the Phillips curve.

2.1 The unconditional projection coefficient

Suppose, first, that all the observable variation in data is explored. The resulting *unconditional* projection coefficient estimate, denoted by $\hat{\gamma}_u$, is given by

$$\begin{aligned}\hat{\gamma}_u &= \frac{\kappa\left(\sigma_m^2+\sigma_u^2\right)-\phi_\pi\left(\sigma+\phi_y\right)\sigma_z^2-\frac{\phi_\pi\kappa^2(\sigma+\phi_y)}{(\sigma+\varphi)^2}\sigma_\psi^2-\frac{\phi_\pi\kappa^2(\sigma+\phi_y)(1+\varphi)^2}{(\sigma+\varphi)^2}\sigma_a^2}{\sigma_m^2+\sigma_u^2+\phi_\pi^2\sigma_z^2+\left(\frac{\phi_\pi\kappa}{\sigma+\varphi}\right)^2\sigma_\psi^2+\left(\frac{\phi_\pi\kappa(1+\varphi)}{\sigma+\varphi}\right)^2\sigma_a^2}\\ &= \frac{\kappa-\phi_\pi\left(\sigma+\phi_y\right)\frac{\sigma_s^2}{\sigma_d^2}}{1+\phi_\pi^2\frac{\sigma_s^2}{\sigma_d^2}}. \end{aligned} \tag{9}$$

Here $\sigma_d^2 = \sigma_m^2 + \sigma_u^2$ 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.

A couple of remarks are in place: first, $\hat{\gamma}_u$ is 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 for reasons unrelated to the Phillips curve. A naive econometrician observing a decline in $\hat{\gamma}_u$ may erroneously conclude that the Phillips curve has flattened even if κ has remained unchanged. The decline in $\hat{\gamma}_u$ could very well be due to stricter inflation targeting ($\phi_\pi \uparrow$), or to a rise in the *relative* volatility of supply-side shocks ($(\sigma_s^2/\sigma_d^2)\uparrow$). The expression for $\hat{\gamma}_u$ essentially reveals that one cannot use unconditional data to separate Phillips curve changes from changes in policy or shocks. Luckily, the good news is that we can exploit the variation in conditional data, a point which we turn to next.

2.2 Conditional projection coefficients

Suppose that we are able—somehow—to purge the data for all variation due to supply-side shocks. This implies setting $\sigma_s^2 = 0$ in (9). The associated projection coefficient estimate follows below:

$$\hat{\gamma}_s = \kappa \tag{10}$$

We interpret $\hat{\gamma}_s$ as *the supply curve slope*, given that it arises solely from shifts in demand. This slope is clearly an unbiased estimate of κ. Thus, if we somehow can trace changes in $\hat{\gamma}_s$, 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 $\hat{\gamma}_s$ in our simple framework.

Next, suppose that we are able—somehow—to purge the data for all variation due to demand side shocks. The implied estimate, this time conditional on supply shocks only, follows below:

$$\hat{\gamma}_d = -\frac{\sigma+\phi_y}{\phi_\pi} \tag{11}$$

Figure 2: Three alternative explanations – simulated data

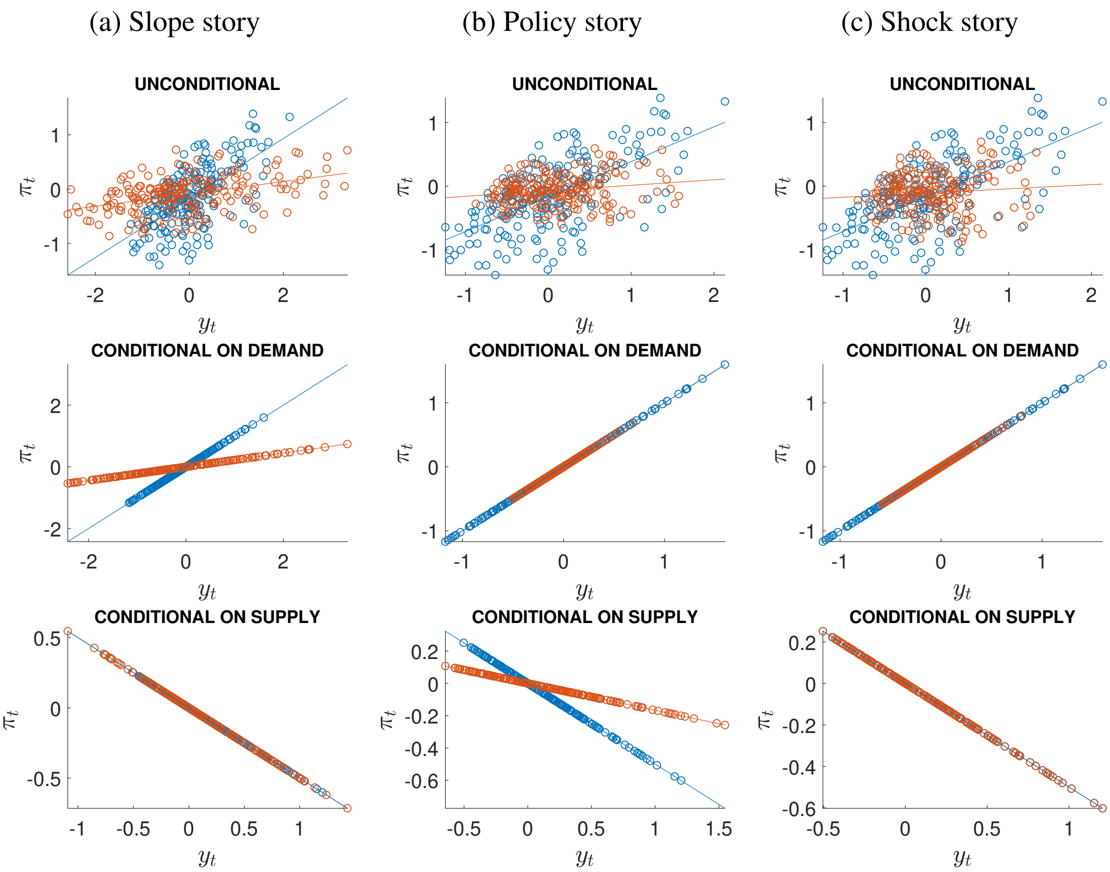

Notes: Each column in the figure represents a separate explanation—the slope story is illustrated in column (a), the policy story in column (b), and the shock story in column (c). The top subplot in each column reports the unconditional (simulated) data augmented with the projection coefficient $\hat{\gamma}_u$. Data conditional only on demand shocks and augmented with $\hat{\gamma}_s$ are reported in the second row. Data conditional on supply shocks and augmented with $\hat{\gamma}_d$ are reported in the third. 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 3 across the samples, while the shock story is a decline in the volatility of demand shocks by 50%.

We refer to the projection coefficient $\hat{\gamma}_d$ as *the demand curve slope*. As such, $\hat{\gamma}_d$ is negative and purely a function of the two policy parameters as well as σ, all of which are important for the responsiveness of *demand* to prices. Stricter inflation targeting, in particular, makes $\hat{\gamma}_d$ flatter (less negative). Changes in the Phillips curve slope or in the composition of shocks play no role for $\hat{\gamma}_d$ in our framework.

Finally, the fact that both $\hat{\gamma}_s$ and $\hat{\gamma}_d$ 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 shock 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] Consider first the slope story (first column). In this case the following applies: if the observed decline in $\hat{\gamma}_u$ (first row) is a consequence predominantly of a flatter Phillips curve slope κ, then we should also see a decline in $\hat{\gamma}_s$ (second row), combined with a relatively stable $\hat{\gamma}_d$ (third row). If, instead, the decline in $\hat{\gamma}_u$ is driven by stricter inflation targeting (second column), then we should also find a relatively stable $\hat{\gamma}_s$ over time (second row), combined with a flattening (towards less negative values) of $\hat{\gamma}_d$ (third row). Finally, if the cause of a decline in $\hat{\gamma}_u$ is that supply side shocks have become more volatile relative to demand side shocks (third column), then we should find relatively stable values of both $\hat{\gamma}_s$ and $\hat{\gamma}_d$ across sub-samples (second and third row). Note that, in each column, the unconditional scatter plot in the top plot is the sum of the two scatter plots below.

In total we have three alternative explanations of an observed decline in $\hat{\gamma}_u$: (i) the slope story $\{\hat{\gamma}_s \downarrow, \hat{\gamma}_d \text{ unchanged}\}$, (ii) the policy story $\{\hat{\gamma}_s \text{ unchanged}, \hat{\gamma}_d \uparrow\}$, and (iii) the shock story $\{\hat{\gamma}_s \text{ and } \hat{\gamma}_d \text{ unchanged}\}$. Together, these *simple arithmetics* constitute a set of testable implications which we pursue empirically in the remainder of the paper.

2.3 Conditional variances

Conditional variances provide additional and complementary implications that can be used to disentangle the slope hypothesis, the policy hypothesis and shock hypothesis. The variance decomposition for the output gap can be expressed with the following analytical 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.[6] 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 $VD\left(y|d\right)$. Intuitively, ceteris paribus the output gap responds more to a given shift in demand (supply) if the supply (demand) curve becomes flatter. These implications for conditional variances are also visible in Figure 2. When comparing sub-samples, we see that demand shocks become relatively more important for output when the slope story applies, while the opposite is true for the policy story. Note that a declining role of demand shocks for output gap fluctuations naturally emerges also as a consequence of the shock story. However, an increase in the relative volatility of supply shocks would imply a constant demand slope $\hat{\gamma}_d$ over time, as opposed to the flattening predicted by stricter monetary policy. It is in this way that variance decompositions and slope inspections can provide complementary insights.

[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 $\sigma_u = 1$, $\sigma_m = \sigma_a = \sigma_\psi = 0.25$, 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]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 would be replaced with either σ_m^2 or σ_u^2, respectively.

2.4 TRANSMISSION LAGS

Are the simple arithmetics discussed above jointly exhaustive, or could important candidate explanations be left out? A fourth alternative may be of particular interest, namely an event that triggers a simultaneous shift in both demand and supply curves. The simple model considered here features a so-called divine coincidence. It results from perfect proportionality between marginal costs and the output gap and allows us to abstract from simultaneous shifts in demand and supply. But proportionality seizes to hold if, for example, wages are rigid or in the presence of an active fiscal policy. Thus, in Appendix B we introduce fiscal policy and lay out the implications for our simple arithmetics. Equipped with the analytical solution, we identify potential biases when simultaneous shifts occur in response to government spending shocks, and argue that our simple arithmetics prove useful even in such cases.

The remainder of this section, instead, focuses on another violation of divine coincidence, namely the presence of transmission lags. Lags of various endogenous variables (or "bells and whistles") are often included in quantitative models to better account for well-documented empirical moments in data (persistent and hump-shaped dynamics, for example). The implications of transmission lags for our simple arithmetics are documented in two complementary ways: first, we add some of the most commonly used bells and whistles to the baseline model, but keep the assumption of Calvo pricing. Second, we disregard Calvo pricing altogether and instead assume sticky information as a source of sluggish price adjustment.

2.4.1 HABITS, INDEXATION, AND INTEREST RATE INERTIA

Here we extend the baseline model along three dimensions: first, it is assumed that households derive utility from their consumption relative to the previous period's aggregate consumption, so-called (external) consumption habits. Second, firms that do not get to update prices optimally are assumed to index prices partially to lagged inflation. Third, we add interest rate inertia to the central bank's reaction function. These extensions imply that equations (1), (6) and (7) can be written as follows:

$$
\begin{aligned}
y_t &= \mathbb{E}_t\left(y_{t+1} - \alpha y_t\right) + \alpha y_{t-1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - u_t\right) \\
\pi_t &= \beta \mathbb{E}_t\left(\pi_{t+1} - \zeta \pi_t\right) + \zeta \pi_{t-1} + \lambda mc_t + z_t \\
i_t &= \rho_i i_{t-1} + (1-\rho_i)\left(\phi_\pi \pi_t + \phi_y y_t\right) + m_t
\end{aligned}
$$

The parameter α governs the importance of habit persistence, ζ the degree of inflation indexation, and ρ_i the degree of interest rate inertia. The rest of the model is unchanged. Thus, $\alpha = \zeta = \rho_i = 0$ brings us back to the baseline.

In Figure 3 we redo the scatter plots in Figure 2, but with the calibration $\alpha = 0.75$, $\zeta = 0.5$, and $\rho_i = 0.75$ as a starting point for the simulations. These numbers are in the ballpark of those typically used in quantitative macro models. The presence of transmission lags imply that conditional scatter plots become clouds rather than being perfectly aligned along the regression line. Past realizations affect current outcomes, and changes in the environment (e.g. changes in structural slopes or volatilities) will in general impact both regression lines. Nevertheless, it is clear that the inclusion of transmission lags does not significantly alter our simple arithmetics: a flattening of $\hat{\gamma}_s$ in

Figure 3: Simulated data from a model with additional rigidities

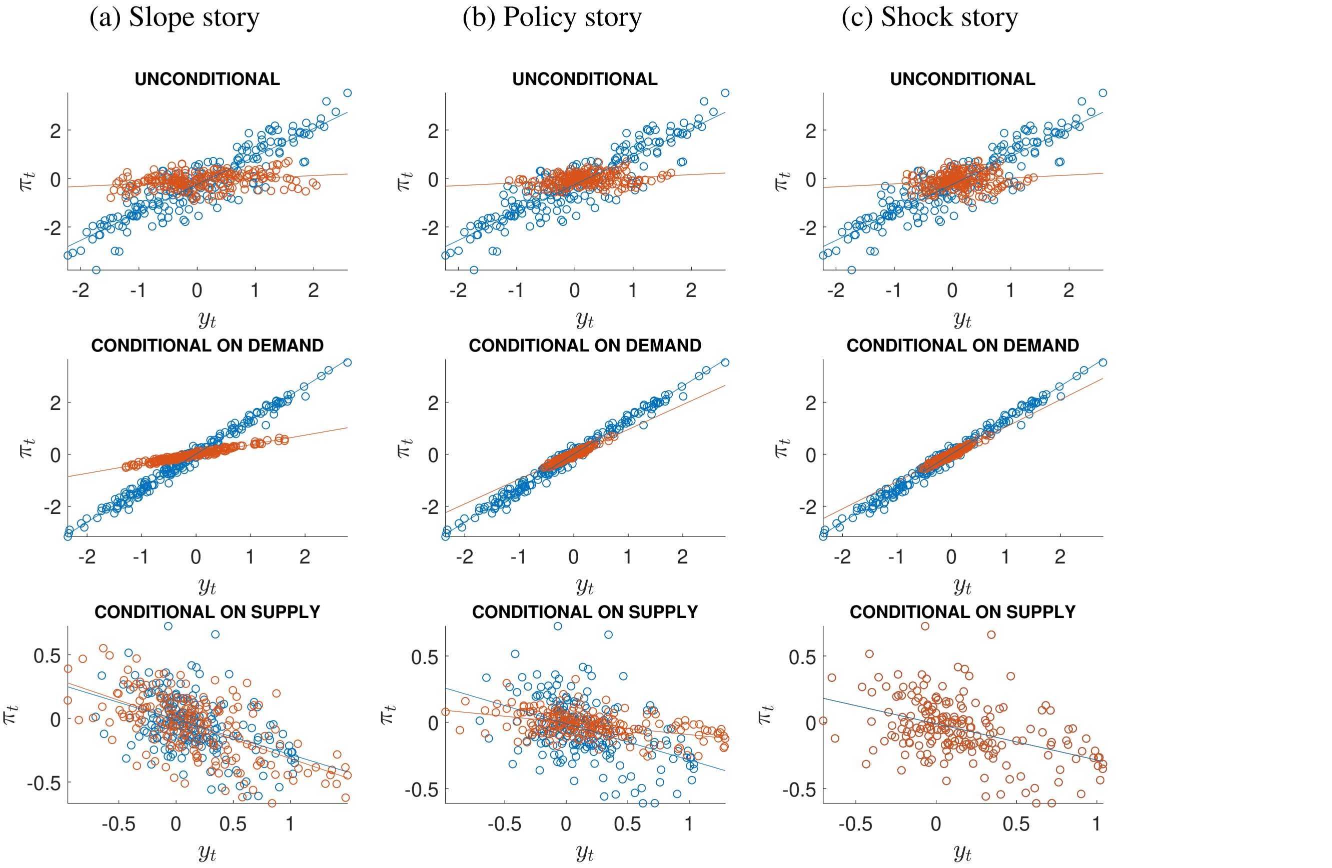

Notes: The slope story in column (a) is illustrated with a rise in the Calvo parameter from 0.75 in sample 1 (blue) to 0.875 in sample 2 (red). The policy story in column (b) is represented by a rise in ϕ_π from 1.5 to 2.5 across the samples, while the shocks story in column (c) arises from a decline in the volatility of demand shocks by 45%. See Figure 2 for additional details.

combination with a relatively stable $\hat{\gamma}_d$ is still a unique feature of the slope story. Stability in $\hat{\gamma}_s$ combined with a flatter (less negative) $\hat{\gamma}_d$ is still a unique feature of the policy story. Finally, the shock story still implies stability in both $\hat{\gamma}_s$ and $\hat{\gamma}_d$.

2.4.2 Sticky information

In this section, we replace Calvo pricing with sticky information á la Mankiw and Reis (2002) and Reis (2009). To this end, it is assumed that agents update their information set infrequently: equilibrium outcomes today depend on past expectations about the current state of affairs, giving rise to delayed and muted responses to shocks. Following Reis (2009) we allow for sticky information both in the Phillips curve and in the IS equation:

$$p_t = \delta_p \sum_{j=0}^{\infty} (1-\delta_p)^j \, \mathbb{E}_{t-j} \left(p_t + mc_t + z_t\right)$$

$$p_t = p_{t-1} + \pi_t$$

Figure 4: Simulated data from a model with sticky information

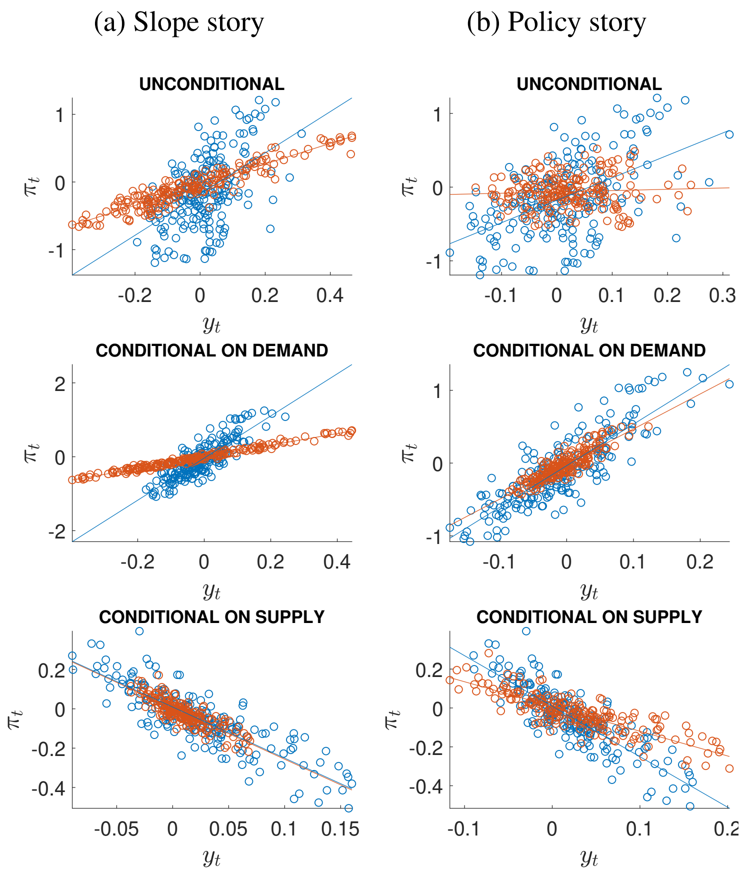

Notes: The slope story in column (a) is illustrated with a 50% decline in the information updating parameter δ_p, from 0.52 in sample one (blue) to 0.26 in sample two (red). The policy story in column (b) is shown as a rise in ϕ_π from 1.5 to 3 across the samples, while the shock story in column (c) is a decline in the volatility of demand shocks by 50%. See Figure 2 for additional details.

$$y_t = -\delta_y \frac{1}{\sigma} \sum_{k=0}^{\infty} (1-\delta_y)^k \, \mathbb{E}_{t-k} R_t$$

$$R_t = \sum_{k=0}^{\infty} \mathbb{E}_t \left[i_{t+k} - \pi_{t+1+k} - u_{t+k} \right]$$

The first two equations replace equation (6) and represent the Phillips curve under sticky information. Each period, only a fraction δ_p of firms update their information set and set prices accordingly. The remaining firms set prices based on what they thought would be optimal the last time they acted on new information. Likewise, the last two expressions above replace equation (1) in the baseline. Here only a fraction δ_y of consumers update their information set each period. R_t represents the long real interest rate after taking into account preference shocks. Note that $\delta_p = \delta_y = 1$ brings us back to full information and flexible prices, i.e. the special case of the baseline model with $\theta = 0$.

In Figure 4 we redo the scatter plots in Figure 2, but with the calibration $\delta_p = 0.52$ and $\delta_y = 0.08$ as a starting point for the simulations.[7] These numbers are taken from

[7] When solving the sticky information model, we follow Verona and Wolters (2013) and truncate the infinite

the posterior mean estimates in Reis (2009). Once again, we find that the inclusion of transmission lags—this time due to sticky information—does not significantly alter the simple arithmetics presented earlier: only the slope story implies a flattening of $\hat{\gamma}_s$ and a stable $\hat{\gamma}_d$. Only the policy story implies stable $\hat{\gamma}_s$ and flatter $\hat{\gamma}_d$. Finally, only the shock story implies unchanged slopes based on the conditional data.

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. We split the sample in 1994:Q4 in order to obtain two samples of equal size and we consider alternative splittings in the sensitivity analysis. The deterministic component of the SVAR, which is obviously different across samples, captures the dynamics of trend inflation over time.

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 notation and the estimation of the SVAR model follow Arias, Rubio Ramirez, and Waggoner (2018). The structural representation of the model can be written as follows:

$$Y_t' A_0 = x_t' A_+ + \varepsilon_t' \qquad (12)$$

where Y_t is a $n \times 1$ vector of endogenous variables, ε_t is an $n \times 1$ vector of exogenous structural shocks, A_0 and $A_+' = [A_1', \ldots, A_p', c']$ are matrices of parameters with A_0 invertible, and $x_t' = [Y_{t-1}', \ldots, Y_{t-p}', 1]$ for $1 \leq t \leq T$, with c a $1 \times n$ vector of parameters, p the lag length, and T the sample size. The vector ε_t, conditional on past information and the initial conditions $Y_0, ..., Y_{1-p}$, is Gaussian with mean zero and covariance matrix I_n, the $n \times n$ identity matrix. The dimension of A_+ is $m \times n$ where $m = np + 1$. The reduced-form representation implied by Equation (12) is

$$Y_t' = x_t' B + u_t' \qquad (13)$$

where $B = A_+ A_0^{-1}$, $u_t' = \varepsilon_t' A_0^{-1}$, and $E[u_t u_t'] = \Sigma = (A_0 A_0')^{-1}$. The matrices B and Σ are the reduced-form parameters, while A_0 and A_+ are the structural parameters.

Our baseline SVAR contains two variables observed at quarterly frequency:

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

π_t represents one-period inflation in the GDP deflator while y_t represents the quarterly output gap computed by the CBO.

sums that arise from sticky information. Verona and Wolters (2013) document reasonably precise results with 16 expectation lags. We include 32 lags to ensure accuracy.

Table 1: Sign restrictions - SVAR models

Panel (A)	Demand ↑	Supply ↓		
Inflation	+	+		
Output gap	+	-		
Panel (B)	Demand ↑	Supply ↓	Residual	
Inflation	+	+	+	
Output gap	+	-	$\star$	
Inflation expectations	+	+	-	
Panel (C)	Demand ↑	Supply ↓	Policy ↓	
Inflation	+	+	+	
Output gap	+	-	+	
Interest (or shadow) rate	+	$\star$	-	
Panel (D)	Demand ↑	Technology ↓	Labor Supply ↓	Policy ↓
Inflation	+	+	+	+
Output gap	+	-	-	+
Shadow rate	+	$\star$	$\star$	-
Real wages	$\star$	-	+	$\star$

Note: Restrictions are imposed only on impact. The notation $\star$ means that no restriction is imposed. In Panel (B) we add data on inflation expectations to the baseline setup. In Panel (C) we add interest rates, while Panel (D) includes both interest rates and real wages as observables.

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 and impose sign restrictions on impact (and only on impact, as recommended by Canova and Paustian (2011)) to identify the structural shocks. The QR decomposition algorithm proposed by Arias et al. (2018) is used for this purpose.[8] The algorithm is continued until we have obtained $10,000$ draws that satisfy the imposed sign restrictions.[9]

Sign restrictions are specified in Table 1. Panel A summarizes our baseline identification scheme, which disentangles demand shocks from supply shocks based on the impact co-movement between inflation and the output gap.

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, at each point in time (i.e.

[8]This algorithm enables us 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.

[9]An important and interesting debate on the choice of priors in sign-restricted SVAR models is at play at the frontier of the literature. Advantages and disadvantages of the conventional approach based on Arias et al. (2018) are discussed in Baumeister and Hamilton (2015), Inoue and Kilian (2020) and Arias, Rubio Ramirez, and Waggoner (2022). We acknowledge the importance of this debate.

each quarter), the median contribution of each shock to fluctuations in inflation and in the output gap. The final outcome of this exercise is one historical decomposition, with each shock's contribution representing the median across the $10,000$ alternative models. Note that such a way to compute historical decompositions is recommended by Bergholt, Canova, Furlanetto, Maffei-Faccioli, and Ulvedal (2024) to account properly for the uncertainty surrounding the estimated deterministic component of the SVAR. Similarly, we obtain a summary measure for the variance decomposition.

At this stage, some additional motivation for our identification strategy is in order. The main advantage of using sign restrictions is that, in the baseline bivariate setup, they form a set of mutually exclusive and jointly exhaustive identification restrictions. Only the *joint* identification of the two shocks in the same framework allows us to make statements on both the supply and the demand curves slopes. Previous papers focusing on specific demand shocks (like monetary policy shocks in Barnichon and Mesters (2020) or shocks to the excess bond premium in Del Negro et al. (2020)) were not equipped to identify changes in the slope of the demand curve and thus could not exploit fully the simple arithmetics implied by the New Keynesian theory. Identification schemes based on the Cholesky decomposition identify usually only one shock of interest on the basis of timing assumptions. Proxy measures of shocks obtained from external sources are not jointly exhaustive since they do not explain fully the variation in the data while our shocks exhibit that property by construction. Identification schemes based on long-run restrictions can, in principle, disentangle supply shocks (defined as shocks with long-run effects on output) and demand shocks (defined as shocks with purely transitory effects on output) in the same set-up. However, Furlanetto, Lepetit, Robstad, Rubio Ramírez, and Ulvedal (2024) show that the shock with long-run effects cannot (and should not) be associated only with a supply shock in a large part of our sample: the shock with permanent effects on output in a Blanchard-Quah decomposition would look like a demand shock in our second sample. Against this background, we see the use of sign restrictions as the preferable way to investigate our simple arithmetics in the data. The disadvantages of using such a framework are that the model is set-identified (and therefore subject to model uncertainty in addition to estimation uncertainty) and to the fact that several demand shocks (like government spending, financial, monetary and discount factor shocks, just to name a few) and several supply shocks (like technology, labor supply, mark-up and many others) are commingled into only two innovations.[10] All in all, we believe that the benefit of joint exaustivity (i.e. the fact that the two shocks explain the entire variation in unconditional data) are larger than the costs paid in terms of uncertainty and commingling of shocks.

4 Results

This section presents estimates of conditional slopes and shock decompositions, and summarizes the main results when we confront data with the simple arithmetics derived from theory.

[10]Note, however, that it would be straightforward to identify several demand and supply shocks by imposing additional restrictions in the context of a larger SVAR including more variables. Such an exercise might significantly complicate the interpretation of the results. Nevertheless, we inspect larger SVAR models with more shocks in various robustness exercises, see section 5.

Figure 5: Empirical scatterplots

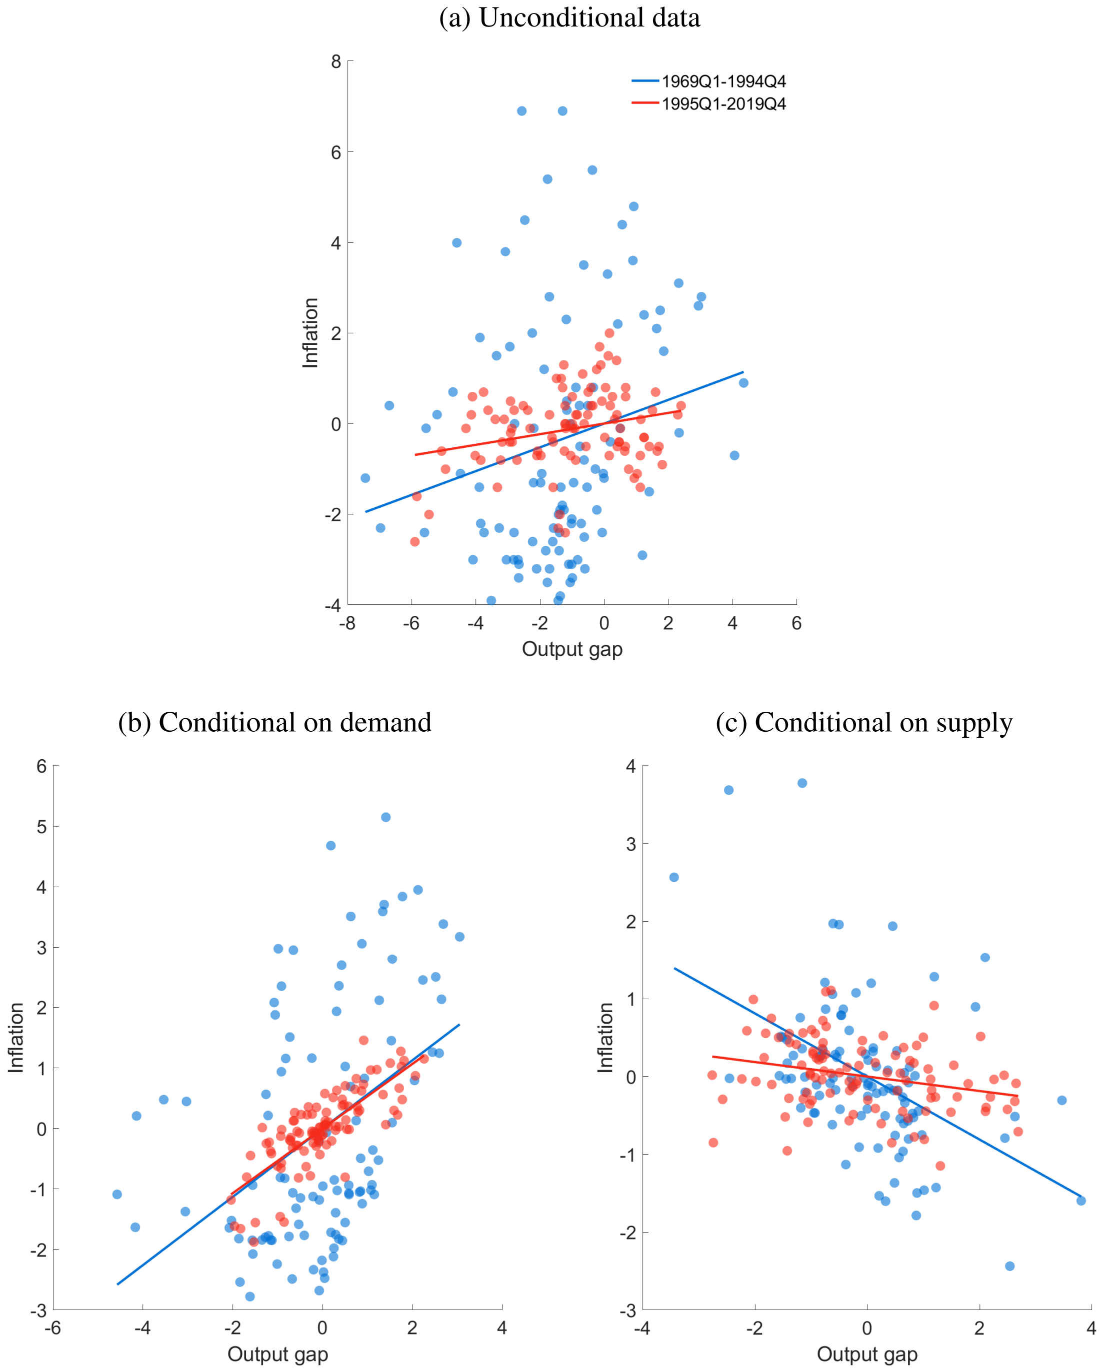

Notes: Unconditional data vs. conditional data obtained from the estimated SVAR model. Corresponding slope estimates are provided in Table 2.

4.1 From unconditional to conditional slopes

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

$$\pi_t = c_1 (1 + D_t \delta_c) + \gamma_1 (1 + D_t \delta_\gamma) y_t + u_t \tag{14}$$

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 projection coefficients for the output gap across samples.[11] In particular, we denote the projection coefficient in the first sample by γ_1, and the projection 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.

To assess the competing explanations for a flatter statistical slope, we estimate equation (14) both on unconditional data and on conditional data generated by the SVAR model in the filtering step. This allows us to contrast the unconditional estimates $\hat{\gamma}_{u,1}$ and $\hat{\gamma}_{u,2}$ with the estimates based on conditional data across the two samples. We obtain $\hat{\gamma}_{s,1}$ and $\hat{\gamma}_{s,2}$ from data purged for supply shocks, and $\hat{\gamma}_{d,1}$ and $\hat{\gamma}_{d,2}$ from the data purged for demand shocks. This leaves us with a set of estimated projection 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 5a. The horizontal axis measures observations of the CBO's output gap, the vertical axis measures inflation in the GDP deflator (plotted in deviation from its mean $\pi_t - c_1 (1 + D_t \delta_c)$). The scatter plot is augmented with the regression lines which represent the sample-specific, unconditional projection coefficients $\hat{\gamma}_{u,1}$ and $\hat{\gamma}_{u,2}$.

Panel A in Figure 5 illustrates a major decline in the projection coefficient estimated on unconditional data. Quantitatively, the estimated slope is $\hat{\gamma}_{u,1} = 0.26$ during the sample period 1969:Q1-1994:Q4, but only $\hat{\gamma}_{u,2} = 0.12$ during the period 1995:Q1-2019:Q4. Thus, we observe a decline of more than 50% in the unconditional slope in the second sample. We note that this decline has limited statistical significance given that the p-value associated with $\hat{\delta}_{u,y}$ is 0.28. Nevertheless, as stated earlier, the decline in the unconditional slope 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 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 (to a lesser extent) 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 5 plots the relationship between output gap and inflation when we condition on the empirical variation attributed solely to identified demand shocks. Several important observations emerge: first, the estimated projection coefficient is substantially higher than its unconditional counterpart. This is reassuring given its interpretation as a supply slope, as opposed to the likely downward bias in $\hat{\gamma}_u$, as emphasized earlier. Second, the supply slope features only a minor decline across samples, from $\hat{\gamma}_{s,1} = 0.56$ to $\hat{\gamma}_{s,2} = 0.53$. This decline is far from significant statistically, with a p-value as high as

[11] Estimation of the interacted regression equation (14) on pooled data is equivalent to splitting the sample and running two separate regressions.

Figure 6: Variance decomposition of CBO's output gap

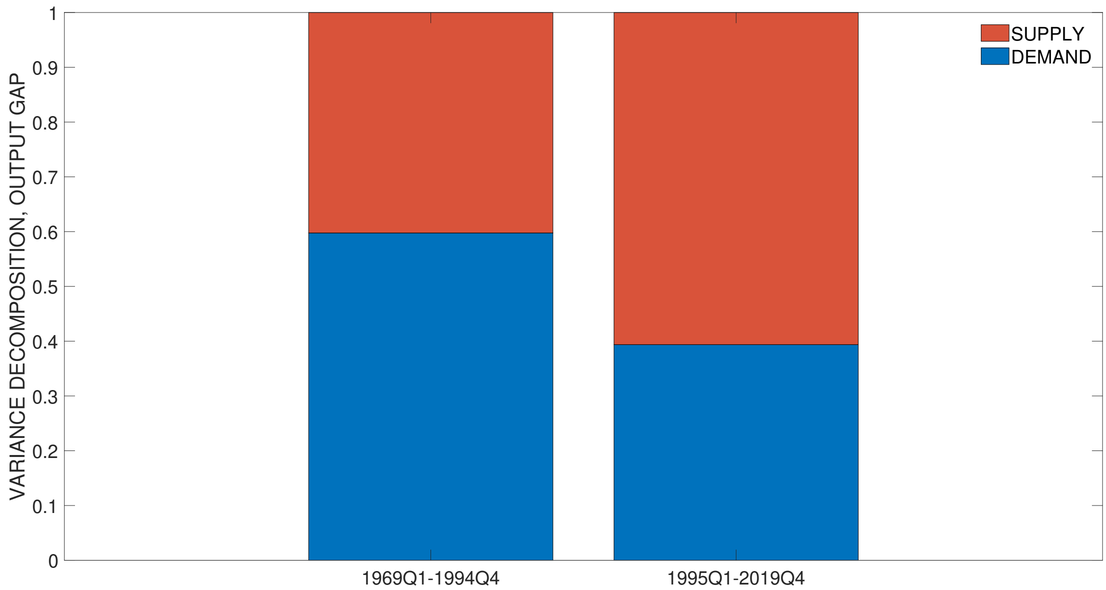

0.85 for $\hat{\delta}_{s,y}$. Thus, the relationship between output and inflation, which has weakened substantially in unconditional data for the US economy, remains almost unchanged once we purge out supply-side variation. In this sense, we do not find empirical evidence of a flattening of the Phillips curve. Rather, the Phillips curve seems to be alive and well. Third, the reduced volatility of both inflation and output gap carries over when we zoom in on demand-driven variation in the data.

Finally, Panel C in Figure 5 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 estimated projection slopes across samples. In fact, the slope goes from $\hat{\gamma}_{d,1} = -0.41$ to $\hat{\gamma}_{d,2} = -0.09$ and the p-value associated with $\hat{\delta}_{d,y}$ is 0.00. Thus, the null hypothesis of no change in the demand slope across samples is rejected at any relevant significance level. It is clear that our estimated demand curve slope changes from highly negative to relatively flat. All in all, we conclude that there seems to be a significant flattening of the demand slope, and this flattening is substantially larger than what we find when conditioning on demand shocks. Thus, the evidence reported here points to a flattening of the demand curve rather than the supply curve.

4.2 Conditional variances

Sample-specific variance decompositions of the CBO output gap are summarized in Figure 6. These variance decompositions are based on the scatterplots presented in Figure 5. Supply shocks, in particular, explain 40% of the output gap in the first sample, but 61% in the second sample. Consistent with the discussion in subsection 2.3, an increasing role of supply shocks for the output gap is consistent with the view that the central bank has focused more on inflation targeting, and at the same time speaks against a flattening of the structural Phillips curve. Ceteris paribus, we would instead have expected to see a more dominant role of demand shocks in the second sample, had the slope of the Phillips curve declined.

4.3 Discussion

Recall our estimates for the regression slope conditional on demand shocks. They are centered around 0.5 in both samples. Importantly, this number should not be interpreted as evidence of a very steep Phillips curve. In fact, $\hat{\gamma}_s$ is unbiased relative to κ only if shocks are i.i.d., as in the stylized model analyzed in section 2. The more general case with persistent shocks is discussed in Appendix A, which shows that persistence renders $\hat{\gamma}_s$ proportional to κ, but with an upward bias (see also Hazell et al. (2022) for a discussion on this point). Just as an illustrative example; if demand shocks exhibit an autocorrelation of 0.8, then $\hat{\gamma}_s = 0.5$ would imply a value of κ around 0.1, similar to values often used in calibrated versions of the New Keynesian model. The only way to justify treating $\hat{\gamma}_s$ as an unbiased estimate of κ when shocks are persistent, is to properly control for expectations as well. Appendix A provides further analytical details and discussions.

Our results on conditional slopes and conditional variances are consistent with a flattening of the demand curve, whose slope in the canonical New Keynesian model is given by the projection coefficient $\hat{\gamma}_d = -\frac{\sigma+\phi_y}{\phi_\pi}$. We believe the most natural interpretation for such a flattening is a more aggressive Fed response to inflation fluctuations (i.e. an increase in the coefficient ϕ_π relative to the coefficient ϕ_y), in keeping with the estimates presented in the seminal paper by Clarida, Galí, and Gertler (2000). At the same time, we acknowledge that an alternative interpretation may appear plausible to some: a higher impact of interest rates on aggregate demand (lower σ) could also explain a flattening of the demand curve. This could reflect better consumption smoothing opportunities for example induced by more widespread asset market participation. However, two additional theoretical implications are in contrast with this alternative interpretation. First, a decline in σ should also translate into a decline in $\hat{\gamma}_s$ (recall that $\kappa = \lambda(\sigma + \varphi)$), Our SVAR does not support this implication, see Figure 5. Second, a decline in σ should leave the variance decomposition unaffected across samples, in contrast with the evidence presented in Figure 6. This latter implication follows if we substitute $\kappa = \lambda(\sigma + \varphi)$ into the variance decomposition of output, given analytically in subsection 2.3. Nonetheless, 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)). In light of this evidence, an increase in stock market participation could rationalize a flattening of the demand curve. However, Albonico, Ascari, and Haque (2023) find that such an increase has been modest (and irrelevant for macroeconomic dynamics), thus casting doubts on this alternative explanation for the flattening of the demand curve.

Finally, our simple arithmetics can be challenged in richer and more complicated models. 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 setting would be to introduce a working capital channel, so that the interest rate affects marginal costs directly (Ravenna and Walsh, 2006). However, this would lead to a steeper Phillips curve slope when ϕ_π increases, an implication that seems less relevant in this context.

Figure 7: Phillips multipliers

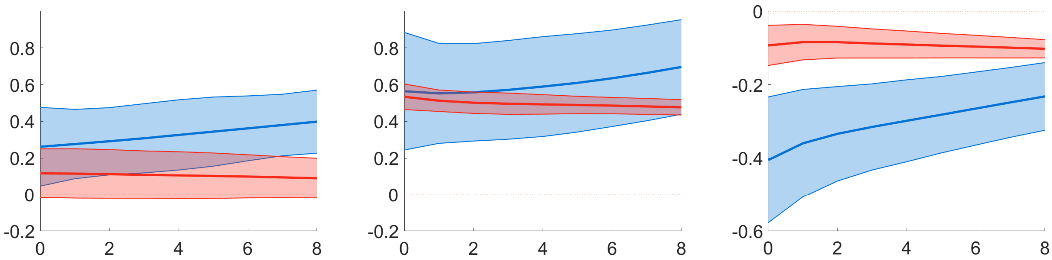

Notes: Phillips multipliers obtained from cumulative regressions over the horizons $h \in [0, ..., 8]$, for sample 1 (blue) and sample 2 (red). Panel A: Unconditional data. Panel B: Conditional on demand shocks only. Panel C: Conditional on supply shocks only. The confidence bands are computed as ± 1.65 x Newey-West robust standard errors.

At this point it should be clear that, in order to quantify the actual, structural Phillips curve slope directly, we would have to assume a specific functional form for the Phillips curve, including the number of lags to include. Here, instead, we compute a variant of the so-called Phillips multiplier, following Barnichon and Mesters (2021). This statistic is a non-parametric characterization of the trade-off between inflation and economic activity. While the Phillips multiplier is reduced-form in nature, it may help us to shed additional light on possible changes in the slope.

Barnichon and Mesters (2021) construct the Phillips multiplier as follows: consider a variable x and its impulse response conditional on a shock (or a set of shocks) ε_t. Let $\mathcal{I}_j^x$ denote the impulse response $j \geq 0$ periods after the shock was realized, and let $\mathcal{I}_h^{\bar{x}} = \frac{1}{h}\sum_{j=0}^{h} \mathcal{I}_j^x$ denote the *average* impulse response at horizon h. The Phillips multiplier is then given by

$$\mathcal{P}_h = \mathcal{I}_h^{\bar{\pi}} / \mathcal{I}_h^{\bar{y}}, \qquad h = 0, 1, 2, ...,$$

where $\mathcal{I}_h^{\bar{\pi}}$ and $\mathcal{I}_h^{\bar{y}}$ represent the average impulse responses of inflation and output, respectively. The authors show that the multiplier can be estimated from the cumulative regression

$$\sum_{j=0}^{h} \pi_{t+j|\varepsilon_t} = \mathcal{P}_h \sum_{j=0}^{h} y_{t+j|\varepsilon_t} + e_{t+h},$$

where $\pi_{t+j|\varepsilon_t}$ and $y_{t+j|\varepsilon_t}$ represent the variation in inflation and output projected by the shock ε_t.

In their application, Barnichon and Mesters (2021) consider a monetary policy shock as an instrument to obtain $\pi_{t+j|\varepsilon_t}$ and $y_{t+j|\varepsilon_t}$. We instead condition on the demand and supply shocks obtained in our baseline SVAR. The cumulative regression equation above is estimated, first, using the unconditional data.[12] Second, we use the data purged for demand shocks. Finally we use the data purged for supply shocks. Estimates of the multiplier are obtained separately for the two sub-samples. Note that we depart from the

[12] This essentially means that we estimate the $\mathcal{P}_h$ in the regression equation $\sum_{j=0}^{h} \pi_{t+j} = \mathcal{P}_h \sum_{j=0}^{h} y_{t+j} + e_{t+h}$.

Table 2: Robustness exercises

	$\hat{\gamma}_u$ S1	$\hat{\gamma}_u$ S2	$\hat{\gamma}_s$ S1	$\hat{\gamma}_s$ S2	$\hat{\gamma}_d$ S1	$\hat{\gamma}_d$ S2	$VD(y\|s)$ S1	$VD(y\|s)$ S2
				Baseline				
(a)	0.26	0.12	0.56	0.53	-0.41	-0.09	0.40	0.61
				Other price and activity measures				
(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
				Different sample periods				
(e)	-0.04	0.19	0.25	0.50	-0.36	-0.04	0.43	0.50
(f)	0.26	-0.01	0.57	0.64	-0.41	-0.14	0.41	0.69
				Controlling for expectations				
(g)	0.26	0.12	0.56	0.39	-0.80	0.06	0.19	0.45
(h)	0.25	0.12	0.50	0.66	-0.46	-0.05	0.19	0.54
(i)	0.44	0.19	0.81	1.02	-0.88	0.04	0.19	0.60
				Additional variables and shocks				
(j)	0.26	0.12	0.40 0.77*	0.38 0.92*	-0.62	-0.07	0.22	0.47
(k)	0.26	0.12	0.39 0.76*	0.35 0.78*	-0.63	-0.11	0.22	0.45
(l)	0.26	0.12	0.38 0.73*	0.41 0.84*	-0.61† -0.43‡	-0.01† -0.22‡	0.37	0.58

*Conditional on monetary policy shocks.
†Conditional on technology shocks.
‡Conditional on labor supply shocks.

procedure in Barnichon and Mesters (2021) by using the SVAR model as a filter to obtain projected data, rather than conditioning on the variation projected by instruments. Thus, we estimate the regression equation above by OLS.

Figure 7 shows the results of this exercise. The unconditional multipliers (Panel A) drop by more than 50% across all horizons, suggesting an empirical disconnect between inflation and output also when we consider averages (or cumulated data) over time. The conditional multipliers in panel B, instead, are much more similar across the two samples. The multiplier conditional on demand shocks declines from 0.7 to 0.5 at horizon 8 in the second sub-sample. Finally, the multiplier conditional on supply shocks increases (becomes less negative) significantly in the second sample across all horizons, consistent with a flattening of the demand curve.

5 ROBUSTNESS

In this section, we evaluate the robustness of our main empirical results from the bi-variate SVAR model. We inspect alternative measures of inflation and economic activity, alternative sample periods, as well as extensions of the SVAR where additional variables and shocks are included. Below we summarize the battery of robustness exercises that we discuss in the main text (in addition to the baseline model presented in the previous section), while an analysis of the posterior distributions is relegated to Appendix C. Results from the robustness exercises are presented in Table 2:

(a) Baseline as a reference

(b) Cyclically sensitive inflation as a measure of inflation (Stock and Watson (2020))

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

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

(e) Sample split in 1998Q4 as in Jørgensen and Lansing (2023)

(f) Second sample ends in 2008Q4

(g) Baseline augmented with SPF expectations

(h) Baseline augmented with Michigan expectations

(i) CPI inflation and Michigan expectations

(j) Three-variable SVAR with the Federal Funds Rate

(k) Three-variable SVAR with the shadow rate as computed by Wu and Xia (2016)

(l) Four-variable SVAR with the shadow rate and the real wage

5.1 ALTERNATIVE INFLATION AND ACTIVITY MEASURES

In specification (b) we estimate the bivariate SVAR using the "cyclically sensitive" (CSI) measure of inflation computed by Stock and Watson (2020). CSI inflation depends little on tradable goods and on poorly measured sectors, but has a high weight on non-tradable goods and services, as well as relatively well-measured sectors. Consistent with Stock and Watson (2020) we find that, unconditionally, the statistical relationship between real economic activity and CSI inflation remains relatively strong in the second sample. Given that this sample presents a period with relatively little volatility, one could perhaps expect to find stable slopes also in the conditional data. In contrast, we confirm a flattening of the projected demand coefficient $\hat{\gamma}_d$ and a larger role of supply shocks in the second part of the sample, as in the baseline model. The combination of stability in the unconditional slope and a flatter demand slope suggests a steepening in the estimated projection coefficient $\hat{\gamma}_s$. In fact, this is what we find: the supply slope $\hat{\gamma}_s$ increases by more than 80%, from 0.43 to 0.79. If anything, this indicates 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.[13] In both

[13]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 com-

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 implies a steepening of the estimated supply slope $\hat{\gamma}_s$, 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 source of the inflation disconnect.[14]

5.2 ALTERNATIVE SAMPLE PERIODS

In specification (e) we consider a sample split in 1998:Q4. We choose this date because, as shown by Jørgensen and Lansing (2023), it results in a negative correlation between the level of inflation and the output gap in the first sample, and a positive correlation in the second sample. Using our data, the slope 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 slope 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 of supply with a flattening of demand to explain the shift in sign in the unconditional slope. Notably, supply shocks explain a larger share of output fluctuations in the second sample, in keeping with our baseline model.

In specification (f) we modify the length of the second sample by ending 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 the propagation of shocks can change substantially when the zero lower bound is binding (although unconventional monetary policies seem to limit the changes in propagation in practice, according to the empirical evidence provided by Debortoli et al. (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 (f): we find a tiny steepening in response to demand shocks, a clear flattening in response to supply shocks, and an increased role for supply shocks in the variance decomposition of output.

5.3 INCLUDING INFLATION EXPECTATIONS

In specifications (g), (h), and (i) 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. In order to match the number of observables with the number of identified shocks, we also include a third disturbance in the system. This shock moves inflation and inflation expectations in opposite directions, see Panel (B) in Table 1. The

parison with our baseline model, we switch the sign of all slopes' coefficients reported for specifications (c) and (d) in Table 2.

[14] A few papers question the validity of the output gap or the unemployment rate as indicators of labor market slack (cf. Ball and Mazumder (2019) and Faccini and Melosi (2023) among others). Alternative measures of labor slack may emerge in the context of more complex models with long-term unemployment, or with search on the job.

third shock, which we do not assign any economic label, plays a minor role in the model (inflation and inflation expectations are positively correlated in the data), perhaps reflecting the fact that 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)).

Specification (g) adds inflation expectations data from the Survey of Professional Forecasters (SPF). We confirm a strong decline in the estimated slope conditional on 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. But we also estimate a non-trivial reduction in the slope conditional on 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. However, such a conclusion is premature. In fact, the expectations relevant for pricing decisions in the context of the Phillips curve are firms' inflation expectations, as discussed in detail in Coibion and Gorodnichenko (2015). Unfortunately, there is no quantitative measure of firms' 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 (h) 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 the supply coefficient $\hat{\gamma}_s$, from 0.5 to 0.66, thus hinting at 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 (i). The results are again stronger than in the baseline model. The estimate for $\hat{\gamma}_d$ (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.

5.4 Extensions with additional shocks

In specification (j) we include the interest rate in the SVAR. Specification (k) includes the shadow rate as computed by Wu and Xia (2016). The latter is tailored to describe interest rates when the nominal policy rate is stuck at the lower bound, a common situation in the later parts of the second sample. These additional variables allow us to identify a monetary policy shock, a second demand shifter which effectively provides us with an additional cross-check of our simple arithmetics. In a way, we are decomposing the baseline demand shock into two components. We assume that the monetary policy shock causes a negative impact co-movement between the interest rate and the inflation rate, see Panel (C) in Table 1. Interestingly, there is no sign of a decline in $\hat{\gamma}_s$ 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 specification (j) and from -0.63 to -0.11 in specification (k). Finally, the role of supply shocks in the variance decomposition of output doubles in both specifications.

Finally, in specification (l) we add a fourth variable (real wages) to the system. Having a larger system also helps to avoid deformation problems, as discussed recently by Canova and Ferroni (2022). The idea is to identify an additional supply-side factor by disentangling labor supply shocks from technology shocks. Panel (D) in Table 1 presents the identification scheme. The slopes conditional on the two demand shocks are almost unchanged over the two samples while we observe a clear flattening in response to both supply shocks, although the change is more pronounced conditional on technology shocks. The share of variance explained jointly by the two supply shocks clearly increases over time.[15]

6 TAKING STOCK

In this section, we review inflation and output dynamics over the last 20 years through the lenses of our empirical results. Finally, we compare our results with selected literature which is particularly related to our paper.

6.1 US INFLATION AND OUTPUT IN THE 21ST CENTURY

Our main result is that the demand curve has flattened over time. A flatter demand curve implies that supply shocks have larger effects on the output gap (relatively to the first sample). Not surprisingly given our previous results on the variance decomposition, this is what we observe when we plot the historical decomposition for the output gap based on our model in Figure 8. In contrast, demand shocks are the main drivers of inflation, in keeping with the idea that the Phillips curve is alive and well. Therefore, a mild dichotomy emerges with inflation being mainly driven by demand shocks and the output gap being mainly driven by supply shocks over the last 20 years.[16]

Some episodes are quite intuitive: for example, our model describes the boom in the second half of the 90s as driven by supply shocks while the pre-Great Recession boom is driven by demand shocks. Perhaps more intriguingly, the model sees the Great Recession

[15] A useful complement to our sensitivity analysis is proposed by our discussant von Schweinitz (2022) who applied our simple arithmetics using the code provided by Baumeister and Hamilton (2018) to estimate a sign-restricted SVAR model driven by demand, supply and monetary policy shocks. When using their data and their priors, von Schweinitz (2022) finds a steepening conditional on demand shocks and a clear flattening in response to supply shocks. He also finds a flattening in response to monetary shocks which, however, have a minor explanatory power in the model.

[16] Angeletos, Collard, and Dellas (2020) find full dichotomy with unemployment (or real economic activity) being driven by a so-called main business cycle shock and inflation by its own shock. Ascari and Fosso (2024) and Bianchi, Nicolò, and Song (2023) extend the analysis in the context of a trend-cycle VAR and find that the main business cycle shock drives also a substantial share of inflation fluctuations at business cycle frequencies.

as driven by a combination of negative supply shocks and negative demand shocks. This does not mean that the demand impulse was small but rather that its propagation was less dramatic than one may think. Once again, the reason is due to a relatively flat demand curve, perhaps reflecting the strong anti-deflationary response of the Federal Reserve at that time. Debortoli et al. (2020) provide empirical evidence in favor of such a view.

Since the fall in output was massive during the Great Recession, the model needs an important role for supply shocks to fit the data during that period. What could these supply shocks be? After all, we are used to the narrative that the Great Recession was predominantly driven by a large and negative demand shock. One possibility is that oil shocks further reinforced the fall in output and mitigated the decline in inflation. 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. A second possibility is that the financial impulse behind the Great Recession (or at least part of it) propagated more like a supply shock than like a demand shock. Gilchrist et al. (2017) provide a theory (and supportive empirical evidence) of why financially constrained firms may have been forced to increase prices during the Great Recession (see also Manea (2020) for complementary theory and empirical evidence). A third possibility is that supply shocks capture in part the increase in wage mark-ups due to downward nominal wage rigidities, as discussed in Galí, Smets, and Wouters (2012).

6.2 Our contribution in perspective

We now discuss our results in connection with a few crucial papers in the literature that have questioned, as we do, the narrative of the Phillips curve flattening.

A natural starting point is the work by McLeay and Tenreyro (2020), who make the case that the Phillips curve can be alive and well even if inflation does not co-move with the output gap in the data. The authors show that optimal monetary policy neutralizes all variation except that due to cost-push shocks (see also Seneca (2018)). In such a scenario, a negative unconditional regression slope emerges automatically. Our identification scheme, in contrast, assumes that monetary policy is represented by a Taylor rule that is unable to mimic optimal monetary policy perfectly. 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 slope in the data. Panel A in Figure 5 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. Our results are fully compatible with their finding on the slope of the Phillips curve. Importantly, we apply the Barnichon and Mesters (2020) recommendation also in using supply shocks to trace the slope of the demand curve. Our focus on the *joint* identification of *both* shocks allows us to exploit fully our simple arithmetics and derive our main results.

Hazell et al. (2022) and Jørgensen and Lansing (2023) 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

Figure 8: A historical shock decomposition of data

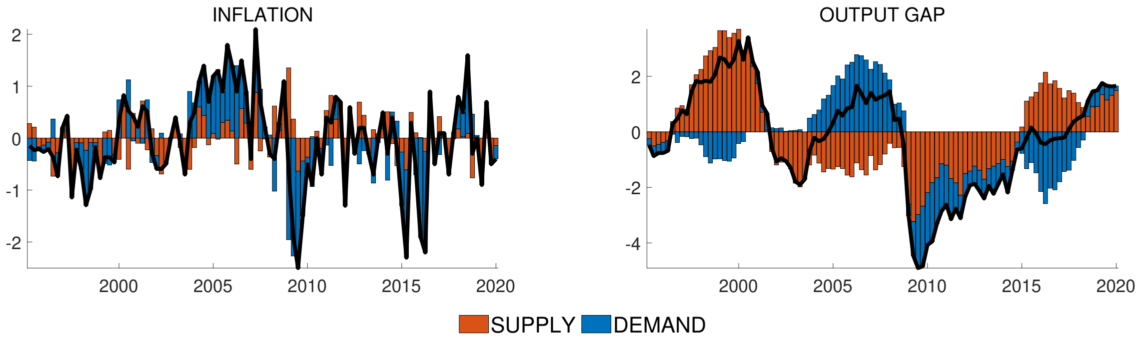

Notes: Historical decompositions of inflation and output gap are presented in deviation from initial conditions.

aggregate data. In our framework, the anchoring of inflation expectations is a by-product of a more aggressive monetary policy response to inflation. Our focus on supply shocks is crucial to providing additional validation of the result obtained in these previous papers.

7 Conclusion

In this paper we have reconsidered the puzzling stability of inflation in spite of 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 regression slope linking inflation to the output gap has declined unconditionally, we find that slopes based on data properly purged for supply shocks have been relatively stable. In contrast, we find substantial support for a flattening of the demand curve recovered by inspecting the propagation of supply shocks over adjacent sample periods. One natural explanation for such a flattening is a more aggressive response of monetary policy to inflation movements in the second part of the sample.

Arguably, one benefit of our paper is its simplicity and the clear mapping between theory and empirics. Nonetheless, it would be interesting to relax some assumptions behind our analysis. Perhaps, the most interesting avenue would be to include non-linearities in our set-up. In fact, while we rely currently on a linearized New Keynesian model and on a linear SVAR, it is conceivable that both the supply curve and the demand curve (perhaps in connection with the zero lower bound on interest rates) may feature non-linearities (cf. Harding, Lindé, and Trabandt (2022) and Harding, Lindé, and Trabandt (2023) for non-linear analysis of the Phillips curve). Disentangling our simple arithmetics from genuine non-linearities in the transmission mechanism of shocks seems of high importance for future research. Unfortunately, the literature has not reached a consensus on how to integrate sign restrictions into non-linear SVAR models, thus making the extension to a non-linear setting far from straightforward.

REFERENCES

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

Albonico, A., G. Ascari, and Q. Haque (2023). The (ir) relevance of rule-of-thumb consumers for us business cycle fluctuations. Journal of Money, Credit and Banking, forthcoming.

Angeletos, G.-M., F. Collard, and H. Dellas (2020). Business-cycle anatomy. *American Economic Review 110*(10), 3030–3070.

Antolín-Díaz, J. and J. F. Rubio-Ramírez (2018). Narrative sign restrictions for SVARs. *American Economic Review 108*(10), 2802–29.

Arias, J. E., J. F. Rubio Ramirez, 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.

Arias, J. E., J. F. Rubio Ramirez, and D. F. Waggoner (2022). Uniform priors for impulse responses. Federal Reserve Bank of Philadelphia Working Paper 22-30.

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

Ascari, G. and L. Fosso (2024). The international dimension of trend inflation. *Journal of International Economics 148*, 103896.

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

Ball, L. and S. Mazumder (2019). A Phillips curve with anchored inflation expectations and short-term unemployment. *Journal of Money, Credit and Banking 51*(1), 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.

Baumeister, C. and J. D. Hamilton (2015). Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica 83*(5), 1963–1999.

Baumeister, C. and J. D. Hamilton (2018). Inference in structural vector autoregressions when the identifying assumptions are not fully believed: Re-evaluating the role of monetary policy in economic fluctuations. *Journal of Monetary Economics 100*, 48–65.

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.

Bergholt, D., F. Canova, F. Furlanetto, N. Maffei-Faccioli, and P. Ulvedal (2024). What drives the recent surge in inflation? the historical decomposition roller coaster. Norges Bank Working Paper 7/2024.

Bianchi, F., G. Nicolò, and D. Song (2023). Inflation and real activity over the business cycle. NBER Working Paper 31075.

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 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.

Canova, F. and F. Ferroni (2022). Mind the gap! Stylized dynamic facts and structural models. *American Economic Journal: Macroeconomics 14*(4), 104–135.

Canova, F. and M. Paustian (2011). Business cycle measurement with some theory. *Journal of Monetary Economics 58*(4), 345–361.

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. NBER Working Paper 27836.

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*.

Faccini, R. and L. Melosi (2023). Job-to-job mobility and inflation. Review of Economics and Statistics, forthcoming.

Fitzgerald, T. J., C. Jones, M. Kulish, and J. P. Nicolini (2020). Is there a stable relationship between unemployment and future inflation? Federal Reserve Bank of Minneapolis Staff Report 614.

Forbes, K. J. (2019). Inflation dynamics: Dead, dormant, or determined abroad? *Brookings Papers on Economic Activity*.

Furlanetto, F., A. Lepetit, Ø. Robstad, J. Rubio Ramírez, and P. Ulvedal (2024). Estimating hysteresis effects. American Economic Journal: Macroeconomics, forthcoming.

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 (2020). Has the us wage phillips curve flattened? a semi-structural exploration. *Castex G, Galí J, Saravia D, editors. Changing Inflation Dynamics, Evolving Monetary Policy. Banco Central de Chile*.

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.

Galí, J., F. Smets, and R. Wouters (2012). Unemployment in an estimated New Keynesian model. *NBER Macroeconomics Annual 26*(1), 329–360.

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. NBER Working Paper 19390.

Harding, M., J. Lindé, and M. Trabandt (2022). Resolving the missing deflation puzzle. *Journal of Monetary Economics 126*, 15–34.

Harding, M., J. Lindé, and M. Trabandt (2023). Understanding post-covid inflation dynamics. *Journal of Monetary Economics 140*, S101–S118.

Hasenzagl, T., F. Pellegrino, L. Reichlin, and G. Ricco (2022). A model of the Fed's view on inflation. *The Review of Economics and Statistics 104*(4), 686–704.

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, and A. Şahin (2022). The missing inflation puzzle: The role of the wage-price pass-through. *Journal of Money, Credit and Banking 54*(S1), 7–51.

Hobijn, B. (2020). The supply-side origins of US inflation. *Castex G, Galí J, Saravia D, editors. Changing Inflation Dynamics, Evolving Monetary Policy. Banco Central de Chile*.

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

Inoue, A. and L. Kilian (2020). The role of the prior in estimating VAR models with sign restrictions. Available at SSRN 3753945.

Inoue, A., B. Rossi, and Y. Wang (2022). Has the Phillips curve flattened and why? Manuscript.

Jørgensen, P. L. and K. J. Lansing (2023). Anchored inflation expectations and the slope of the Phillips curve. Federal Reserve Bank of San Francisco Working Paper 2019-27.

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.

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. Manuscript.

Manea, C. (2020). Monetary policy with financially constrained and unconstrained firms. *Essays in Monetary Economics, PhD Thesis, Universitat Pompeu Fabra*.

Mankiw, G. and R. Reis (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics 117*(4), 1295 – 1328.

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.

Ravenna, F. and C. E. Walsh (2006). Optimal monetary policy with the cost channel. *Journal of Monetary Economics 53*(2), 199–216.

Reis, R. (2009, March). A Sticky-information General Equilibrium Model for Policy Analysis. In K. Schmidt-Hebbel, C. E. Walsh, and N. Loayza (Eds.), *Monetary Policy under Uncertainty and Learning*, Volume 13 of *Central Banking, Analysis, and Economic Policies Book Series*, Chapter 8, pp. 227–283. Central Bank of Chile.

Reis, R. and M. W. Watson (2010). Relative goods' prices, pure inflation, and the Phillips correlation. *American Economic Journal: Macroeconomics 2*(3), 128–157.

Rubbo, E. (2023). Networks, Phillips curves, and monetary policy. *Econometrica 91*(4), 1417–1455.

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

Seneca, M. (2018). A Graphical Illustration of Optimal Monetary Policy in the New Keynesian Framework. Manuscript available at www.seneca.dk.

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

Verona, F. and M. H. Wolters (2013). Sticky Information Models in Dynare. *Computational Economics 43*, 357–370.

von Schweinitz, G. (2022). Discussion of Did the Fed Kill the Phillips Curve? Some Simple Arithmetics. 23rd IWH-CIREQ-GW Macroeconometric Workshop, Halle.

Wong, B. (2013). The evolution of the US output-inflation tradeoff. Technical report, CAMA Working Paper 70/2013.

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

A THE ROLE OF SHOCK PERSISTENCE

Here we discuss the case where shocks to demand and supply are persistent. Persistence implies pass-through from current innovations to agents' forward-looking expectations. In turn, these expectations shift the structural Phillips curve and the aggregate demand curve, summarized by equations (7) and (8) in the main text. Importantly, these shifts come on top of the direct effects from current shocks. The question, then, is how shock persistence and expectations influence our simple arithmetics, as well as what econometric biases we might introduce by ignoring expectations when they actually matter.

To fix ideas, consider the simple model discussed in the main text, summarized by equations (7) and (8). However, suppose that instead of being i.i.d., the structural shocks follow separate AR(1) processes:

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

For simplicity, we assume that demand (supply) shocks share a common demand-specific (supply-specific) autoregressive parameter ρ_d (ρ_s). One can guess and verify the following closed-form solutions for output and inflation:

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

The two auxiliary parameters Ψ_d and Ψ_s are decreasing in ρ_d and ρ_s, respectively:

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

Importantly, $\rho_d > 0$ and $\rho_s > 0$ give rise to a non-zero correlation between current shocks and future realizations of the output gap and inflation, implying an explicit expectations channel when shocks are realized. The quantitative relevance of this channel depends both on shocks' persistence and on the policy coefficients ϕ_π and ϕ_y. Note, also, that the special case with $\rho_d = \rho_s = 0$ brings us back to the expressions stated in the main text.

Next, we state the closed-form expressions for the OLS estimator when shocks are persistent. These objects will depend on whether or not the econometrician accounts properly for the role of expectations. Let us briefly discuss both situations.

A.1 OLS ESTIMATES IF WE IGNORE THE ROLE OF EXPECTATIONS

Suppose that the econometrician ignores expectations and estimate the simple regression equation from the main text:

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

Then $\gamma \equiv \frac{cov(\pi_t, y_t)}{var(y_t)}$ and we obtain the following closed-form expressions for the estimated projection coefficients $\hat{\gamma}_u$, $\hat{\gamma}_s$ and $\hat{\gamma}_d$:

$$\hat{\gamma}_u = \frac{\frac{\kappa(1-\beta\rho_d)}{\Psi_d^2} - \frac{(\phi_\pi-\rho_s)[\sigma(1-\rho_s)+\phi_y]}{\Psi_s^2}\frac{\sigma_s^2}{\sigma_d^2}}{\left(\frac{1-\beta\rho_d}{\Psi_d}\right)^2 + \left(\frac{\phi_\pi-\rho_s}{\Psi_s}\right)^2 \frac{\sigma_s^2}{\sigma_d^2}} \tag{A.1}$$

$$\hat{\gamma}_s = \frac{\kappa}{1-\beta\rho_d} \tag{A.2}$$

$$\hat{\gamma}_d = -\frac{\sigma(1-\rho_s)+\phi_y}{\phi_\pi - \rho_s} \tag{A.3}$$

Here we have introduced the normalizations $\sigma_d^2 = \frac{\sigma_{\varepsilon,d}^2}{1-\rho_d^2}$ and $\sigma_s^2 = \frac{\sigma_{\varepsilon,s}^2}{1-\rho_s^2}$. A few remarks are in place: first, the projection coefficient based on unconditional data, $\hat{\gamma}_u$, remains a function of both the structural Phillips curve slope, the stance of monetary policy, and the volatility of supply disturbances relative to demand disturbances. Thus, just as in the main text, an observed decline in $\hat{\gamma}_u$ over time does not necessarily imply that the structural Phillips curve has flattened and we need more information to identify underlying causes.

Second, $\hat{\gamma}_s$ is increasing in κ while $\hat{\gamma}_d$ is independent of κ. At the same time, $\hat{\gamma}_d$ is increasing in ϕ_π (i.e. becoming less negative) while $\hat{\gamma}_s$ is independent of ϕ_π. Importantly, these observations justify the simple arithmetics presented in the main text: (i) a decline in both $\hat{\gamma}_u$ and $\hat{\gamma}_s$, combined with a stable estimate of $\hat{\gamma}_d$, is what we expect to see if the structural Phillips becomes flatter. (ii) a decline in $\hat{\gamma}_u$ and a flattening of $\hat{\gamma}_d$ (towards less negative values), coupled with a stable estimate of $\hat{\gamma}_s$, would instead point to a flattening of the demand curve. (iii) if $\hat{\gamma}_u$ declines while $\hat{\gamma}_s$ and $\hat{\gamma}_d$ remain relatively stable, then it seems likely that supply shocks have become more volatile relative to demand shocks.

Finally we note that $\hat{\gamma}_s$, if interpreted as a structural Phillips curve slope, is biased upwards relative to κ when shocks are persistent. This bias is strictly increasing in ρ_d. Intuitively, κ captures the elasticity of π_t with respect to y_t *keeping expectations fixed.* But expectations are not fixed when $\rho_d > 0$. If the econometrician ignores expectations and attribute the entire rise in π_t directly to the contemporaneous rise in y_t, then the included regressor y_t will be positively correlated with the error term ε_t. As a result one obtains a too high estimate of κ. Put differently, the estimates of $\hat{\gamma}_s$ presented in the main text are only proportional to κ and should, if anything, be interpreted as potential upper limits of the structural Phillips curve slope.

A.2 OLS ESTIMATES ACCOUNTING FOR EXPECTATIONS

Finally we consider a case where expectations are taken explicitly into account. The regression equation which is estimated follows below:

$$\pi_t = \beta\pi_{t+1}^e + \gamma y_t + \varepsilon_t$$

The variable π_{t+1}^e represents inflation expectations. An econometrician may obtain variation in inflation expectations from instruments or observe them directly. Given our focus on the simple arithmetics associated with conditional data, we will assume that expectations are measured properly.[17] Now the OLS estimator is given by $\gamma \equiv \frac{cov\left(\pi_t - \beta\pi^e_{t+1}, y_t\right)}{var(y_t)}$, and the three datasets considered result in the following closed-form estimates of projection slopes:

$$\hat{\gamma}_u = \frac{\kappa\left(\frac{1-\beta\rho_d}{\Psi_d}\right)^2 - \frac{(\phi_\pi-\rho_s)[\sigma(1-\rho_s)+\phi_y](1-\beta\rho_s)}{\Psi_s^2}\frac{\sigma_s^2}{\sigma_d^2}}{\left(\frac{1-\beta\rho_d}{\Psi_d}\right)^2 + \left(\frac{\phi_\pi-\rho_s}{\Psi_s}\right)^2\frac{\sigma_s^2}{\sigma_d^2}} \tag{A.4}$$

$$\hat{\gamma}_s = \kappa \tag{A.5}$$

$$\hat{\gamma}_d = -\frac{\left[\sigma\left(1-\rho_s\right)+\phi_y\right]\left(1-\beta\rho_s\right)}{\phi_\pi - \rho_s} \tag{A.6}$$

Our simple arithmetics derived earlier remain valid even in this case: (i) the slope story is associated with a decline in $\hat{\gamma}_u$ and $\hat{\gamma}_s$, but not in $\hat{\gamma}_d$. (ii) the policy story implies a decline in $\hat{\gamma}_u$ and a rise (towards less negative values) in $\hat{\gamma}_d$ while $\hat{\gamma}_s$ remains stable. (iii) the shocks story implies a decline in $\hat{\gamma}_u$, combined with unchanged $\hat{\gamma}_s$ and $\hat{\gamma}_d$. Finally, given that expectations are dealt with properly we also obtain $\hat{\gamma}_s$ as an unbiased estimate of the structural Phillips curve slope κ when data are purged for supply-side shocks.

B PUBLIC SECTOR DEMAND AND JOINT SHIFTS IN DEMAND AND SUPPLY CURVES

This exercise illustrates a situation in which demand and supply curves shift simultaneously in response to the same shock. This situation may arise if the output gap seizes to be proportional to marginal costs, an assumption we used in the main text. The purpose is to highlight implications for our simple arithmetics, and especially what kind of biases one might expect. To this end we introduce fiscal policy, although one can think of many other extensions with similar implications, e.g. sticky wages, the use of intermediate inputs in production, etc.

We model a public sector in the simplest way possible: suppose that the public sector accounts for a fraction $1-\alpha$ of the aggregate economy in steady state. Market clearing in the goods market is given by $y_t = \alpha c_t + (1-\alpha)\, g_t$, where c_t represents private consumption and g_t represents public sector consumption. For simplicity, we abstract from the potential role of public debt dynamics and suppose that public demand follows a simple rule; $g_t = \tau y_t + \varepsilon_{g,t}$. The parameter τ determines the cyclical stance of fiscal policy. We refer to the policy rule as counter-cyclical (pro-cyclical) if $\tau < 1$ (> 1), since this implies that the public sector share $g_t - y_t$ is higher (lower) in recessions than in expansions. $\varepsilon_{g,t}$ is a discretionary fiscal policy shock. Combining the fiscal rule with the market clearing condition just stated we note that

$$c_t = \xi_\tau y_t - \frac{1-\alpha}{\alpha}\varepsilon_{g,t},$$

where $\xi_\tau = 1 + \frac{(1-\alpha)(1-\tau)}{\alpha}$. Our setup in the main text, $c_t = y_t$, emerges as a special case when $\alpha = 1$. The rest of the model is as before (see equations (1)-(6) in the main text).

[17]This means that $\pi^e_{t+1} = \mathbb{E}_t\pi_{t+1}$ when expectations are rational and in the absence of further information frictions.

Following the same steps as earlier we once again arrive at an upward-sloping supply curve in the (y_t, π_t)-space, as well as a downward-sloping demand curve:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa_\tau y_t + s_t - \lambda x_{g,t} \tag{B.1}$$

$$y_t = \frac{1}{\sigma_\tau + \phi_y} \left(\sigma_\tau \mathbb{E}_t y_{t+1} - \phi_\pi \pi_t + \mathbb{E}_t \pi_{t+1} + d_t + x_{g,t} \right) \tag{B.2}$$

The auxiliary parameters σ_τ and κ_τ are given by $\sigma_\tau = \sigma \xi_\tau$ and $\kappa_\tau = \lambda (\sigma_\tau + \varphi)$ respectively, while $x_{g,t} = \sigma \frac{1-\alpha}{\alpha} \varepsilon_{g,t}$ captures the pass-through from fiscal shocks. d_t and s_t are defined as in the main text ($d_t = u_t - m_t$ and $s_t = z_t + \lambda \psi_t - \lambda (1 + \varphi) a_t$).

The introduction of a public sector changes the model in two ways: first, the cyclicality of fiscal policy affects both demand and supply curve slopes. A more countercyclical policy rule raises σ_τ (and, therefore, κ_τ), implying a flatter demand curve and a steeper supply curve. Second, both curves shift out in response to an expansionary fiscal shock: the aggregate demand curve shifts outwards because increased public demand, ceteris paribus, stimulates both output and inflation. But increased public demand also operates through income effects on households' labor supply, causing an outward shift in the aggregate supply curve as well. This further stimulates output but reduces inflation. The net effect on inflation is, therefore, ambiguous. Moreover, the joint shift in both curves presents an identification challenge for the econometrician trying to identify demand and supply curve slopes in the data.

To better understand this identification challenge, we bring back the assumption that all shocks are white noise and derive analytical solutions for output and inflation:

$$y_t = \frac{1}{\sigma_\tau + \phi_y + \kappa_\tau \phi_\pi} \left[d_t - \phi_\pi s_t + (1 + \phi_\pi \lambda) x_{g,t} \right]$$

$$\pi_t = \frac{1}{\sigma_\tau + \phi_y + \kappa_\tau \phi_\pi} \left[\kappa_\tau d_t + (\sigma_\tau + \phi_y) s_t + \left(1 - \frac{\sigma_\tau + \phi_y}{\sigma_\tau + \varphi} \right) \kappa_\tau x_{g,t} \right]$$

An exogenous public sector expansion is inflationary if $\varphi > \phi_y$, i.e. if income effects on labor supply are not too large. We suppose that this is the most relevant case empirically.

One can still derive closed-form expression for the OLS estimator $\gamma \equiv \frac{cov(\pi_t, y_t)}{var(y_t)}$, associated with the simple regression specification $\pi_t = \gamma y_t + \varepsilon_t$. The projection coefficient based on unconditional data admits the following closed-form solution:

$$\hat{\gamma}_u = \frac{\kappa_\tau \sigma_d^2 - \phi_\pi (\sigma_\tau + \phi_y) \sigma_s^2 + (1 + \phi_\pi \lambda) \left(1 - \frac{\sigma_\tau + \phi_y}{\sigma_\tau + \varphi} \right) \kappa_\tau \sigma_g^2}{\sigma_d^2 + \phi_\pi^2 \sigma_s^2 + (1 + \phi_\pi \lambda)^2 \sigma_g^2} \tag{B.3}$$

The volatility of fiscal shocks is captured by $\sigma_g^2 = \left(\sigma \frac{1-\alpha}{\alpha} \right)^2 \sigma_{\varepsilon_g}^2$ while σ_s^2 and σ_d^2 are defined in the main text ($\sigma_d^2 = \sigma_m^2 + \sigma_u^2$ and $\sigma_s^2 = \sigma_z^2 + \lambda^2 \sigma_\psi^2 + \lambda^2 (1 + \varphi)^2 \sigma_a^2$ respectively).

It is clear that fiscal shocks further limit the scope for $\hat{\gamma}_u$ to represent an accurate estimate of κ_τ. Moreover, the implications for an econometrician hoping to exploit conditional variation depend critically on the correlation structure induced by $\varepsilon_{g,t}$. Here we suppose that fiscal shocks generate co-movement between y_t and π_t (as implied by $\varphi > \phi_y$) so that our econometrician, who follows a sign restriction approach, treats them as conventional demand shifters. The estimated projection coefficient conditional on supply shocks only is naturally unchanged in this case:

$$\hat{\gamma}_d = -\frac{\sigma_\tau + \phi_y}{\phi_\pi} \tag{B.4}$$

Thus, stricter inflation targeting still flattens the demand curve while changes in the Phillips curve slope play no role for the slope of demand.

Turning to the supply curve, i.e. the estimated projection coefficient conditional on *perceived* demand shocks, we define $\hat{\gamma}_{\tilde{s}}$ to highlight the econometrician's inclusion of fiscal shocks in the vector of demand disturbances. The closed-form solution when fiscal shocks are present can be derived as follows:

$$\hat{\gamma}_{\tilde{s}} = \left[1 - \frac{(1+\phi_\pi\lambda)\left(\phi_\pi\lambda + \frac{\sigma_\tau+\phi_y}{\sigma_\tau+\varphi}\right)\frac{\sigma_g^2}{\sigma_d^2}}{1+(1+\phi_\pi\lambda)^2\frac{\sigma_g^2}{\sigma_d^2}}\right]\kappa_\tau < \kappa_\tau = \hat{\gamma}_s \tag{B.5}$$

It is clear that $\hat{\gamma}_{\tilde{s}}$ (i) tends to be biased downward relative to the Phillips curve slope κ_τ, and (ii) may change because of shifts in ϕ_π as well. In fact, suppose that we condition on fiscal shocks only:

$$\hat{\gamma}_g = \frac{1 - \frac{\sigma_\tau+\phi_y}{\sigma_\tau+\varphi}}{1+\phi_\pi\lambda}\kappa_\tau \tag{B.6}$$

Both a flatter Phillips curve and stricter inflation targeting lead to a decline in $\hat{\gamma}_g$. It is only when labor supply becomes infinitely inelastic (φ goes to infinity), and at the same time the Phillips curve becomes infinitely flat (λ goes to zero), that $\hat{\gamma}_g$ approaches κ_τ. This highlights the danger of using fiscal shocks (or shocks with similar nature) as instruments when trying to estimate the structural Phillips curve slope.

Nevertheless, while the fiscal shock at least in principle may represent a potential threat to our identification strategy, in practice we do not believe this to be a quantitatively relevant issue. First, it seems implausible that fiscal policy shocks are responsible for a dominant share of the business cycle volatility in output and inflation. Second, recall that in section 4 we find essentially no decline in $\hat{\gamma}_s$, but a substantial flattening of $\hat{\gamma}_d$, when inspecting actual data. Given the discussion here, these findings seem consistent with stricter inflation targeting even if fiscal policy shocks would turn out to be quantitatively important. A flattening of the Phillips curve, instead, should have led to a flattening of $\hat{\gamma}_s$ combined with a stable $\hat{\gamma}_d$, i.e. the opposite of what we find in the data.

C Inspecting the posterior distribution

Following common practice, all of the results so far rely on one historical decomposition based on the median contribution of each shock at each quarter in the SVAR model. As an additional robustness exercise, in this appendix we instead evaluate the conditional regression slopes across a wide range of the posterior distribution, thus taking into account the model uncertainty which is an inherent feature of set-identified models. Recall that the Bayesian approach based on Arias et al. (2018) provides us with $10,000$ different

Table 3: Robustness to posterior credibility sets

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

posterior models for each of the two samples under consideration. All 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 $\hat{\gamma}_s$ and $\hat{\gamma}_d$. This is done both for the baseline specification and 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 $\hat{\gamma}_s$, we confirm the general picture established earlier. Projected supply slopes (i.e. conditional on demand shocks) do not tend to decrease, at least not to a large extent. The 68% credible set for $\hat{\gamma}_s$ in sample 2 spans its posterior mean in sample 1 in all specifications except (e) and (f). Specification (e) uses the same sample split as in Jørgensen and Lansing (2023), implying in fact a steepening of the slope–both unconditionally and conditional on demand. Specification (f), instead, is where we discard observations after 2008:Q4. This is the only significant exception pointing to a flattening of the Phillips curve. Second, regarding $\hat{\gamma}_d$, we find a major flattening (i.e. a less negative value) 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 $\hat{\gamma}_d$ in sample 2 does not span the posterior mean in sample 1 in any of the specifications except (h), where we add inflation expectations from the Michigan survey to the analysis. Finally, it seems that the biggest change in the posterior distribution of $\hat{\gamma}_d$ 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 support 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.

While our results survive broadly after the inspection of the posterior distribution of admissible models, it is fair to admit that model uncertainty (summarized by the credible intervals) is large. This is an inherent feature of SVAR models identified with sign restrictions, and the price to pay when using only a few identification assumptions. These credible intervals could be shrunk significantly by imposing the sign restrictions over a longer horizon (and not only on impact) and by imposing a few narrative restrictions (as in Antolín-Díaz and Rubio-Ramírez (2018)) in addition to the impact sign restrictions. In the context of the simplicity that characterizes both the theoretical and empirical part of this paper, our goal is to show that some solid results can be obtained even when imposing just a couple of impact sign restrictions (cf. Canova and Paustian (2011)) while admitting that these results could be refined by imposing more structure on the empirical model.