The Long-Run Phillips Curve is... a Curve

Guido Ascari, Paolo Bonomolo and Qazi Haque

DeNederlandscheBank

EUROSYSTEEM

The Long-Run Phillips Curve is… a Curve

Guido Ascari, Paolo Bonomolo and Qazi Haque*

The Long-Run Phillips Curve is ... a Curve[*]

Guido Ascari [†]
De Nederlandsche Bank
University of Pavia

Paolo Bonomolo[‡]
De Nederlandsche Bank

Qazi Haque [§]
The University of Adelaide
Centre for Applied Macroeconomic Analysis

August 2, 2023

Abstract

In U.S. data, inflation and output are negatively related in the long run. A Bayesian VAR with stochastic trends generalized to be piecewise linear provides robust reduced-form evidence in favor of a threshold level of trend inflation of around 4%, below which potential output is independent of trend inflation, and above which, instead, potential output is negatively affected by trend inflation. Moreover, this negative relationship is quite substantial: above the threshold every percentage point increase in trend inflation is related to about 1% decrease in potential output per year. A New Keynesian model generalized to admit time-varying trend inflation and estimated via particle filtering provides theoretical foundations to this reduced-form evidence. The structural long-run Phillips Curve implied by the estimated New Keynesian model is not statistically different from the one implied by the reduced-form piecewise linear BVAR model.

JEL Classification Numbers: C32, C51, E30, E31, E52
Keywords: Long-Run Phillips Curve, Inflation, Bayesian VAR, DSGE; Particle Filter.

---

1   Introduction

Is there a tradeoff between inflation and output in the long-run?

The answer is no, according to undergraduate macroeconomics textbooks. These textbooks explain that while there is a short-run tradeoff between inflation and output (or the unemployment rate), this tradeoff disappears in the long run, so that the long-run Phillips curve is vertical at the natural level of output (or the natural rate of unemployment). The long-run relationship between inflation and economic activity, dubbed the long-run Phillips curve (LRPC), can shift if real forces shift this natural level, but inflation and monetary factors do not affect the LRPC, so that inflation and real economic activity are unrelated in the long-run. From the seminal works of Friedman (1968) and Phelps (1967) onwards, the idea that *"inflation is a monetary phenomenon"* is a central tenet of macroeconomic theory and of the inflation targeting monetary policy strategy of most western central banks.

The relationship between inflation and economic activity is therefore of paramount importance for monetary policymaking as most central banks, including the Federal Reserve and the European Central Bank, perceive price stability as the basis for long-term economic growth. Inflation, both headline and core, is on the rise in the world economy at large. Inflation in the U.S. is at at 40-year highs, sparking a debate about whether high inflation is on the way back after years of playing dead. If inflationary pressures turn out to be more permanent, this may lead to higher underlying trend inflation. What would be the impact of higher trend inflation on real economic activity in the long-run?

While considerable effort has been devoted in the economics literature to investigate this relationship in the short-run, it might be very surprising to realize that (see the related literature paragraph below): (i) little econometric work has been devoted to estimating the LRPC, and (ii) many theoretical frameworks would not imply a vertical LRPC.[1] This paper tackles both issues by first trying to establish the nature of the long-run relationship between inflation and economic activity in U.S. data and testing the natural rate hypothesis. Then, it provides a possible theoretical structural interpretation of the empirical finding.

Regarding (i), a first contribution of the paper is to develop a new empirical framework to investigate the existence of a potential non-linear relationship between inflation and output in the long-run. The framework generalizes the Bayesian VAR with stochastic trends (see Del Negro et al., 2017; Johannsen and Mertens, 2021) to a piecewise linear case. From a

---

[1]The property of the absence of a long-run relationship between the rate of growth of money - which is equal to the rate of inflation in steady state - and real variables is often referred to as *superneutrality of money.* Macroeconomic models with optimizing agents generally do not satisfy this property, where by 'generally' we mean 'without some further assumptions'. Already Sidrauski (1967) showed that money is *not* superneutral in a Ramsey model of growth with money in the utility function. The same applies to a real business cycle with a cash-in-advance constraint, see Cooley and Hansen (1989), or to almost any model with a shopping time technology, as well as to overlapping generations monetary models, see, e.g., Drazen (1981), just to give some basic examples.

methodological point of view, the functional form of the piecewise linear model depends on the latent processes, which in our case is trend inflation. Our theoretical contribution is to show that both the likelihood function and the posterior distribution of the latent states can be derived analytically. Therefore, in terms of efficiency our estimator is comparable to linear models. More importantly, the piecewise linear framework allows us to test the idea that the long-run relationship between inflation and output can change nature depending on the level of trend inflation. In other words, we think about the piecewise linear model as an efficient way to approximate a possible underlying non-linear relation. The main result is the evidence in favor of a threshold level of trend inflation below which potential output is independent of trend inflation, and above which, instead, potential output is *negatively* affected by trend inflation. This threshold level of inflation is around 4%. Moreover, the slope of the LRPC is sizeable. Every percentage point increase in inflation above the threshold level of inflation is related to about 1% decrease in potential output per year. Given these findings, we define a new concept: *"the long-run output gap"*, which is defined as the deviation of potential output under positive trend inflation from its counterfactual level under zero trend inflation. The long-run output gap has been about negative 2% per year on average during the Great Inflation, implying sizeable output costs. Related to this point, we also discuss the implications of a negatively-sloped LRPC for the measurement of business cycles. Specifically, neglecting the long-run relationship between inflation and output leads to more negative short-run output gap estimates in periods of high inflation, particularly during the Great Inflation, thereby overstating the cyclical component of output fluctuations.

Regarding (ii), we then look for a possible theoretical interpretation of this empirical reduced-form result. It is natural to start by asking whether the most standard workhorse New Keynesian (NK) framework can quantitatively reproduce the LRPC estimates of the BVAR. The canonical NK model does imply a non-linear LRPC (see Ascari and Sbordone, 2014) because positive trend inflation creates inefficient price dispersion due to nominal rigidities and hence reduces the natural level of output.[2] The relevance of this non-linearity and the magnitude of the negative effect depend on the parameters of the model. Then, the question becomes empirical. Moreover, to verify the extent to which the New Keynesian model can reproduce the main features of the long-run relationship between inflation and output found in the BVAR analysis, we need to extend the model by allowing for time variation in steady state inflation. Allowing trend inflation to vary every period is a non-trivial modification of the baseline model, both because the steady state of the model becomes time-varying, and because the dynamics of the model is affected non-linearly by the level of trend inflation. This paper, thus, generalizes to a full NK model the work in Cogley and Sbordone (2008), who estimate the New Keynesian

---

[2]We use the terms 'natural level of output', 'potential output', 'steady state output' as indicating the same object: the long-run level of output. Furthermore, we will use as synonymous the terms 'trend inflation', 'inflation target' and 'steady state inflation'.

Phillips Curve (NKPC) allowing for time variation in trend inflation, and thus in the NKPC coefficients.[3] A second contribution of the paper is, therefore, to estimate the structural NK model generalized by adding time-varying trend inflation and stochastic volatility. We develop an econometric strategy suited for this problem, allowing us to jointly estimate the short-run dynamics and the long-run relationship implied by the model. The model parameters and the latent states are estimated using a Bayesian approach based on Sequential Monte Carlo methods. In particular, we use the econometric strategy for parameter learning that combines the approach of Carvalho et al. (2010), and the particle filter of Liu and West (2001), as in Ascari et al. (2019).[4]

The estimated Generalized New Keynesian (GNK) model reproduces very well the evidence of the reduced-form BVAR model of a negative long-run relationship between inflation and output. The LRPC is not vertical but negatively sloped and non-linear. In particular, it is vertical for very low levels of inflation and then it exhibits an increasingly negative slope as the long-run inflation rate increases above 3-4%. In terms of output losses, going from 2% to 4% inflation target causes an output loss of roughly about 0.6% per year. The effect is highly non-linear such that a 5% and a 6% inflation target would imply an output loss (relative to 2% target) of roughly 1.2% and 2% per year, respectively. The estimates are quite precise and they are not statistically different from the one implied by the reduced-form piecewise linear BVAR model, i.e., the estimated structural LRPC is within the credibility bands of the estimated long-run relation between trend inflation and potential output from the BVAR. In addition, the long-run output gap estimate from the structural model is quantitatively similar to the one from the BVAR, with output cost estimates of about $1-3\%$ per year during the Great Inflation. From a medium to long-run perspective, these numbers are not negligible, even for low levels of trend inflation if one looks at the cumulative losses over the years.

Related Literature. The famous correlation unveiled by Phillips (1958) was initially thought to imply a long-run negative tradeoff between (wage) inflation and unemployment (Phillips, 1958; Samuelson and Solow, 1960). As is well-known, the idea of a long-run tradeoff disappeared with the seminal papers by Friedman (1968) and Phelps (1967) that introduce the keystone concept of a natural rate of unemployment and a vertical LRPC. Early tests of the natural rate hypothesis (NRH) (e.g., Sargan, 1964; Solow, 1969; Gordon, 1970) were based on

---

[3]Cogley and Sbordone (2008) structurally decompose inflation dynamics into a time-varying long-run component (i.e., trend inflation) and a short-run one (i.e., the inflation gap given by the difference between inflation and trend inflation). Their main finding is that time-varying trend inflation captures the low frequency variation in the dynamics of inflation, while the short-run inflation gap fits well a purely forward-looking NKPC without the need of any *ad hoc* intrinsic inertia.

[4]Fernández-Villaverde and Rubio-Ramírez (2007) present pioneering work on the estimation of non-linear or non-Gaussian DSGE models, based on particle filtering within a Markov Chain Monte Carlo scheme. The use of Sequential Monte Carlo methods is less common in the literature. Exceptions are Creal (2007), Chen et al. (2010) and Herbst and Schorfheide (2014).

estimating a Phillips Curve using some distributed lags of inflation to capture expectations and then look at whether the sum on the inflation coefficients would add up to one.[5] After these early times, the literature on testing the natural rate hypothesis is surprisingly slim, compare to its pivotal role in macroeconomics. King and Watson's (1994) influential paper find the inflation and unemployment series to be $I(1)$ but no evidence of cointegration between them. Karanassou et al. (2005) is one of the first papers to cast doubt about the NRH.[6] Beyer and Farmer (2007) cannot reject the assumption of $I(1)$ for inflation and unemployment, but, unlike King and Watson (1994), they find that the low frequency comovements are stable and cointegrated across the whole sample. They interpret their finding as evidence against the NRH. Even more surprisingly, they find that the cointegrating vector in their VECM model implies a positive long-run relationship between inflation and unemployment, contrary to the famous Phillips (1958) negative correlation. Berentsen et al. (2011) report a positive correlation between the low frequency (filtered) component of inflation and unemployment. Haug and King (2014) corroborate this suggestive evidence using more advanced time-series methods for filtering. A recent paper by Ait Lahcen et al. (2021) uses cross-country panel data from the OECD countries to document that the positive correlation between long-run anticipated inflation and unemployment is state-dependent, i.e., it is higher when unemployment is higher. This is consistent with our findings. Benati (2015) conducts SVAR analysis for several advanced economies including the U.S. and concludes that there is no evidence in favour of a non-vertical LRPC. However, the uncertainty surrounding the estimates is so large that is not possible to reject an alternative view, where he meant a negative relationship.

We add to this literature in many dimensions. First, we employ a different methodology based on the BVAR analysis with stochastic trends, thereby providing a multivariate trend-cycle decomposition. Second, we provide a methodological contribution as we generalize this approach to a non-linear setting. While the non-linear approach is necessary to identify a threshold value of trend inflation that tilts the long-run relationship between inflation and output, it is also justified by the difficulties in estimating this relationship, as flagged by Beyer and Farmer (2007) and Benati (2015). Beyer and Farmer (2007) estimate the model over two different subsamples because of parameter shifts. Benati (2015) discusses the difficulties in identifying this long-run relationship because of changing inflation dynamics due to different monetary policy regimes (Benati, 2008). The possibility of identifying the LRPC depends on the inflation process displaying permanent variations, i.e., a unit root. However, inflation persistence changed quite dramatically during the post-WWII sample in the U.S. data, and the Great Inflation might be

---

[5]See King and Watson (1994) and King (2008) for a comprehensive survey of the history of the debate over the nature of the Phillips curve in macroeconomic history in the '70s and'80s. See Karanassou and Sala (2010) and Svensson (2015) for a very recent investigation using a similar approach.

[6]Karanassou and Sala have a series of papers investigating the NRH for various countries and using different methods - GMM, VAR and chain-reaction theory (CRT) - see Karanassou and Sala (2010) and Karanassou et al. (2010) for a survey of these works.

the only period that allows identification of the LRPC. Finally, we also estimate a structural model providing theoretical underpinnings to the empirical analysis.

Regarding theory, first, it is well-known that the GNK model delivers a negative relationship between steady state inflation and output (Ascari, 2004; Ascari and Sbordone, 2014). Hence, it is natural to work with the workhorse NK model which is at the core of the modern analysis of business cycle and monetary policy.[7] Second, given the complexity of the estimation procedure due to our assumption of time-varying trend inflation, we estimate a relatively small-scale version of this model with flexible wages and no role for capital. Third, recently, some papers have used micro data to try to infer the welfare or output costs of misallocation due to nominal rigidities. In a dataset on pricing behavior during the Great Inflation period, Nakamura et al. (2018) find no evidence of larger absolute price changes and evidence of an increase in the frequency of price changes, suggesting that state-dependent sticky price models might be a more plausible mechanism to describe pricing frictions. Nakamura et al. (2018) further show that the positive relationship between inflation and price dispersion is very weak in their state-dependent pricing model.[8] Sheremirov (2020) shows that microdata exhibit a positive comovement between inflation and the dispersion of regular prices - that is, excluding temporary sales - and that the Calvo model overstates this comovement, while the standard fixed menu cost model understates it. Moreover, Sheremirov (2020) suggests that a Calvo model with sales is the only one able to replicate the relation between inflation and price dispersion in the microdata. Importantly for us, he shows that: (i) the inflation cost of business cycles is 40% higher in his favourite model compared to the standard Calvo model, leading to a lower optimal inflation rate; (ii) the shape of the output response to monetary policy shock in the Calvo model with sales is similar to the standard one without sales, suggesting that the implied short-run dynamics of the Calvo model is a good approximation for aggregate variables. Adam et al. (2023) assume that efficient prices follow (product-specific) trend inflation and use this assumption to identify changes in the inefficient price dispersion in the UK data. Contrary to Nakamura et al. (2018), they show that, at the aggregate level, fluctuations in inefficient price dispersion are sizable and covary positively with aggregate inflation. The fact that suboptimally high (or low) inflation is associated with distortions in relative prices that are quantitatively large provide empirical support for the mechanism in standard sticky price models, as the one in our paper. A recent paper by Cavallo et al. (2021) estimates the level of the cost of misallocation using granular price level data for the Euro Area and a state dependent menu cost model. They provide evidence for sizeable costs: in the low-inflation environment that prevailed before 2022, the efficiency

---

[7]Berentsen et al. (2011) use an alternative approach based on search-and-matching frictions both in the goods and labor markets to explain the positive correlation between long-run anticipated inflation and unemployment. Ait Lahcen et al. (2021) build on this model to explain the non-linearity in this relationship they find in the OECD data. None of these papers estimate the model.

[8]The flat relationship between price dispersion and inflation in menu costs model heavily depends on the fact that the model needs large idiosyncratic shocks to fit the microdata.

cost amounts to roughly 2% of GDP. From a modelling perspective, the literature recognizes that the cost of steady state inflation is higher in the standard Calvo model compared to alternative sticky price models. However, these recent works just mentioned above, provide evidence of much more sizeable costs of price dispersion than previously thought using detailed and granular micro data, and thus more similar to the one implied by the standard NK model.

Most importantly, we are mainly concerned here with the steady state relationship, not the short-run costs. In this regards, our results show that our estimated GNK model is able to reproduce the LRPC estimated from the reduced-form BVAR analysis. Therefore, it is able to capture the long-run tradeoff between inflation and output in the aggregate data, despite not capturing the richness of the microdata behaviour. Estimating a LRPC in a DSGE model with state-dependent prices and time-varying trend inflation is computationally challenging, if not infeasible. Future research might tell whether a menu cost model is able to match the reduced form empirical evidence in aggregate data.

Finally, closer to our approach, a recent paper by Abbritti et al. (2021) augment a standard Calvo-type New Keynesian (NK) framework with endogenous growth, frictional labor market and downward wage rigidity. The calibrated model yields a long-run trade-off between output growth and inflation, and consumption equivalent welfare losses of deviation from the optimal inflation target that are a multiple of those associated with traditional models, because endogenous growth magnifies the trade-off between price distortions and output hysteresis.

The paper proceeds as follows. The next section presents the reduced form BVAR methodology along with the estimated long-run Phillips curve. The section introduces the notion of the long-run output gap and shows its estimates from the BVAR and also discusses the implications for business cycle measurement arising from a non-linear LRPC. Section 3 presents the structural GNK model, the estimation methodology and the estimation results. The section documents that a canonical NK model with time-varying trend inflation implies an estimated LRPC that is both qualitatively and quantitatively in line with the BVAR analysis. Finally, Section 4 concludes.

2   A time series approach

We propose a time series model that is tailored to the purpose of estimating the long-run Phillips curve. As in Del Negro et al. (2017) and Johannsen and Mertens (2021), we express a VAR in deviations from time varying trends that we interpret as the long-run components of the respective variables. The methodology is a generalization of the steady state VAR by Villani (2009) and is a trend-cycle decomposition where the dynamics of the cyclical components are described by an unrestricted VAR, but the long-run trends have a structure inspired by economic

theory.[9]

More formally, indicate with X_t a $n \times 1$ vector of observed variables at time t. We define $\bar{X}_t$ as the long-run component of X_t. This interpretation follows from the assumption that the deviations $(X_t - \bar{X}_t)$ have stable dynamics and unconditional expectations equal to zero. In particular these deviations are described by the following stable VAR:

$$A(L)\left(X_t - \bar{X}_t\right) = \epsilon_t \tag{1}$$

where $A(L)$ is a polynomial in the lag operator L and $\epsilon_t \sim N\left(\mathbf{0}, \Sigma_{\epsilon,t}\right)$. We assume that the reduced-form shocks ϵ_t have stochastic volatility:

$$\Sigma_{\epsilon,t} = B^{-1} S_t (B^{-1} S_t)' \tag{2}$$

where S_t is diagonal and B is lower triangular. Collecting the elements in the main diagonal of S_t in the vector s_t, we follow the well-established literature (see, for example, Cogley and Sargent, 2005; Primiceri, 2005) by modeling the time variation in the volatilities as:

$$\log s_t = \log s_{t-1} + \nu_t \qquad \nu_t \sim N\left(\mathbf{0}, \Sigma_\nu\right) \tag{3}$$

and we restrict Σ_ν to be diagonal.

The focus of our analysis is the long-run component $\bar{X}_t$ which is assumed to depend on a $(q \times 1)$ vector of latent variables θ_t:

$$\begin{cases} \bar{X}_t = h\left(\theta_t\right) \\ \theta_t = f\left(\theta_{t-1}, \eta_t\right) \end{cases} \tag{4}$$

where $h(\theta_t)$ and $f\left(\theta_{t-1}, \eta_t\right)$ are generic (potentially non-linear) functions, and η_t is a vector of exogenous Gaussian shocks. In this way we can specify the dynamics of the long-run component in a sufficiently general way, and in particular we are going to use equation (4) to define a long-run Phillips curve.

2.1 The model

We build a model for GDP per capita y_t, the inflation rate π_t and the nominal interest rate i_t. We use a bar over each respective variable to indicate its time-varying long-run component, e.g., $\bar{\pi}_t$ is the long-run component of inflation (i.e., trend inflation) at time t.

---

[9]This approach has been recently used by Maffei-Faccioli (2020) and Ascari and Fosso (2021).

We assume that potential output $\bar{y}_t$ can be decomposed as the sum of two components:

$$\bar{y}_t = y_t^* + \delta\left(\bar{\pi}_t\right) \tag{5}$$

where y_t^* is a trending component and $\delta\left(\bar{\pi}_t\right)$ is a function of trend inflation such that $\delta\left(0\right) = 0$. Then, we can interpret y_t^* as the long-run level of output in case of zero trend inflation, and we assume it has the following dynamics:

$$y_t^* = y_{t-1}^* + g_{t-1} + \eta_t^y \qquad \eta_t^y \sim N\left(0, \sigma_y^2\right) \tag{6}$$

$$g_t = g_{t-1} + \eta_t^g \qquad \eta_t^g \sim N\left(0, \sigma_g^2\right). \tag{7}$$

The assumption about the trend component y_t^* is quite standard in the literature, and as in Harvey and Todd (1983) and Clark (1987), we allow for both the slope and the level to change over time.[10]

We depart from the existing literature by adding the explicit possibility of a relation between the long-run level of output and trend inflation: equation (5) is the long-run Phillips curve. In particular, the function $\delta\left(\bar{\pi}_t\right)$ measures the long-run costs or benefits from having a positive trend inflation. We model it as a piecewise linear function:

$$\delta(\bar{\pi}_t) = \begin{cases} k_1 \bar{\pi}_t & \text{if } \bar{\pi}_t \leq \tau \\ k_2 \bar{\pi}_t + c_k & \text{if } \bar{\pi}_t > \tau. \end{cases} \tag{8}$$

With this assumption we allow for the slope of the long-run Phillips curve to change depending on trend inflation being higher or lower than a certain threshold τ. k_1 and k_2 are the slope parameters below and above the threshold, respectively, and c_k is a constant that assures continuity between the two piecewise lines. The main advantage of using a piecewise linear setting is the availability of an analytical expression for the likelihood function, so the efficiency of the estimator we propose is comparable to the one we use in case of linear models. Moreover, equation (8) is easy to interpret, and the posterior distribution of τ is a natural statistic to consider when reasoning about the potential costs or benefits of a positive level of trend inflation. We describe more formally how we propose to treat this simple class of models in Section 2.2, where we also discuss the pros and cons of this approach.

The long-run components of the other two variables in the model evolve as follows: trend

---

[10]In our specification the process for GDP is by assumption integrated of order 2. The more parsimonious option with $\sigma_g = 0$ has been extensively used in the literature (Watson, 1986; Kuttner, 1994; Planas et al., 2008). However, in the sample considered we find convenient to capture the slowdown in GDP as a slow moving decline in the growth rate of potential output (see also Maffei-Faccioli, 2020).

inflation dynamics are described by a random walk:[11]

$$\bar{\pi}_t = \bar{\pi}_{t-1} + \eta_t^\pi \qquad \eta_t^\pi \sim N\left(0, \sigma_\pi^2\right), \tag{9}$$

and the nominal interest rate obeys the long-run Fisher equation:

$$\bar{i}_t = \bar{\pi}_t + cg_t + z_t. \tag{10}$$

As in Laubach and Williams (2003), we assume that the long-run real interest rate is a function of the growth rate of potential output g_t and of a component z_t that captures all the slow moving trends that might affect the natural rate of interest, but are not directly included in the model. In particular we assume that z_t also evolves as a random walk:

$$z_t = z_{t-1} + \eta_t^z \qquad \eta_t^z \sim N\left(0, \sigma_z^2\right). \tag{11}$$

The model described above belongs to a general class of piecewise linear specifications in which equation (4) is written as:

$$\begin{cases} \bar{X}_t = H_t \theta_t \\ \theta_t = \bar{M}_t + \bar{G}_t \theta_{t-1} + \bar{P}_t \eta_t \end{cases} \tag{12}$$

with $\eta_t \sim N\left(\mathbf{0}, \Sigma_{\eta,t}\right)$ and $H_t, \bar{M}_t, \bar{G}_t, \bar{P}_t$ are matrices of appropriate dimensions that are functions of $\bar{\pi}_t$, that is an element of the latent vector θ_t. In particular, at each time t we have two possibilities depending on the region to which $\bar{\pi}_t$ belongs:

$$\left(H_t, \bar{M}_t, \bar{G}_t, \bar{P}_t\right) = \begin{cases} \left(H_{1,t}, \bar{M}_{1,t}, \bar{G}_{1,t}, \bar{P}_{1,t}\right) & \text{if } \bar{\pi}_t \leq \tau \\ \left(H_{2,t}, \bar{M}_{2,t}, \bar{G}_{2,t}, \bar{P}_{2,t}\right) & \text{if } \bar{\pi}_t > \tau. \end{cases} \tag{13}$$

In the general specification above, we allow the matrices in equation (13) to be time varying as commonly done for state space models - e.g., they can depend on data - even though in our particular empirical exercise based on (8) they will not.

For the model we just described, the vector $\bar{X}_t$ contains the long-run components of GDP per capita, inflation rate and the nominal interest rate, that is, potential output $\bar{y}_t$, trend inflation $\bar{\pi}_t$ and the long-run nominal interest rate $\bar{i}_t$. $H_t = H$ is a constant matrix. Finally, the latent vector $\theta_t = (\bar{\pi}_t, \bar{y}_t, g_t, z_t)'$, so it contains trend inflation $\bar{\pi}_t$, potential output $\bar{y}_t$, the growth rate of long-run output g_t and the residual component of the long-run real interest rate z_t. For a

---

[11]A large part of the literature also assumes stochastic volatility for the shock to trend inflation (see Stock and Watson, 2007, 2016; Mertens, 2016; Mertens and Nason, 2020). We make this assumption for the structural model in Section 3.

detailed description of the state space representation of our model, see Appendix A.4.

2.2 Econometric strategy

We use a Bayesian approach to estimate the joint posterior distribution of the unknown parameters and latent processes. Note that we have two sources of non-linearity: the piecewice linear nature of the state space model and the presence of stochastic volatility. Then, we find it convenient to base our econometric strategy on particle filtering. Our approach is tailored for the specific class of models we consider in this paper. In particular, there are two important aspects to stress. The first concerns the posterior distribution of the latent states which is approximated in an efficient way making use of the so called "Rao-Blackwellization" principle. The second aspect is about the posterior distributions of the parameters which are approximated jointly with the posterior of the latent processes through particle filtering. For clarity of exposition, we describe these two aspects one at a time.

2.2.1 Inference on the latent processes

Let's assume, for the moment, that the parameters are known and we only need to make inference on the latent processes, which are the elements of the latent vector θ_t and the stochastic volatilities in s_t. Appendix A.2.1 shows that we can write our time series model in the following state-space form:

$$Y_t = F_t \vartheta_t + \epsilon_t \tag{14}$$

$$\vartheta_t = M_t + G_t \vartheta_{t-1} + P_t \eta_t \tag{15}$$

where Y_t is an observed vector of dimension $n \times 1$ and the latent vector ϑ_t is equal to $\left(\theta_t' \quad \theta_{t-1}' \cdots \theta_{t-p}' \right)'$, where p is the number of lags in the VAR (1). Equations (14) and (15) define a piecewise linear state-space model with stochastic volatility, where the matrices of the state space are generally allowed to be a functions of ϑ_t as in equation (13), and the covariance matrix of the shocks ε_t is modeled as in equation (2).

Particle filtering allows to approximate the posterior distribution of the latent processes recursively. Assume that at time $t-1$ we have a set of N particles $\left\{ \vartheta_{t-1}^{(i)}, \log s_{t-1}^{(i)} \right\}_{i=1}^{N}$ with associated weights $\left\{ \omega_{t-1}^{(i)} \right\}_{i=1}^{N}$ that approximate the joint distribution of $(\vartheta_{t-1}, \log s_{t-1})$, and we want to get an analogous set of particles $\left\{ \vartheta_t^{(i)}, \log s_t^{(i)} \right\}_{i=1}^{N}$ and weights $\left\{ \omega_t^{(i)} \right\}_{i=1}^{N}$ to approximate:

$$p\left(\vartheta_t, \log s_t | Y_{1:t} \right), \tag{16}$$

where $p(\cdot)$ is a generic density function. The presence of stochastic volatility makes direct sampling from the joint posterior above not possible. Particle filtering is based on importance

sampling: we can choose an alternative distribution $q(\vartheta_t, \log s_t)$ called the "importance distribution" from which we draw the new set of particles. Standard derivations show that if the support of the posterior distribution is included in the support of the importance density, the associated weights for the respective particles will be proportional to:

$$\omega_t^i \propto \omega_{t-1}^i \frac{p\left(\vartheta_t^{(i)}, \log s_t^{(i)} | \vartheta_{t-1}^{(i)}, \log s_{t-1}^{(i)}, Y_{1:t}\right)}{q\left(\vartheta_t^{(i)}, \log s_t^{(i)}\right)} \qquad \text{for } i = 1, ... N. \tag{17}$$

The choice of the importance distribution is crucial for accuracy of the approximation. The best choice in terms of statistical efficiency is the density in the numerator of (17), that is the posterior distribution of $(\vartheta_t, \log s_t)$, conditional on the data and on the past draws. This case would imply perfectly equal weights. As we already mentioned, we are not able to draw directly from the full conditional posterior, however we can get close by writing it as:

$$p\left(\vartheta_t, \log s_t | \vartheta_{t-1}^{(i)}, \log s_{t-1}^{(i)}, Y_{1:t}\right) = p\left(\vartheta_t | \log s_t, \vartheta_{t-1}^{(i)}, \log s_{t-1}^{(i)}, Y_{1:t}\right) p\left(\log s_t | \vartheta_{t-1}^{(i)}, \log s_{t-1}^{(i)}, Y_{1:t}\right).$$

The first factor on the right hand side is the posterior distribution of the latent vector ϑ_t conditional on stochastic volatility. A methodological contribution of this paper is to show that this distribution is available analytically. Then, given a set of particles for $\log s_t$, we can directly draw the particles for ϑ_t from its full conditional posterior distribution. In this way we increase the efficiency of our estimator as an implication of the Rao-Blackwell theorem. In practice this is a crucial improvement.

Note that this procedure is analogous for conditionally linear models: if we consider the particular case in which the state space form (14) and (15) has linear propagation dynamics (when the matrices of the system do not depend on trend inflation), the full conditional distribution of ϑ_t is also available analytically through the well-known Kalman filter recursions. Our contribution is to show that similar recursions exist for the class of piecewise linear models we consider, such that we can write both the likelihood and the posterior distribution of ϑ_t in explicit forms. Below we give an intuitive derivation of both densities, referring to Appendix A.2.2 for the details.

The likelihood. The contribution to the likelihood of the observation at time t, for particle i, is:

$$p\left(Y_t | \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) = p\left(Y_t | \bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) \Pr\left(\bar{\pi}_t \leq \tau | \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right)$$
$$+ p\left(Y_t | \bar{\pi}_t > \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) \Pr\left(\bar{\pi}_t > \tau | \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right), \tag{18}$$

where we omit the conditioning on $\log s_t$ to simplify the notation, and we indicate with $\Pr(x)$ the probability of an event x. Start considering the first addend of equation (18): it is the product of a density function scaled by the probability that trend inflation is below the threshold. The density is given by the integral:

$$p\left(Y_t|\bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) = \int_{\bar{\pi}_t \leq \tau} p\left(Y_t, \bar{\pi}_t|\vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) d\bar{\pi}_t . \tag{19}$$

Equation (19) can be solved analytically: it corresponds to the definition of the Unified Skew Normal (SUN) distribution introduced by Arellano-Valle and Azzalini (2006).[12] Following an analogous reasoning, we recognise that the density function of the second addend of equation (18) is also a SUN.

The probabilities that appear in the two addends are scaling factors that make sure the likelihood is a proper density that integrates to one. These probabilities are easily computed, given the random walk assumption for trend inflation. Then, the likelihood is a combination of two Unified Skew Normal densities with parameters that are recursively updated: the exact formulas are shown in Appendix A.2.2.

The posterior distribution of ϑ_t. The posterior distribution of ϑ_t is also the sum of two pieces, depending on trend inflation being above or below the threshold:

$$\begin{aligned} p\left(\vartheta_t|\vartheta_{t-1}^{(i)}, Y_{1:t}\right) &= p\left(\vartheta_t|\bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t}\right) \Pr\left(\bar{\pi}_t \leq \tau|\vartheta_{t-1}^{(i)}, Y_{1:t}\right) \\ &+ p\left(\vartheta_t|\bar{\pi}_t > \tau, \vartheta_{t-1}^{(i)}, Y_{1:t}\right) \Pr\left(\bar{\pi}_t > \tau|\vartheta_{t-1}^{(i)}, Y_{1:t}\right). \end{aligned} \tag{20}$$

Again, let's focus on the first addend of equation (20). Since trend inflation is an element of ϑ_t, the density is a multivariate truncated Normal with the truncation that only applies to $\bar{\pi}_t$. The probability that scales this distribution is proportional to:

$$\Pr\left(\bar{\pi}_t \leq \tau|\vartheta_{t-1}^{(i)}, Y_{1:t}\right) \propto p\left(Y_t|\bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) \Pr\left(\bar{\pi}_t \leq \tau|\vartheta_{t-1}^{(i)}, Y_{1:t-1}\right), \tag{21}$$

that is exactly the first addend of equation (18). An analogous reasoning hold for the second piece of the posterior distribution.

The particle filter. Using the derivations above we can get draws of the latent processes in two steps: first, we get a particle approximation for the volatilities using their law of motion (3). This proposal is labeled a "blind" distribution because it is not conditional on the data at time t. Then, we can use the full conditional distribution to obtain the corresponding particles

---

[12]The Unified Skew Normal is defined by Arellano-Valle and Azzalini (2006) as a generalization of the Skew Normal distribution by Azzalini (1985).

for ϑ_t.

We choose to implement a "resample - propagation" scheme based on the Auxiliary particle filter by Pitt and Shephard (1999). The idea is to start selecting the "best" particles through a resampling step of $\left\{\vartheta_{t-1}^{(i)}, \log s_{t-1}^{(i)}\right\}_{i=1}^{N}$ using weights proportional to the predictive likelihood. Then, the resampled particles are propagated to time t using the two steps procedure just described. All the details of the filter are described in Appendix A.2.2.

Discussion. The state-space form (14) and (15) is quite general and it can be useful for a wide range of applications using likelihood-based methods. A common trade-off in the choice of the model specification is the following: on one hand it is desirable to estimate a fully non-linear model in order to reduce the misspecification. However, this task might be too difficult, or the approximation of the likelihood can be poor due to computational constraints. On the other hand, a linear approximation may suffer from model misspecifications but has the advantage that the likelihood function can be computed analytically. In this trade-off, we propose a third option: the piecewise linear specification. The advantage of this choice is that it reduces the misspecification with respect to the linear case, while keeping the analytical availability of the likelihood. The cost in this case is in the number of parameters to estimate, as clear from equation (13). The more the number of intervals, the better the approximation, but the more the number of parameters to estimate. We lack a formal criterion to choose an appropriate number of intervals. While we think this is an interesting question for future research, in this paper we opt for the simple choice of a single break.

2.2.2 Inference on the parameters

For the estimation of non-linear macroeconomic models there is a strong tradition that makes use of particle filters to get an approximation of the likelihood function in the context of Markov Chain Monte Carlo methods, as pioneered by Fernández-Villaverde and Rubio-Ramírez (2007). In this paper, instead, the use of a particle filtering strategy directly aims at approximating the joint posterior distribution of the latent processes and the parameters.

We combine two approaches. We primarily use the particle learning scheme by Carvalho et al. (2010). The methodology consists of augmenting the vector of latent processes with sufficient statistics for the full conditional distributions of the different parameters. This idea uses the same "Rao-Blackwellization" principle to increase the efficiency of the estimator. Unfortunately, we are not able to use it for all the parameters: in particular sufficient statistics are not available for the posterior distribution of the threshold τ. To estimate the latter, we use a mixture of Normal distributions, following Liu and West (2001).

The use of particle filters, and in general sequential Monte Carlo methods, to estimate the parameters of macroeconomic models is becoming more common (e.g., Ascari et al., 2019;

Mertens and Nason, 2020). With respect to the more traditional approaches based on Markov Chain Monte Carlo (MCMC), sequential Monte Carlo (SMC) methods do not have the problems related to the convergence of the chain (which can be severe in case of non-linear models), and are much better at approximating multi-modal posterior distributions. Moreover, it is easy to exploit computational advantages from parallelization, especially in this era of multi-core processors (see Herbst and Schorfheide, 2014).

2.2.3 Data and priors

We estimate the model using three U.S. quarterly time series, available from the St. Louis Fed FRED database: per capita real GDP, (annualized) quarterly growth rate of the GDP deflator, and the Federal Funds rate, over the period $1960Q1 - 2008Q2$. We choose to exclude post-2008 data from our sample to avoid technical issues related to the lower bound on the nominal interest rate.

The prior distribution of the parameters in the model for the long-run component $\bar{X}$ are reported in Table 1. According to our prior information the long-run Phillips curve is vertical: this is summarized by the choice of a Normal distribution with mean equal to 0 and standard deviation equal to 0.6 for both k_1 and k_2. The prior for the threshold τ is centered at 4, which is close to the average of inflation in our sample. This prior is quite informative because we want to avoid wasting effort in exploring unrealistic region of the support, especially considering the range of trend inflation estimates in the literature (see, e.g., Cogley and Sbordone, 2008; Cogley et al., 2010; Stock and Watson, 2016; Mertens, 2016; Mertens and Nason, 2020). The parameter c governs the relation between the growth rate of potential output and the natural interest rate. While the empirical evidence in favor of this link has been debated (Hamilton et al., 2016), this relation can be derived from the Euler equation in a micro-founded structural model, and we make our prior consistent with logarithmic utility (the nominal interest rate is expressed in annual terms). The priors for the variances of the shocks to the long-run components are assumed to follow standard Inverse-Gamma distributions whose parameters are shown in Table 1. The short-run dynamics are described by the VAR in equation (1) for which we choose 4 lags. For the 36 parameters in $A(L)$, we use a standard Minnesota prior with the hyperparameter governing the overall tightness equal to 0.2, the one for the cross-variable tightness equal to 0.5 and the hyperparameter for the lag length decay equal to 1. The prior for the matrices in equation (2) that decompose the covariance matrix $\Sigma_{\epsilon,t}$ is centered at the OLS estimates of the corresponding VAR with constant volatility. In particular, we assume an Inverse-Wishart distribution with 5 degrees of freedom and we consider the implied distributions for each coefficient. Finally, the variances of the shocks to the stochastic volatilities have an Inverse-Gamma prior with mean 0.025^2 and 5 degrees of freedom.

Table 1: Prior and posterior distributions for the long-run component of the time series model

	Prior			Posterior	
Parameter	Density	Mean	Standard Deviation	Model L	Model PWL
k_1	Normal	0.0	0.6	-0.15 $[-0.49\ 0.19]$	-0.07 $[-0.51\ 0.38]$
k_2	Normal	0.0	0.6		-0.92 $[-1.35\ -0.47]$
τ	Normal	4.0	0.3		4.09 $[3.88\ 4.29]$
c	Normal	4.0	0.75	3.53 $[3.28\ 3.78]$	2.93 $[2.68\ 3.18]$
	Density	Mean	Degrees of freedom		
σ_π^2	Inverse Gamma	0.25^2	15	0.2^2 $[0.18^2\ 0.23^2]$	0.23^2 $[0.21^2\ 0.26^2]$
σ_y^2	Inverse Gamma	0.5^2	15	0.49^2 $[0.45^2\ 0.54^2]$	0.59^2 $[0.54^2\ 0.66^2]$
σ_g^2	Inverse Gamma	0.05^2	15	0.043^2 $[0.039^2\ 0.048^2]$	0.05^2 $[0.042^2\ 0.058^2]$
σ_z^2	Inverse Gamma	0.15^2	15	0.14^2 $[0.13^2\ 0.16^2]$	0.17^2 $[0.14^2\ 0.19^2]$

Posterior median and the 90% probability interval in brackets

2.3 Results

We estimate two specifications of our model: the linear case (Model L) in which we assume that the slope of the long-run Phillips curve is constant, and the piecewise linear case (Model PWL) where we allow for two different slopes, as extensively explained in the previous Section.[13] Table 1 reports the median and the 90% credibility interval of the posterior estimates for the parameters of the long-run components. We discuss three main results. First, our evidence suggests that the long-run Phillips curve is not vertical, but it is non-linear and negatively sloped, and substantially so. Hence, trend inflation is an important determinant of potential output. Second, we define the concept of the "long-run output gap (LROG)" and quantify the losses due to high trend inflation in U.S. data. Third, our model implies that failing to take into account the LROG results in biased estimates of the short-run output gap.

2.3.1 The long-run Phillips curve is non-linear and negatively sloped

Considering the linear specification, the estimates of Model L are in line with the results in the existing literature, in particular with Benati (2015), providing no evidence against a vertical long-run Phillips curve. The posterior distribution of the slope parameter k_1 is concentrated around zero with a negative median.

---

[13]We employ $500,000$ particles for each time iteration of our particle filter, for both Model L and Model PWL.

Figure 1: Online inference of the slope k_1 - Linear model.

Estimating the parameters of the model through particle filtering allows us to check how the parameters' estimates evolve over the sample, depending on the available information. Figure 1 shows sequential inference for the slope coefficient k_1 of Model L. As a reference, we report the time series of annualized inflation in the top panel. While at the very beginning of the sample the median estimate moves more toward the positive side, when inflation starts fluctuating around 4% during the Seventies, then the posterior estimate of k_1 concentrates its mass more toward negative values. During the Great Moderation, the posterior distribution of k_1 remains stable due to little variation in the estimated trend inflation series.

This result suggests that there is specific informational content on the slope of the LRPC in different subsamples. In particular, when inflation is persistently high during the Great Inflation period, then the model captures a negative correlation between the long-run components of output and inflation. As stressed by Benati (2015), since this is the period in which inflation clearly exhibits a unit root, the relevant information for the identification of the slope of the LRPC comes from this sample. We deal with this identification problem allowing for non-linearity in Model PWL.

When we allow for the slope of the long-run Phillips curve to change, the estimation prefers to use this feature to interpret the data. While the estimate of k_1 remains around zero, the posterior distribution of k_2 in Table 1 has a median of -0.92 and the 90% probability interval lies entirely on the negative side. Note that the piecewise linear specification admits the linear model as a particular case. However, the estimation rejects this option as evident from Figure 2, which compares the prior and the posterior distributions of the slopes k_1 and k_2 for Model PWL.

Figure 2: Estimates of the slopes of the long-run Phillips curve - Piecewise linear model.
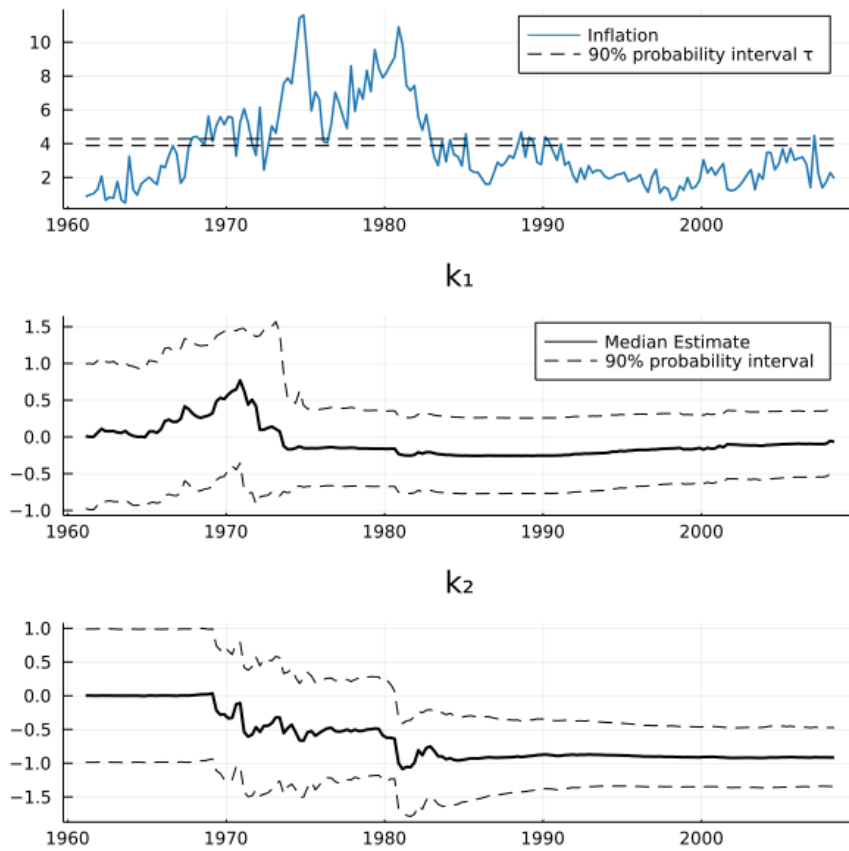

Figure 3: Online inference of the slopes k_1 and k_2 - Piecewise linear model.

Then, the first important result of our analysis is that the long-run Phillips curve is non-linear and negatively sloped. The negative effects on output materialize when trend inflation is above the threshold value τ, estimated to be about 4%. The second important takeaway from these results is that the absolute value of the LRPC slope is sizeable and roughly equals one. Hence, every percentage point increase in trend inflation above the threshold value of 4% is related to about 1% decrease in potential output per year.

It is informative to look at the sequential inference about the slopes of the LRPC in the non-linear model just as we did for the linear case. Figure 3 shows how the median and the 90% probability interval of the posterior distributions of k_1 and k_2 evolve recursively over the sample. In the first panel, we plot the (annualized) inflation rate as well as the posterior probability interval of τ (dotted lines). The pattern of the posterior distribution of k_1 initially resembles the one in the linear model: it becomes slightly positive until the beginning of the Seventies when it reverts back toward zero. Now the model has the option to let k_2 to capture the negative correlation between output and inflation. The persistent high inflation during this period makes trend inflation cross the threshold and k_2 is confidently estimated to be negative. Figure 3 makes it clear why non-linearity is important in finding evidence in favour of a non-vertical (and negatively sloped) LRPC.

Figure 4 shows the estimated long-run Phillips curve by plotting the deviation of potential output from its zero trend inflation counterfactual as a function of trend inflation based on (5). The 90% probability interval reflects the uncertainty around the parameters estimates. The LRPC is vertical when trend inflation is below the threshold τ and negatively sloped above. Note that the uncertainty around k_1 makes our results consistent both with models in which there is a positive optimal level of trend inflation (see Adam and Weber, 2019; Abbritti et al., 2021) and with frameworks in which the best value for trend inflation is zero.

Finally, it is important to stress that we interpret our piecewise linear model as an approximation to an underlying non-linear relation as the one we estimate in Section 3. This means that the value of the estimated threshold, while giving an important indication, does not have to be taken literally: trend inflation can imply potential output losses even below τ, as clear from the figure.

2.3.2 The long-run output gap

Under the assumption of a vertical long-run Phillips curve $\bar{y}_t$ is exactly equal to y_t^* in (5). However, our estimates suggest that when trend inflation is above the threshold τ, potential output $\bar{y}_t$ is different from y_t^*, the potential output under zero trend inflation. We call this difference the long-run output gap (see also the discussion at the end of Section 3.3). The estimates allow us to quantify the long-run output gap in our sample and to answer the following question: how much was the loss in potential output due to high trend inflation in the U.S.

Figure 4: Long-run Phillips curve estimated through the piecewise linear model: median (continuous line) and 90% probability interval (dashed lines).

data?

Our estimate of trend inflation is reported in Figure 5. During the Great Inflation period the median reaches almost 7.5%, which induces substantial losses in potential output as shown in Figure 6. During the Great Inflation, the median of the output cost associated with the long-run Phillips curve had been on average about 2% per year and the maximum reached 3%.[14]

An alternative explanation of the high inflation and low output during the Great inflation period is related to the important role of supply shocks. These are shocks to the short-run/cyclical component of the variables. To allow our empirical framework to account for this possibility, we allow for stochastic volatility in the reduced-form shocks in the VAR (1). Hence, the model could capture large shocks, and eventually interpret the high inflation of the '70s as short-run deviations due to large supply shocks. Despite this possibility, the high inflation and the low output are so persistent that the model prefers to attribute the negative correlation between these two variables to their respective permanent components.

This finding suggests that the non-linearity in the LRPC has important implications for the measurement of the costs associated with higher trend inflation, being related to a decline in potential output. We discuss this point in further details below.

---

[14]Both trend inflation and the long run output gap are filtered estimates obtained by conditioning on the median estimates of the parameters. The choice of showing the filtered series is motivated by an easier comparison with the structural model (see next section) for which smoothed estimates are computationally difficult to obtain.

Figure 5: Inflation and trend inflation - Piecewise linear model.

Figure 6: Long-run output gap estimated through the piecewise linear model.

Table 2: Parameters in the long-run component of the time series model, case with $k_1 = 0$

	Posterior	
	Model L	Model PWL
k_2		-1.04 [-1.52 -0.51]
τ		4.08 [3.64 4.52]
c	3.08 [2.88 3.28]	4.17 [3.85 4.49]
σ_π^2	0.25^2 [0.23^2 0.27^2]	0.23^2 [0.21^2 0.26^2]
σ_y^2	0.56^2 [0.51^2 0.61^2]	0.53^2 [0.48^2 0.59^2]
σ_g^2	0.035^2 [0.032^2 0.04^2]	0.059^2 [0.044^2 0.068^2]
σ_z^2	0.16^2 [0.14^2 0.18^2]	0.16^2 [0.14^2 0.19^2]

Posterior median and the 90% probability interval in brackets

2.3.3 Business cycle measurement: implications for the short-run output gap

A negatively-sloped LRPC has important consequences for the measurement of business cycles. According to our estimates, assuming a vertical LRPC, that is, imposing the absence of a relation between output and inflation in the long-run, leads to biased estimates of the short-run output gap. In particular, during the Great Inflation period, these traditional estimates of the short-run output gap tend to overstate the negative development of the business cycle. First, we estimate a version of Model L presented in Section 2 but now we impose $k_1 = 0$, that is a vertical LRPC, rather than estimating k_1 as in Table 1. The posterior estimates of the parameters of the long-run component are reported in the first column of Table 2.

In Figure 7 we show the potential output and the implied output gap measures estimated under the two models, focusing on the Great inflation period.[15] The upper-left panel compares the potential output estimated assuming a vertical LRPC with the corresponding CBO measure: the inference is extremely similar implying almost the same values for the output gap during the high trend inflation period (bottom-left panel). Let's stress again that by assumption we are offsetting the role of the long-run output gap which is calibrated to be zero.

We now relax this assumption by estimating a piecewise linear version of this model in which we still calibrate $k_1 = 0$, but we allow for k_2 to be different from zero (we use the same Gaussian prior as before). Calibrating k_1 makes this model directly comparable to the linear version we just described which is a nested case. The posterior estimates are reported in the

---

[15]Note that these measure are smoothed estimates.

Figure 7: Potential output (upper panels) and implied output gap (lower panels) measures under the linear model, the piecewice linear model and the official CBO estimate.

second column of Table 2. The first result to highlight is that k_2 is confidently estimated to be negative: the data rejects the assumption of a vertical LRPC. The inference on the parameters is in line with our previous estimates, in particular about the posterior distribution of k_2 and the resulting costs related to high trend inflation.

The upper-right panel of Figure 7 plots the measures of potential output estimated by the two versions of our VAR. During the period of high trend inflation we notice a sizable decline in the PWL estimate of potential output with respect to the corresponding measure under a vertical LRPC. This decline captures the loss in potential output related to the non-linear and negatively sloped LRPC. Then, while the linear model (similarly to the CBO) interprets the low output as a drop in the output gap, Model PLW attributes part of this loss to potential output because of the high trend inflation. This affects the measure of the output gap, resulting in a different interpretation of the cyclical component of GDP, as clear from the bottom-right panel. For example, under Model L the output gap turns negative already in the first half of 1980 and reaches its trough at the end of 1982 with a value of -8.2%; under Model PWL, the output gap turns negative only at the end of 1981, and its lowest value in 1982 is equal to -5.7%. In other words, assuming a vertical LRPC leads to overstating the negative development of the business cycle in times of high inflation.

Finally, we note that the time series model does not give any causal interpretation about the LRPC: we can not state if the long-run losses in output are the consequence of high trend inflation, or vice versa. In order to get such an interpretation we need a structural model. We discuss this point in the next section.

3   A structural approach

We have shown above that the U.S. data are better described by a time-series model implying a negatively sloped LRPC in times of high trend inflation. However, such a model is not able to give a structural interpretation and to explain the causal link in the long-run relationship between output and inflation. Which type of structural model could deliver a theoretical interpretation while at the same time be quantitatively in line with our empirical results? Fortunately, we do not have to look very far: the standard workhorse New Keynesian model would do it.[16]

Therefore, we embed time-varying positive trend inflation in a simple dynamic stochastic general equilibrium (DSGE) model. As in the time-series model, the variables in the DSGE model are decomposed into short-run and long-run components, and we estimate the two com-

---

[16]While there could be other possible model that could provide alternative structural interpretation to a negative sloped LRPC, it seemed to us natural to start with the most popular model in the literature. An interesting question for future research would be to ask if other alternative structural interpretation are theoretically and quantitatively coherent with our time-series results.

ponents together with the parameters of the model. Our objective is to get a model-based estimate of how trend inflation is related to potential output in the long-run, in other words, estimate the LRPC using a structural DSGE model.

3.1 The model

The artificial economy is a variant of the Generalized New Keynesian (GNK) model in Ascari and Sbordone (2014). The model consists of a representative household, a representative final-good firm, a continuum of intermediate-good firms, and a monetary authority. The model is very standard, so here we describe the main features, while Appendix B contains more details. The novelty comes from the assumption of a time-varying steady state or trend inflation. Hence, we need to take particular care of how we log-linearize the model around a time-varying steady state.

The representative agent maximizes the following expected utility function where preferences are additively separable in individual consumption of final goods, C_t, and labor, N_t:

$$E_0 \sum_{t=0}^{\infty} \beta^t d_t \left[\ln \left(C_t - h\tilde{C}_{t-1} \right) - d_n \frac{N_t^{1+\varphi}}{1+\varphi} \right] \qquad 0 < \beta < 1, d_n > 0, \ \varphi \geq 0, \ 0 \leq h < 1, \quad (22)$$

where φ is the inverse of the Frisch labor supply elasticity, d_n governs the steady state disutility of work, and h is the degree of habit persistence in consumption. Habit persistence is 'external', meaning that the consumer is concerned with the level of her current consumption C_t relative to the aggregate consumption in the previous period $\tilde{C}_{t-1}$. The term d_t stands for a shock to the discount factor, β, which follows the stationary autoregressive process:

$$\ln d_t = (1 - \rho_d)\, \overline{d} + \rho_d \ln d_{t-1} + \sigma_{d,t}\epsilon_{d,t}, \quad (23)$$

where $\epsilon_{d,t}$ is i.i.d $N(0,1)$ and $\sigma_{d,t}$ denotes time-varying standard deviation of the preference shock. The period budget constraint is given by:

$$P_t C_t + R_t^{-1} B_t = W_t N_t - T_t + D_t + B_{t-1}, \quad (24)$$

where P_t is the price level, R_t is the gross nominal interest rate on bonds, B_t is one-period bond holdings, W_t is the nominal wage rate, T_t is lump sum taxes, and D_t is the profit income.

Firms come in two forms. A final-good firm produces output for consumption. This output is made from the range of differentiated goods that are supplied by intermediate-goods firms who have market power. Each intermediate-good firm i produces a differentiated good $Y_{i,t}$ under monopolistic competition using the production function $Y_{i,t} = A_t N_{i,t}^{1-\alpha}$. Here A_t denotes the level of aggregate technology that is non-stationary and its growth rate $g_t \equiv A_t/A_{t-1}$ follows

the process:

$$\ln g_t = \ln \overline{g} + \sigma_{g,t} \epsilon_{g,t}, \tag{25}$$

where $\overline{g}$ is the steady-state gross rate of technological progress which is also equal to the steady-state balanced growth rate, $\epsilon_{g,t}$ is a i.i.d. $N(0,1)$ and $\sigma_{g,t}$ is the time-varying standard deviation of the technology shock. Intermediate-goods producers are subject to nominal rigidities in the form of Calvo (1983) with partial indexation. Hence, they face a constant probability, $0 < (1-\theta) < 1$, of being able to adjust their price while the price of a firm that cannot change the price is automatically indexed to past-inflation with a degree χ.

The central bank's monetary policy follows a Taylor rule featuring inertia and responding to the inflation gap, the output gap and output growth gap:

$$\ln R_t = \rho \ln R_{t-1} + (1-\rho) \ln \bar{R}_t + (1-\rho) \psi_\pi \ln \left(\frac{\pi_t}{\overline{\pi}_t}\right) + (1-\rho) \psi_x \ln \left(\frac{X_t}{\overline{X}_t}\right) + (1-\rho) \psi_{\Delta y} \ln \left(\frac{g_t^y}{g^y}\right) + \sigma_{r,t} \epsilon_{r,t}, \tag{26}$$

where X_t is the output gap defined as the deviation of output from its natural (or flexible price) level, $\overline{X}_t$ is the steady state output gap, g_t^y is the growth rate of output, $g^y = \overline{g}$ is the steady state growth rate of output, and $\epsilon_{r,t}$ is a i.i.d. $N(0,1)$ monetary policy shock with time-varying standard deviation $\sigma_{r,t}$. The parameters ψ_π, ψ_x and $\psi_{\Delta y}$ govern the central bank's responses to the inflation gap, output gap and output growth, respectively. Note that the output gap is defined as usual as the deviation of the level of output from the flexible price output level. Thus, it enters the Taylor rule relative to its long-run level, that is, the long-run output gap, which emerges in this model (see the discussion in Subsection 3.3.2), as in the non-structural model of the Section 2. The idea is that monetary policy responds, as usually assumed, to short-run deviation of variables from their steady state values. The inflation gap is the deviation of the inflation rate from time-varying trend inflation, i.e., $\overline{\pi}_t$, which represents the central's banks (time-varying) inflation target and follows a unit root process:

$$\ln \overline{\pi}_t = \ln \overline{\pi}_{t-1} + \sigma_{\overline{\pi},t} \epsilon_{\overline{\pi},t}, \tag{27}$$

where $\epsilon_{\overline{\pi},t}$ is i.i.d. $N(0,1)$ and $\sigma_{\overline{\pi},t}$ denotes time-varying standard deviation of the inflation target shock.

Following Justiniano and Primiceri (2008), we allow for stochastic volatility by assuming that each element of σ_t evolves independently according to the following stochastic process:

$$\ln \sigma_{i,t} = \ln \sigma_{i,t-1} + \nu_{i,t} \qquad \nu_{i,t} \sim N\left(0, \delta_i^2\right) \qquad for \ i = d, g, r, \overline{\pi}. \tag{28}$$

The steady state of the system is stochastically changing because it is characterized by time-varying trend inflation, $\overline{\pi}_t$, and also because of stochastic (unit-root) technology process.

As a result, we first de-trend the real variables of the model to remove the trend in technology and then log-linearize the resulting non-linear model around a drifting steady state. Here, we describe heuristically the state-space form for the estimation, composed of the following elements.[17]

1. A set of equations that define the detrended variables as deviations from steady state $Z_t = \overline{Z}_t Z_t / \overline{Z}_t = \overline{Z}_t \widetilde{Z}_t$. In logs: $\ln Z_t = \ln \overline{Z}_t + \widehat{Z}_t$, where $\widehat{Z}_t \equiv \ln \widetilde{Z}_t$.

2. A law of motion for $\overline{\pi}_t$: $\overline{\pi}_t = \overline{\pi}_{t-1} \exp(\sigma_{\overline{\pi},t} \epsilon_{\overline{\pi},t})$.

3. A set of equations that define the steady state of the variables as a function of $\overline{\pi}_t$: $\overline{Z}_t = f(\overline{\pi}_t, \overline{\pi}_{t-1})$.

We can then write the usual system for the dynamics of the log-linearized variables in canonical form, but now the system will have time-varying parameters as they are functions of $\overline{\pi}_t$:

$$\Gamma_0(\overline{\pi}_t)\hat{Z}_t = \Gamma_1(\overline{\pi}_t)\hat{Z}_{t-1} + \Psi(\overline{\pi}_t)\varepsilon_t + \Pi(\overline{\pi}_t)\eta_t, \tag{29}$$

where ε_t is a vector of exogenous disturbances and η_t is a vector of one-step ahead forecast errors. However, for any given value of $\overline{\pi}_t$ (and for a given realization of stochastic volatility), the system (29) is conditionally linear and can be solved with standard methods.

At each time t, we observe a vector of data denoted by y_t. Then, the solution of model (29) has the following state-space representation:

$$
\begin{aligned}
y_t &= c_1 + F\hat{Z}_t \\
\hat{Z}_t &= c_{2,t} + M_{z,t}\hat{Z}_{t-1} + M_{\varepsilon,t}\varepsilon_t \qquad \varepsilon_t \sim N(0, \Sigma_{\varepsilon,t})
\end{aligned}
\tag{30}
$$

where $\Sigma_{\varepsilon,t}$ is a diagonal matrix with $\sigma_{i,t}$ of time-varying standard deviations on the main diagonal. Note that the terms that appear in the state equations, $c_{2,t}$, $M_{z,t}$, $M_{\varepsilon,t}$, depend on t due to time-varying trend inflation. In other words, when trend inflation drifts, the coefficients of the state equation in (30) also drift, even if the underlying structural parameters are constant.

3.2 Econometric strategy

We follow a Bayesian approach to make inference regarding the parameters and the latent processes of the DSGE model. The presence of time-varying trend inflation as well as stochastic volatility leads to a non-Gaussian and analytically intractable likelihood function. We use the same particle filtering strategy as for the time-series model to directly approximate the joint

---

[17]Appendix B presents the log-linearized equations of the model. As in Cogley and Sbordone (2008), the steady state of the model is time-varying because of drifts in trend inflation. As such, care must be taken when log-linearizing the model. As in Cogley and Sbordone (2008) - see footnote 5 therein - we assume 'anticipated utility' following Kreps (1998).

posterior distribution of both the parameters and the latent state variables. In the context of DSGE models this approach has been use by Chen et al. (2010) and Ascari et al. (2019).

The main idea of using SMC methods is to get an approximation of a complicated posterior through the sequential approximation of simpler distributions. Two approaches have been proposed for DSGE models: (i) a likelihood tempering scheme (Herbst and Schorfheide, 2014) in which the simpler sequential distributions are obtained by tempering the likelihood function; (ii) a filtering scheme in which the intermediate distributions are obtained by sequentially adding observations to the likelihood function.[18] In this paper, in analogy with the estimation of the time series model in the previous section, we opt for the second approach.

As for the time series model, we can get higher efficiency through Rao-Blackwellization: conditional on $\bar{\pi}_t$ and the realization of stochastic volatility $\sigma_{i,t}$, the state space (30) is linear and Gaussian. This implies that, given a set of particles for $\bar{\pi}_t$ and $\sigma_{i,t}$, both the predictive likelihood and the full conditional distribution of the other latent states are analytically available through the standard Kalman filter recursion.

Inference on the parameter is obtained combining the same methods as for the time series model. The parameters are divided into two sets: one with the variances of the disturbance to the stochastic volatility processes, and one with all the other structural parameters. For the former, we assume Inverse-Gamma priors, allowing us to characterize the posterior distribution analytically using sufficient statistics computed as functions of the data and the latent processes of the model. We make inference on these parameters using the particle learning approach of Carvalho et al. (2010). For the latter, we approximate the posterior distribution through mixtures of Normal distributions, following Liu and West (2001).[19]

3.2.1 Data and priors

We estimate the model using the same U.S. data as in the time-series analysis: per capita real GDP growth rate, (annualized) quarterly growth rate of the GDP deflator and the Federal Funds rate, over the period $1960Q1 - 2008Q2$.

As customary when taking DSGE models to the data, we calibrate a small number of parameters. In particular, we set the discount factor β to 0.997, the steady state markup to 10 per cent (i.e., $\varepsilon = 11$), the inverse of the labor supply elasticity φ to 1, the quarterly net steady state output growth rate $\bar{g}$ to 0.5, and the degree of decreasing returns to scale α to 0.3. In light of the result of Cogley and Sbordone (2008) regarding the lack of support for intrinsic inertia in the GNK Phillips curve, the model is estimated without backward-looking price indexation, i.e., $\chi = 0$. The remaining parametes are estimated. Table 3 summarizes the specification of the prior distributions. The prior for the inflation coefficient ψ_π follows a Gamma distribution

---

[18]See also Creal (2007) and Herbst and Schorfheide (2016).

[19]For a more detailed description of the SMC algorithm, we refer to the online appendix of Ascari et al. (2019).

Table 3: Prior and Posterior Distributions

Parameter		Prior		Posterior
	Density	Mean	St Dev	
ψ_π	Gamma	1.5	0.5	2.36 [2.04 2.7]
ψ_x	Gamma	0.125	0.05	0.13 [0.08 0.21]
$\psi_{\Delta y}$	Gamma	0.125	0.05	0.34 [0.20 0.56]
ρ	Beta	0.7	0.1	0.75 [0.72 0.79]
h	Beta	0.5	0.1	0.39 [0.33 0.45]
r^*	Gamma	2	0.5	2.12 [1.87 2.41]
θ	Beta	0.5	0.1	0.49 [0.45 0.52]
ρ_d	Beta	0.7	0.1	0.79 [0.74 0.83]
	Density	Mean	Degree of freedom	
δ_d^2	Inverse Gamma	0.02^2	5	0.119^2 [0.105^2 0.135^2]
δ_g^2	Inverse Gamma	0.02^2	5	0.028^2 [0.023^2 0.032^2]
δ_r^2	Inverse Gamma	0.02^2	5	0.057^2 [0.046^2 0.065^2]
δ_π^2	Inverse Gamma	0.02^2	5	0.029^2 [0.026^2 0.034^2]

Posterior median and 90% credibility interval in brackets

centered at 1.50 with a standard deviation of 0.50 while the response coefficient to the output gap and output growth are centered at 0.125 with standard deviation 0.05. We employ a Beta distribution with mean 0.70 for the interest rate smoothing parameter ρ and the persistence of the discount factor shock ρ_d, while the Calvo probability θ and habit persistence in consumption h are centered around 0.50. The steady state real interest rate follows a Gamma distribution centered at 2. For the variances of the shocks to the volatilities δ_i^2, we assume an Inverse Gamma distribution with mean equal to 0.02 and 5 degrees of freedom. Our estimation assumes a unique rational expectations equilibrium, i.e., we do not allow for indeterminacy.[20]

3.3 Results

Table 3 reports the posterior medians and the 90% posterior density intervals based on one million particles from the final stage in the SMC algorithm. The Taylor rule's response to the inflation gap is strongly active as the estimated response lies mostly above 2. We also find a moderate response to the output gap and a strong response to output growth along with high degree of interest rate smoothing. The degree of habit formation is somewhat low and is close to 0.4. The posterior mean for the degree of price stickiness θ turns out to be around 0.5, which is smaller than the estimates reported in Smets and Wouters (2007) and Justiniano et al. (2010) and implies an expected price duration of six months.

Figure 8 plots the model-implied evolution of trend inflation along with the 90% posterior density interval and the actual GDP deflator inflation rate. Trend inflation began rising in the mid-1960s and jumped higher in the aftermath of the 1973 oil crisis.[21] Subsequently, it dropped remarkably during the Volcker-disinflation period and somewhat settled around $2 - 2.5\%$ since the mid-1990s. Overall, visual inspection suggests that the estimated trend inflation is similar to others in the literature (e.g., Ireland, 2007; Cogley and Sbordone, 2008; Cogley et al., 2010; Ascari and Sbordone, 2014, among others). Moreover, it is also very similar to the estimate of trend inflation from the reduced-form piecewise linear model in Figure 5.

To the best of our knowledge, we are the first ones to estimate a DSGE model with time-varying steady state or trend inflation using full-system Bayesian estimation. Most papers in the literature either assume that steady state inflation is fixed (mostly at zero). One exception is Cogley and Sbordone (2008) who derive a generalized NKPC (GNKPC) with time-varying

---

[20]This stands in contrast to the evidence on passive monetary policy in the pre-Volcker period proposed by, among others, Clarida et al. (2000) and Lubik and Schorfheide (2004), that eventually led to non-fundamental sunspot fluctuations, which these authors argued to be one of the drivers of the Great Inflation. Nevertheless, Justiniano and Primiceri (2008) find that a model with active monetary policy and stochastic volatility fits the post-war U.S. data better than one with indeterminacy. In addition, Haque (2020) in an estimated NK model with exogenous time-varying inflation target finds that the evidence for indetermincay in the Great Inflation period dissapears once the model allows for time variation in the Federal Reserve's inflation target.

[21]The upward trend in inflation in the 1970s may be interpreted as "[...] *a systematic tendency for Federal Reserve policy to translate the short-run price pressures set off by adverse supply shocks into more persistent movements in the inflation rate itself - part of an effort by policymakers to avoid at least some of the contractionary impact those shocks would otherwise have had on the real economy.*" (Ireland, 2007, p. 1853)
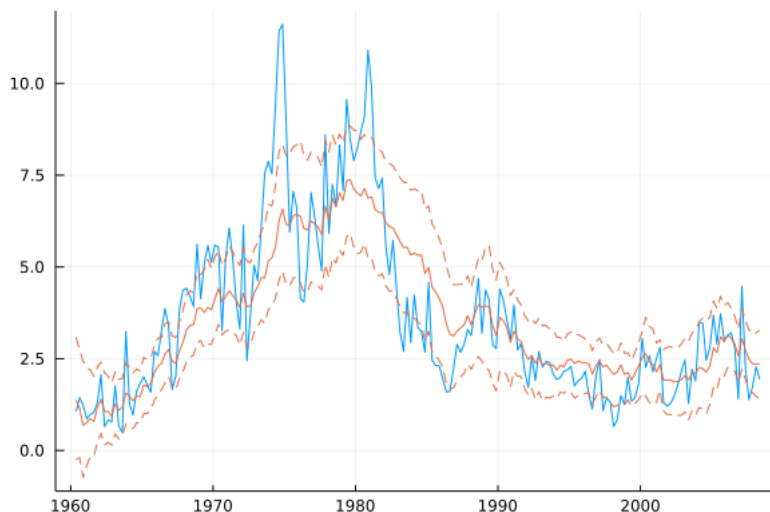
Figure 8: Inflation and trend inflation - GNK model.

trend inflation and document that inflation persistence results mainly from variation in the long-run trend component of inflation and that a purely forward-looking GNKPC fits the data quite well.

Figure 9 shows the estimated pattern of the time-varying standard deviations of the shocks. Despite the fact that we work with a much smaller model with respect to Justiniano and Primiceri (2008), the main conclusions remain very similar. First, the model accounts for the reduction in the volatility of U.S. macroeconomic variables, dubbed the Great Moderation, due to a substantial decrease in the volatility of exogenous disturbances. The pattern of stochastic volatility of monetary policy shocks is remarkably similar to that in Justiniano and Primiceri (2008) - our estimates capture the Volcker disinflation episode as well as the reduction in the volatility of monetary policy shocks during the Greenspan period. Other shocks also exhibit fluctuations in their standard deviations. The standard deviation of the technology shocks exhibit an inverted-U shaped pattern, which is consistent with the observed reduction in the volatility of GDP during the Great Moderation period. The volatility of preference shocks have also declined since the 1980s, possibly capturing the role that technological progress or financial innovations may have played in easing households' consumption smoothing. As in Cogley et al. (2010), we find an increase in the volatility of trend inflation shocks during the Great Inflation period and a subsequent decline in the post-Volcker period, although there is higher uncertainty around the estimates.

3.3.1 The long-run Phillips curve is non-linear and negatively sloped

Figure 10 plots the estimated LRPC from the structural model, expressed as percentage deviations from the zero inflation steady state, and compares it with the corresponding estimate coming from the BVAR. The structural model is able to capture the negative long-run empirical

30

Figure 9: Stochastic volatility of the structural shocks

relationship between output and inflation as observed in the data, both qualitatively and quantitatively. The estimated GNK LRPC and its 90% probability interval lies entirely within the 90% probability interval of the BVAR non-structural estimated LRPC. The LRPC is non-linear and downward sloping. There is a flat part of the curve for low level of trend inflation, but for trend inflation levels roughly above $3 - 4\%$ the slope increases sharply in absolute value with trend inflation. In terms of output losses, going from 2% to 4% inflation target causes an output loss of roughly about 0.6% per year. The effect is highly non-linear such that a 5% and a 6% inflation target would imply an output loss (relative to 2% target) of roughly 1.2% and 2% per year, respectively.

As is well-known in the literature, the negative steady state relationship between inflation and output in the GNK model is due to the negative effect of higher price dispersion on aggregate output. Higher trend inflation increases price dispersion by causing a greater difference between the price set by the resetting firms and the average price level. Higher price dispersion works like a negative aggregate productivity shock, as it increases the amount of input required to produce a given level of output, which in turn translates into an output loss (see Ascari, 2004; Yun, 2005; Ascari and Sbordone, 2014). Therefore, long-run superneutrality breaks down and a negative long run relationship emerges between inflation and output.

The Calvo model generally features larger costs of inflation relative to other price setting model, both time dependent (e.g., the Rotemberg or the Taylor models) or state dependent (menu cost models). Hence, the literature often stresses that the Calvo model tends to exaggerate the inflation costs, which are then considered as excessively large. It is important to stress that our analysis rebuts this interpretation. Figure 10 shows that the LRPC implied by the

Figure 10: Long-run Phillips curve: median (continuous line) and 90% probability interval (dashed lines) - comparison between VAR (blue) and GNK (black) estimates.

Calvo model is in line with the costs from the non-structural model of Section 2, at least up to trend inflation levels we historically observed in the U.S. data, i.e., up to 7-8% trend inflation.

3.3.2 The long-run output gap

As for the BVAR, a time-varying trend inflation generates a long-run wedge alongside the usual short-run wedge for the variables in the model. In the standard NK model the output gap (i.e., short-run wedge) is usually defined as deviation of output from its flexible price counterpart:

$$X_t = \frac{Y_t}{Y_t^n} \qquad \text{in logs} \rightarrow \qquad x_t = y_t - y_t^n. \tag{31}$$

Normally, the steady state output of the sticky price NK model and the one of its flexible price counterpart are the same.[22] If that is the case, (31) could also be written as:[23]

$$X_t = \frac{Y_t}{\bar{Y}_t} \frac{\bar{Y}_t}{Y_t^n} \qquad \text{in logs} \rightarrow \qquad x_t = \hat{y}_t - \hat{y}_t^n, \tag{32}$$

where a '*hat*' on the variables indicates log-deviations from the trend level, i.e., the steady state output level of the NK model and of the flexible price model, which are one and the same. According to the GNK model, instead, the trend $\bar{Y}_t$ for the sticky price model is different from the trend of the flexible price model, $\bar{Y}_t^n$. The latter does not depend on trend inflation, because flexible prices imply a vertical LRPC. Instead, trend inflation affects the long-run level of output under sticky prices.[24] The model then implies another wedge with respect to the flexible model counterpart: a *long-run output gap* arises from comparing the long-run behavior of the flexible price and sticky price models, as in the BVAR specification, see (5). One possibly useful decomposition is:

$$X_t = \frac{Y_t}{Y_t^n} = \frac{Y_t}{\bar{Y}_t} \frac{\bar{Y}_t}{\bar{Y}_t^n} \frac{\bar{Y}_t^n}{Y_t^n} \qquad \text{in logs} \rightarrow \qquad x_t = \hat{y}_t + \hat{\tilde{x}}_t - \hat{y}_t^n = \tilde{y}_t^{SR} + \tilde{y}_t^{LR}. \tag{33}$$

The output gap is divided into a short-run and a long-run component. As before in (32), the short-run component is the difference in the log-deviations of current output from its trend ($\bar{Y}_t$) and the flexible price output and its trend ($\bar{Y}_t^n$), that is $\tilde{y}_t^{SR} = \hat{y}_t - \hat{y}_t^n$. The long-run component is the log-deviation between the GNK-output trend ($\bar{Y}_t$) and the flexible price output trend ($\bar{Y}_t^n$). Assuming a vertical LRPC then $\tilde{y}_t^{LR} = 0$ and $x_t = \tilde{y}_t^{SR}$. However, in the GNK model $\tilde{y}_t^{LR} \neq 0$ because $\bar{Y}_t \neq \bar{Y}_t^n$. Further, the term $\overline{X}_t = (\bar{Y}_t/\bar{Y}_t^n)$ enters the Taylor rule (26) because the short-term interest rate reacts to deviation of the short-run output gap to its trend, that is: $(X_t/\overline{X}_t) = \frac{Y_t}{\bar{Y}_t} \frac{\bar{Y}_t^n}{Y_t^n}$, in logs, $x_t - \hat{\tilde{x}}_t = \hat{y}_t - \hat{y}_t^n = \tilde{y}_t^{SR}$.

Figure 11 plots the estimated long-run output gap implied by the GNK model and compares it to the one from the BVAR. The two estimates are very similar suggesting that the two models measure the actual costs of higher trend inflation in a consistent way.

---

[22]This is because usually NK models assume either zero inflation in steady state - so they do not even consider the long-run relationship between inflation and output - or full indexation of the prices that cannot be adjusted to some combination of past inflation and trend inflation. The latter assumption yields a vertical LRPC.

[23]The level of steady state output $\bar{Y}_t$ can be time varying if there is technological growth (either deterministic or stochastic) as in our model. In solving the model, variables are stationarized so that the steady state levels in stationarized variables are constant along a balanced growth path.

[24]This notion of a flexible price equilibrium complicates the analysis with respect to a non-structural one where one has just potential output as an unobservable to filter out. The somewhat "normative" notion of comparing the model with the flexible counterpart introduces other two non-observable variables: the flexible price output, Y_t^n and the flexible price trend output level, $\bar{Y}_t^n$.
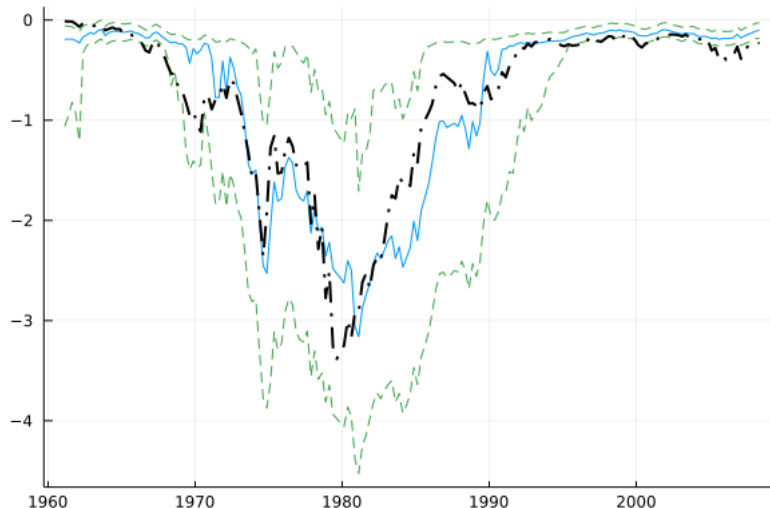
Figure 11: Comparison between long-run output gap estimates: VAR (blue) and GNK (black).

4 Conclusion

The relationship between inflation and economic activity in the long-run is of paramount importance for monetary policymaking as most central banks perceive price stability as the basis for long-term economic growth. The prominent tenet of the inflation targeting framework and of the monetary policy practice is the absence of a long-run trade-off between output and inflation. However, there are relatively few empirical investigations of this long-run relationship. This paper aims to develop an empirical methodology which is tailored to the purpose of providing more precise estimates of the LRPC.

We develop a Bayesian vector autoregression (BVAR) framework with stochastic trends, and provide a sophisticated trend-cycle decomposition of the data. A key methodological contribution is to generalize this BVAR-based trend-cycle decomposition to a piecewise linear model and show that both the likelihood function and the posterior distribution of the latent state variables can be derived analytically. Relative to existing studies, the framework allows for simultaneous estimation of both the short-run business cycle and the long-run trend components, such that the estimated LRPC is also consistent with the cyclical properties of the data.

Our empirical results call for a reassessment of the conventional wisdom according to which the LRPC is vertical and there is no trade-off between output and inflation in the long-run. First, they show that the long-run relationship between output and inflation is non-linear. While the LRPC is vertical for relatively low levels of trend inflation, it becomes negatively sloped when trend inflation is above a certain threshold estimated to be around 4%. Second, the slope of the LRPC is substantial: every percentage point increase in trend inflation above the threshold level is related to about 1% decrease in potential output per year. Third, it follows that trend inflation above the threshold carries material output costs. According to our

estimated model, the long-run output gap, which captures the deviation of potential output under positive trend inflation from its counterfactual level under zero trend inflation, has been on average about negative 2% per year during the Great Inflation. Fourth, neglecting the negative long-run relationship between output and inflation leads to mismeasurement of the size of the short-run output gap, hence to a misinterpretation of the business cycle dynamics. For example, not considering the negative slope of the LRPC leads to a more negative short-run output gap estimates in periods of high inflation during the Great Inflation, thereby overstating the cyclical component of output fluctuations.

Finally, a New Keynesian model generalized to admit time-varying trend inflation and estimated via particle filtering provides theoretical foundations to this reduced-form evidence coming from the BVAR. We show that the structural long-run Phillips Curve implied by the estimated New Keynesian model is not statistically different from the one implied by the reduced-form piecewise linear BVAR model. The structural model is able to provide, therefore, a theoretical rationale for the empirical non-structural results, suggesting that the costs of inflation due to price dispersion in the standard NK model are in line with the empirical evidence.

References

Abbritti, M., A. Consolo, and S. Weber (2021, December). Endogenous growth, downward wage rigidity and optimal inflation. Working Paper Series 2635, European Central Bank.

Adam, K., A. Alexandrov, and H. Weber (2023). Inflation distorts relative prices: Theory and evidence. Technical Report CEPR Discussion Paper No. 18088.

Adam, K. and H. Weber (2019). Optimal trend inflation. *American Economic Review 109*(2), 702–37.

Ait Lahcen, M., G. Baughman, S. Rabinovich, and H. van Buggenum (2021). Nonlinear unemployment effects of the inflation tax. Technical Report mimeo.

Arellano-Valle, R. B. and A. Azzalini (2006). On the unification of families of skew-normal distributions. *Scandinavian Journal of Statistics 33*(3), 561–574.

Ascari, G. (2004). Staggered prices and trend inflation: Some nuisances. *Review of Economic Dynamics 7*, 642–667.

Ascari, G., P. Bonomolo, and H. F. Lopes (2019). Walk on the wild side: Temporarily unstable paths and multiplicative sunspots. *American Economic Review 109*(5), 1805–42.

Ascari, G. and L. Fosso (2021, April). The Inflation Rate Disconnect Puzzle: On The International Component of Trend Inflation and The Flattening of The Phillips Curve. Technical report, Working paper.

Ascari, G. and A. M. Sbordone (2014). The macroeconomics of trend inflation. *Journal of Economic Literature 52*(3), 679–739.

Azzalini, A. (1985). A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics 12*(2), 171–178.

Benati, L. (2008). Investigating Inflation Persistence Across Monetary Regimes. *The Quarterly Journal of Economics 123*(3), 1005–1060.

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

Berentsen, A., G. Menzio, and R. Wright (2011). Inflation and unemployment in the long run. *American Economic Review 101*(1), 371–98.

Beyer, A. and R. E. Farmer (2007). Natural rate doubts. *Journal of Economic Dynamics and Control 31*(3), 797–825.

Calvo, G. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*, 383–398.

Carvalho, C. M., M. S. Johannes, H. F. Lopes, and N. G. Polson (2010, 02). Particle learning and smoothing. *Statistical Science 25*(1), 88–106.

Cavallo, A., F. Lippi, and K. Miyahara (2021). Inflation and misallocation in New Keynesian models. Technical Report 2023 ECB Sintra Forum on central banking.

Chen, H., F. Petralia, and H. F. Lopes (2010). Sequential Monte Carlo Estimation of DSGE Models. Working papers.

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

Clark, P. K. (1987). The cyclical component of U.S. economic activity. *The Quarterly Journal of Economics 102*(4), 797–814.

Cogley, T., G. E. Primiceri, and T. Sargent (2010). Inflation-gap persistence in the u.s. *American Economic Journal: Macroeconomics 2(1)*, 43–69.

Cogley, T. and T. J. Sargent (2005). Drifts and volatilities: Monetary policies and outcomes in the post WWII U.S. *Review of Economic Dynamics 8*(2), 262–302.

Cogley, T. and A. M. Sbordone (2008). Trend inflation, indexation and inflation persistence in the New Keynesian Phillips Curve. *American Economic Review 98*(5), 2101–2126.

Cooley, T. F. and G. D. Hansen (1989). The inflation tax in a real business cycle model. *American Economic Review 79*, 733–748.

Creal, D. (2007). Sequential Monte Carlo samplers for Bayesian DSGE models. Working papers, Vrije Universitiet Amsterdam.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017). Safety, Liquidity, and the Natural Rate of Interest. *Brookings Papers on Economic Activity 48*(1 (Spring), 235–316.

Drazen, A. (1981). Inflation and capital accumulation under a finite horizon. *Journal of Monetary Economics 8*(2), 247–260.

Fernández-Villaverde, J. and J. Rubio-Ramírez (2007). Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies 74*, 1059–1087.

Friedman, M. (1968). The role of monetary policy. *The American Economic Review 58*(1), 1–17.

Gordon, R. J. (1970). The Recent Acceleration of Inflation and Its Lessons for the Future. *Brookings Papers on Economic Activity 1*(1), 8–47.

Hamilton, J. D., E. S. Harris, J. Hatzius, and K. D. West (2016). The Equilibrium Real Funds Rate: Past, Present, and Future. *IMF Economic Review 64*(4), 660–707.

Haque, Q. (2020). Monetary Policy, Inflation Target and the Great Moderation: An Empirical Investigation. mimeo, The University of Western Australia, Centre for Applied Macroeconomic Analysis.

Harvey, A. C. and P. H. J. Todd (1983). Forecasting economic time series with structural and box-jenkins models: A case study. *Journal of Business & Economic Statistics 1*(4), 299–307.

Haug, A. A. and I. King (2014). In the long run, US unemployment follows inflation like a faithful dog. *Journal of Macroeconomics 41*(C), 42–52.

Herbst, E. and F. Schorfheide (2014). Sequential Monte Carlo Sampling For Dsge Models. *Journal of Applied Econometrics 29*(7), 1073–1098.

Herbst, E. P. and F. Schorfheide (2016). *Bayesian Estimation of DSGE Models*. Number 10612 in Economics Books. Princeton University Press.

Ireland, P. (2007). Changes in Federal Reserve's inflation target: Causes and consequences. *Journal of Money, Credit and Banking 39(8)*, 1851–1882.

Johannsen, B. K. and E. Mertens (2021). A time-series model of interest rates with the effective lower bound. *Journal of Money, Credit and Banking 53*(5), 1005–1046.

Justiniano, A. and G. E. Primiceri (2008). The Time-Varying Volatility of Macroeconomic Fluctuations. *American Economic Review 98(3)*, 604–641.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010). Investment Shocks and Business Cycles. *Journal of Monetary Economics 57(2)*, 132–145.

Karanassou, M. and H. Sala (2010). The US inflation-unemployment trade-off revisited: New evidence for policy-making. *Journal of Policy Modeling 32*(6), 758–777.

Karanassou, M., H. Sala, and D. Snower (2005). A reappraisal of the inflation-unemployment tradeoff. *European Journal of Political Economy 21*, 1–32.

Karanassou, M., H. Sala, and D. J. Snower (2010). Phillips Curves And Unemployment Dynamics: A Critique And A Holistic Perspective. *Journal of Economic Surveys 24*(1), 1–51.

King, R. G. (2008). The Phillips curve and U.S. macroeconomic policy : snapshots, 1958-1996. *Economic Quarterly 94*(Fall), 311–359.

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

Kreps, D. M. (1998). Anticipated utility and dynamic choice (1997). In D. P. Jacobs, E. Kalai, M. I. Kamien, and N. L. Schwartz (Eds.), *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997*, Econometric Society Monographs, pp. 242–274. Cambridge University Press.

Kuttner, K. (1994). Estimating potential output as a latent variable. *Journal of Business & Economic Statistics 12*(3), 361–68.

Laubach, T. and J. C. Williams (2003). Measuring the Natural Rate of Interest. *The Review of Economics and Statistics 85*(4), 1063–1070.

Liu, J. and M. West (2001). Combined parameters and state estimation in simulation-based filtering. In Springer (Ed.), *Sequential Monte Carlo Methods in Practice*, Chapter 10, pp. 197–223. New York: Arnaud Doucet and Nando de Freitas and Neil Gordon.

Lubik, T. and F. Schorfheide (2004). Testing for Indeterminacy: An Application to U.S. Monetary Policy. *American Economic Review 94(1)*, 190–217.

Maffei-Faccioli, N. (2020, October). Identifying the Sources of the Slowdown in Growth: Demand vs. Supply. 2020 Papers pma2978, Job Market Papers.

Mertens, E. (2016). Measuring the Level and Uncertainty of Trend Inflation. *The Review of Economics and Statistics 98*(5), 950–967.

Mertens, E. and J. M. Nason (2020). Inflation and professional forecast dynamics: An evaluation of stickiness, persistence, and volatility. *Quantitative Economics 11*(4), 1485–1520.

Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018). The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation*. *The Quarterly Journal of Economics 133*(4), 1933–1980.

Phelps, E. S. (1967). Phillips curves, expectations of inflation and optimal unemployment over time. *Economica 34*(135), 254–281.

Phillips, A. W. (1958). The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861-1957. *Economica 25*(100), 283–299.

Pitt, M. K. and N. Shephard (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association 94*(446), 590–599.

Planas, C., A. Rossi, and G. Fiorentini (2008, January). Bayesian Analysis of the Output Gap. *Journal of Business & Economic Statistics 26*, 18–32.

Primiceri, G. E. (2005). Time-varying structural vector autoregressions and monetary policy. *Review of Economic Studies 72*, 821–852.

Samuelson, P. A. and R. M. Solow (1960). Analytical aspects of anti-inflation policy. *The American Economic Review 50*(2), 177–194.

Sargan, J. D. (1964). Wages and prices in the united kingdom: a study in econometric methodology. *Econometric analysis for national economic planning 16*, 25–54.

Sheremirov, V. (2020). Price dispersion and inflation: New facts and theoretical implications. *Journal of Monetary Economics 114* (C), 59–70.

Sidrauski, M. (1967). Inflation and economic growth. *Journal of Political Economy 75*(6), 796–810.

Sims, C. A. (2002). Solving linear rational expectations models. *Computational Economics 20(1-2)*, 1–20.

Smets, F. and R. Wouters (2007). Shocks and frictions in US business cycles: A bayesian DSGE approach. *American Economic Review 97*(3), 586–606.

Solow, R. M. (1969). *Price Expectations and the Behavior of the Price Level.* Manchester, UK, Manchester University Press.

Stock, J. H. and M. W. Watson (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking 39*, 3–33.

Stock, J. H. and M. W. Watson (2016). Core Inflation and Trend Inflation. *The Review of Economics and Statistics 98*(4), 770–784.

Svensson, L. E. O. (2015). The Possible Unemployment Cost of Average Inflation below a Credible Target. *American Economic Journal: Macroeconomics 7*(1), 258–296.

Villani, M. (2009). Steady-state priors for vector autoregressions. *Journal of Applied Econometrics 24*(4), 630–650.

Watson, M. (1986). Univariate detrending methods with stochastic trends. *Journal of Monetary Economics 18*(1), 49–75.

Yun, T. (2005). Optimal monetary policy with relative price distortions. *American Economic Review 95*, 89–108.

A   Appendix: The piecewise linear Bayesian VAR

A.1   The time varying equilibrium VAR

Indicate with X_t a $n \times 1$ vector of observed variables at time t and with $\bar{X}_t$ the long-run component of X_t. The deviations $(X_t - \bar{X}_t)$ are described by the following stable VAR:

$$X_t - \bar{X}_t = A_1(X_{t-1} - \bar{X}_{t-1}) + A_2(X_{t-2} - \bar{X}_{t-2}) + ...A_p(X_{t-p} - \bar{X}_{t-p}) + \epsilon_t \qquad \text{(A1)}$$

with $\epsilon_t \sim N(\mathbf{0}, \Sigma_{\epsilon,t})$. We assume that the reduced-form shocks ϵ_t have stochastic volatility:

$$\Sigma_{\epsilon,t} = B^{-1}S_t(B^{-1}S_t)' \qquad \text{(A2)}$$

where

$$S_t = \begin{pmatrix} s_{1t} & 0 & 0 \\ 0 & s_{2t} & 0 \\ 0 & 0 & s_{3t} \end{pmatrix} \qquad \text{(A3)}$$

and

$$B = \begin{pmatrix} 1 & 0 & 0 \\ \beta_{21} & 1 & 0 \\ \beta_{31} & \beta_{32} & 1 \end{pmatrix} \qquad \text{(A4)}$$

Collect the elements in the main diagonal of S_t in the vector s_t. We follow the well established literature assuming:

$$\log s_t = \log s_{t-1} + \nu_t \qquad \nu_t \sim N(\mathbf{0}, \Sigma_\nu) \qquad \text{(A5)}$$

and we restrict Σ_ν to be diagonal.

The long-run component $\bar{X}_t$ depends on a $(q \times 1)$ vector of latent variables θ_t:

$$\bar{X}_t = \bar{D}_t + H_t\theta_t \qquad \text{(A6)}$$

and we assume that θ_t has the following dynamics

$$\theta_t = \bar{M}_t + \bar{G}_t\theta_{t-1} + \bar{P}_t\eta_t \qquad \text{(A7)}$$

with $\eta_t \sim N(\mathbf{0}, \Sigma_{\eta,t})$.

The first element of the latent vector θ_t is trend inflation $\bar{\pi}_t$, and the matrices in equations (A6) and (A7) depend on it. In particular, at each time t, conditioning on $\bar{\pi}_{t-1}$ we have two possibilities depending on trend inflation at time t:

$$\left(\bar{D}_t, H_t, \bar{M}_t, \bar{G}_t, \bar{P}_t\right) = \begin{cases} \left(\bar{D}_{1,t}, H_{1,t}, \bar{M}_{1,t}, \bar{G}_{1,t}, \bar{P}_{1,t}\right) & \text{if } \bar{\pi}_t \leq \tau \\ \left(\bar{D}_{2,t}, H_{2,t}, \bar{M}_{2,t}, \bar{G}_{2,t}, \bar{P}_{2,t}\right) & \text{if } \bar{\pi}_t > \tau \end{cases} \qquad \text{(A8)}$$

where τ is a threshold value. The subscript $1, t$ and $2, t$ indicate that the two groups of matrices do not have to be the same at each time t: the important assumption is that we always have a finite number of options (in our case two), so the model is piecewise linear.

A.2 Inference on the latent states θ_t

A.2.1 The state space form

Define $Y_t = X_t - A_1 X_{t-1} - A_2 X_{t-2} - ...A_p X_{t-p}$ and substitute equation (A6) in (A1) to get:

$$
\begin{aligned}
Y_t =& \bar{D}_t - A_1 \bar{D}_{t-1} - A_2 \bar{D}_{t-2}... - A_p \bar{D}_{t-p}+ \\
& H_t \theta_t - A_1 H_{t-1} \theta_{t-1} - A_2 H_{t-2} \theta_{t-2} - ... - A_p H_{t-p} \theta_{t-p} + \epsilon_t.
\end{aligned} \tag{A9}
$$

Define the latent vector $\vartheta_t = \begin{pmatrix} \theta'_t & \theta'_{t-1} & \cdots & \theta'_{t-p} \end{pmatrix}'$. Using equations (A7) and (A9) we can define the state space representation of our model:

$$
Y_t = D_t + F_t \vartheta_t + \epsilon_t \tag{A10}
$$
$$
\vartheta_t = M_t + G_t \vartheta_{t-1} + P_t \eta_t \tag{A11}
$$

The system can be written more explicitly as:

$$
\underset{n \times 1}{Y_t} = \left(\underset{n \times 1}{\bar{D}_t} - \sum_{i=1}^{p} \underset{n \times n}{A_i} \underset{n \times 1}{\bar{D}_{t-i}} \right) + \left(\underset{n \times q}{H_t} \quad \underset{n \times n\ n \times q}{-A_1 H_{t-1}} \quad \cdots \quad \underset{n \times n\ n \times q}{-A_p H_{t-p}} \right) \begin{pmatrix} \underset{q \times 1}{\theta_t} \\ \underset{q \times 1}{\theta_{t-1}} \\ \vdots \\ \underset{q \times 1}{\theta_{t-p}} \end{pmatrix} + \underset{n \times 1}{\epsilon_t} \tag{A12}
$$

and

$$
\begin{pmatrix} \underset{q \times 1}{\theta_t} \\ \underset{q \times 1}{\theta_{t-1}} \\ \vdots \\ \underset{q \times 1}{\theta_{t-p}} \end{pmatrix} = \begin{pmatrix} \underset{q \times 1}{\bar{M}_t} \\ \underset{q \times 1}{0} \\ \vdots \\ \underset{q \times 1}{0} \end{pmatrix} + \begin{pmatrix} \underset{q \times q}{\bar{G}_t} & \underset{q \times q}{0} & \cdots & \underset{q \times q}{0} \\ \underset{q \times q}{I} & \underset{q \times q}{0} & \cdots & \underset{q \times q}{0} \\ \vdots & \ddots & & \vdots \\ \underset{q \times q}{0} & \cdots & \underset{q \times q}{I} & \underset{q \times q}{0} \end{pmatrix} \begin{pmatrix} \underset{q \times 1}{\theta_{t-1}} \\ \underset{q \times 1}{\theta_{t-2}} \\ \vdots \\ \underset{q \times 1}{\theta_{t-p-1}} \end{pmatrix} + \begin{pmatrix} \underset{q \times h}{\bar{P}_t} \\ \underset{q \times h}{0} \\ \vdots \\ \underset{q \times h}{0} \end{pmatrix} \underset{h \times 1}{\eta_t} \tag{A13}
$$

Note that the state space is non-linear due to the dependency of the matrices of the state space form on one of the elements of θ_t, and to the presence of stochastic volatility. It is important to distinguish these two sources of non linearity: conditionally on the volatility processes, the model is piecewise linear and we present a fully adapted particle filter to estimate this class of models.

A.2.2 A fully adapted particle filter for ϑ_t

We derive the full conditional posterior distribution of the latent vector ϑ_t given all the parameters and the stochastic volatilities.

We tackle the curse of dimensionality described in the main text with a particle filtering strategy. Assume that at time $t-1$ we have a set of N particles $\left\{ \vartheta_{t-1}^{(i)} \right\}_{i=1}^{N}$ that approximate $p(\vartheta_{t-1}|Y_{1:t-1})$, and we want to get an analogous set of particles $\left\{ \vartheta_t^{(i)} \right\}_{i=1}^{N}$ to approximate $p(\vartheta_t|Y_{1:t})$. We use the following fully adapted particle filter:

```
┌─────────────────────────────────────────────────────────────────────────────┐
│                      Fully Adapted Particle Filter                          │
│                                                                             │
│ At  t − 1:   {ϑ_{t−1}^{(i)}}_{i=1}^N  approximate  p(ϑ_{t−1}|Y_{1:t−1})      │
│                                                                             │
│ 1) RESAMPLE                                                                 │
│    a) Compute  w̃_t^{(i)} ∝ p(Y_t|ϑ_{t−1}^{(i)}, Y_{1:t−1})                   │
│                                                                             │
│    b) Resample  {ϑ_{t−1}^{(i)}}_{i=1}^N  using  {w̃_t^{(i)}}_{i=1}^N  and     │
│       get  {ϑ̃_{t−1}^{(i)}}_{i=1}^N                                           │
│                                                                             │
│ 2) PROPAGATE                                                                │
│    Draw  ϑ_t^{(i)} ∼ p(ϑ_t|ϑ̃_{t−1}^{(i)}, Y_{1:t})                           │
│                                                                             │
└─────────────────────────────────────────────────────────────────────────────┘
```

In order to implement the filter we need to find two distributions: the predictive density $p\left(Y_t|\vartheta_{t-1}^{(i)}, Y_{1:t-1}\right)$ for the resample step, and the posterior density $p\left(\vartheta_t|\vartheta_{t-1}^{(i)}, Y_{1:t}\right)$ for the propagation step.

The predictive density $p\left(Y_t|\vartheta_{t-1}^{(i)}, Y_{1:t-1}\right)$

Given that our model is piece-wise linear, we write the predictive density as the sum of the two pieces:

$$p\left(Y_t|\vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) = p\left(Y_t|\bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) \Pr\left(\bar{\pi}_t \leq \tau|\vartheta_{t-1}^{(i)}, Y_{1:t-1}\right)$$
$$+ p\left(Y_t|\bar{\pi}_t > \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) \Pr\left(\bar{\pi}_t > \tau|\vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) \quad \text{(A14)}$$

The two addends on the right hand side of equation (A14) are analogous, so we concentrate on the first one, and a similar reasoning applies to the other one.

Start partitioning the latent vector as:

$$\vartheta_t = \begin{pmatrix} \bar{\pi}_t \\ \vartheta_t^x \end{pmatrix}$$

We are going to proceed in two steps: first consider the distribution:

$$p\left(Y_t, \vartheta_t^x|\bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) = \int_{\bar{\pi}_t \leq \tau} p\left(Y_t, \vartheta_t^x, \bar{\pi}_t|\vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) d\bar{\pi}_t \quad \text{(A15)}$$

that is a Unified Skew Normal (SUN) density, as defined by Arellano-Valle and Azzalini (2006). In order to derive it we start from the joint distribution of (ϑ_t, Y_t) under the assumption of a linear model: we set the matrices of the state space form equal to $(\bar{D}_{1,t}, H_{1,t}, \bar{M}_{1,t}, \bar{G}_{1,t}, \bar{P}_{1,t})$ independently from trend inflation, and we subsequently apply the truncation on $\bar{\pi}_t$. In case of linear and unrestricted model we have:

$$\begin{pmatrix} \vartheta_t \\ Y_t \end{pmatrix} \sim N(a_t, R_t) \quad \text{(A16)}$$

where:

$$
\underset{l+n\times 1}{a_t} = \begin{pmatrix} \underset{l\times 1}{M_t} \\ \underset{n\times 1}{D_t} + \underset{n\times 1}{F_t M_t} \end{pmatrix} + \begin{pmatrix} \underset{l\times l}{G_t} \\ \underset{n\times l}{F_t G_t} \end{pmatrix} \underset{l\times 1}{\vartheta_{t-1}^{(i)}} \tag{A17}
$$

and

$$
\underset{l+n\times l+n}{R_t} = \begin{pmatrix} \underset{l\times h}{P_t} & \underset{l\times n}{\mathbf{0}} \\ \underset{n\times h}{F_t P_t} & \underset{n\times n}{I} \end{pmatrix} \begin{pmatrix} \underset{h\times h}{\Sigma_\eta} & \underset{h\times n}{\mathbf{0}} \\ \underset{n\times h}{\mathbf{0}} & \underset{n\times n}{\Sigma_\epsilon} \end{pmatrix} \begin{pmatrix} \underset{h\times l}{P_t'} & \underset{h\times n}{P_t' F_t'} \\ \underset{n\times l}{\mathbf{0}} & \underset{n\times n}{I} \end{pmatrix} \tag{A18}
$$

Now consider that π_t is restricted below τ: the truncation of the first element below the threshold makes the distribution of the remaining elements a Unified Skew Normal:

$$
\left(\vartheta_t^x, Y_t | \bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1} \right) \sim SUN\left(\xi_t, \tau - \bar{\pi}_{t-1}, R_t \right) \tag{A19}
$$

where ξ_t is the $(l+n-1 \times 1)$ vector that contains all the elements in a_t except the first one.

The next step is to find the marginal distribution of $\left(Y_t | \bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1} \right)$. From the properties of the SUN we know that the marginal distribution is still a SUN. In particular, make the following partitions:

$$
R_t = \begin{pmatrix} \underset{1\times 1}{\Gamma} & \underset{1\times l+n-1}{\Delta'} \\ \underset{l+n-1\times 1}{\Delta} & \underset{l+n-1\times l+n-1}{\bar{\Omega}} \end{pmatrix}; \qquad \Delta = \begin{pmatrix} \underset{l-1\times 1}{\Delta_1} \\ \underset{n\times 1}{\Delta_2} \end{pmatrix}; \qquad \xi_t = \begin{pmatrix} \underset{l-1\times 1}{\xi_{1t}} \\ \underset{n\times 1}{\xi_{2t}} \end{pmatrix}
$$

and define:

$$
\Omega^* = \begin{pmatrix} \Gamma & \Delta_2' \\ \Delta_2 & \bar{\Omega}_{22} \end{pmatrix}
$$

where $\bar{\Omega}_{22}$ is the $n \times n$ lower block of $\bar{\Omega}$. We have that:

$$
\left(Y_t | \bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1} \right) \sim SUN\left(\xi_{2t}, \tau - \bar{\pi}_{t-1}, \Omega^* \right), \tag{A20}
$$

or, in explicit form:

$$
p\left(Y_t | \bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1} \right) = \phi_n\left(Y_t; \xi_{2t}, \bar{\Omega}_{22} \right) \frac{\Phi\left(\tau; \bar{\pi}_{t-1} + \Delta_2' \bar{\Omega}_{22}^{-1}\left(Y_t - \xi_{2t} \right), \Gamma - \Delta_2' \bar{\Omega}_{22}^{-1} \Delta_2 \right)}{\Phi\left(\tau; \bar{\pi}_{t-1}, \Gamma \right)} \tag{A21}
$$

where the denominator of the right hand side of equation (A21) is equal to $\Pr\left(\bar{\pi}_t \leq \tau | \vartheta_{t-1}^{(i)}, Y_{1:t-1} \right)$. Then, the first addend in equation (A14) is:

$$
\begin{aligned}
p\left(Y_t | \bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1} \right) \Pr\left(\bar{\pi}_t \leq \tau | \vartheta_{t-1}^{(i)}, Y_{1:t-1} \right) = \\
\phi_n\left(Y_t; \xi_{2t}, \bar{\Omega}_{22} \right) \Phi\left(\tau; \bar{\pi}_{t-1} + \Delta_2' \bar{\Omega}_{22}^{-1}\left(Y_t - \xi_{2t} \right), \Gamma - \Delta_2' \bar{\Omega}_{22}^{-1} \Delta_2 \right)
\end{aligned} \tag{A22}
$$

The second addend of (A14) can be derived with an analogous procedure.

The posterior distribution $p\left(\vartheta_t|\vartheta_{t-1}^{(i)}, Y_{1:t}\right)$

Also the posterior distribution can be written as the sum of two pieces:

$$p\left(\vartheta_t|\vartheta_{t-1}^{(i)}, Y_{1:t}\right) = p\left(\vartheta_t|\bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t}\right) \Pr\left(\bar{\pi}_t \leq \tau|\vartheta_{t-1}^{(i)}, Y_{1:t}\right)$$
$$+ p\left(\vartheta_t|\bar{\pi}_t > \tau, \vartheta_{t-1}^{(i)}, Y_{1:t}\right) \Pr\left(\bar{\pi}_t > \tau|\vartheta_{t-1}^{(i)}, Y_{1:t}\right) \qquad \text{(A23)}$$

As we did for the predictive density, we concentrate on the first row of (A23) (an analogous reasoning will apply to the second row).

Operate the following partitions:

$$a_t = \begin{pmatrix} a_{1t} \\ l \times 1 \\ a_{2t} \\ n \times 1 \end{pmatrix}; \qquad R_t = \begin{pmatrix} R_{11} & R_{12} \\ l \times l & l \times n \\ R_{21} & R_{22} \\ n \times l & n \times n \end{pmatrix}.$$

The posterior distribution of ϑ_t, conditioning on trend inflation below the threshold is a multivariate truncated normal:

$$\left(\vartheta_t|\bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t}\right) \sim TN\left(m_t, C_t; \bar{\pi}_t \leq \tau\right) \qquad \text{(A24)}$$

where:

$$m_t = a_{1t} + R_{12}R_{22}^{-1}\left(Y_t - a_{2t}\right) \qquad \text{(A25)}$$
$$C_t = R_{11} - R_{12}R_{22}^{-1}R_{21} \qquad \text{(A26)}$$

Finally, we need to compute the second term in the first line of equation (A23):

$$\Pr\left(\bar{\pi}_t \leq \tau|\vartheta_{t-1}^{(i)}, Y_{1:t}\right) \propto p\left(Y_t|\bar{\pi}_t \leq \tau, \vartheta_{t-1}^{(i)}, Y_{1:t-1}\right) \Pr\left(\bar{\pi}_t \leq \tau|\vartheta_{t-1}^{(i)}, Y_{1:t-1}\right)$$
$$= \phi_n\left(Y_t; \xi_{2t}, \bar{\Omega}\right) \Phi\left(\tau; \bar{\pi}_{t-1} + \Delta_2'\bar{\Omega}_{22}^{-1}\left(Y_t - \xi_{2t}\right), \Gamma - \Delta_2'\bar{\Omega}_{22}^{-1}\Delta_2\right) \qquad \text{(A27)}$$

that is exactly equation (A22).

A.3   Stochastic volatility

We now augment our algorithm to estimate the time varying standard deviations of the shocks in (A1). As in the Section above, given a set of particles $\left\{\vartheta_{t-1}^{(i)}, \log s_{t-1}^{(i)}\right\}_{i=1}^N$ that approximate the joint distribution of $(\vartheta_{t-1}, \log s_{t-1})$, we want to get a new set $\left\{\vartheta_t^{(i)}, \log s_t^{(i)}\right\}_{i=1}^N$ to approximate

$$p\left(\vartheta_t, \log s_t|Y_{1:t}\right)$$

where we omitted the dependencies on all the parameters to simplify the notation. Since in the previous Section we derived the posterior distribution of ϑ_t conditional on $\log s_t$, it is convenient to write the full conditional posterior at time t as:

$$p\left(\vartheta_t, \log s_t|\vartheta_{t-1}, \log s_{t-1} Y_{1:t}\right) = \underbrace{p\left(\vartheta_t|\log s_t, \vartheta_{t-1}, \log s_{t-1}, Y_{1:t}\right)}_{\text{Full conditional posterior}} \underbrace{p\left(\log s_t|\vartheta_{t-1}, \log s_{t-1}, Y_{1:t}\right)}_{\text{Blind proposal}}.$$
$$\text{(A28)}$$

We can get draws using an importance distribution that operates in two steps: we first use equation (3) to get particles of $\log s_t$: this is called a "blind" proposal because it is not conditioned on observed data. Then, we can condition on these draws and get values for ϑ_t using

the full conditional distribution that we derived analytically above.

Then, our particle filter is "partially" adapted: we use the so called "Rao-Blackwellization" to improve the efficiency of the estimator. With respect to the a fully adapted particle filter, we need to compute the final weights attached to each particle to get the approximation of the target distribution.

The complete algorithm is:

<div style="border:1px solid black; padding:1em;">

<div align="center">

Partially Adapted Particle Filter

</div>

At $t - 1$: the set of particles $\left\{\vartheta_{t-1}^{(i)}, \log s_{t-1}^{(i)}\right\}_{i=1}^{N}$ with corresponding weights $\left\{w_{t-1}^{(i)}\right\}_{i=1}^{N}$ approximate $p\left(\vartheta_{t-1}, \log s_{t-1} | Y_{1:t-1}\right)$

1) RESAMPLE

   a) Compute $\tilde{w}_t^{(i)} \propto w_{t-1}^{(i)} \; p\left(Y_t | \vartheta_{t-1}^{(i)}, g\left(\log s_{t-1}^{(i)}\right), Y_{1:t-1}\right)$

   b) Resample $\left\{\vartheta_{t-1}^{(i)}, \log s_{t-1}^{(i)}\right\}_{i=1}^{N}$ using $\left\{\tilde{w}_t^{(i)}\right\}_{i=1}^{N}$

Let the new particles be $\left\{\tilde{\vartheta}_{t-1}^{(i)}, \log \tilde{s}_{t-1}^{(i)}\right\}_{i=1}^{N}$.

2) PROPAGATE

   a) Draw $\log s_t^{(i)} \sim N\left(\log \tilde{s}_{t-1}^{(i)}, \Sigma_\nu\right)$

   b) Draw $\vartheta_t^{(i)} \sim p\left(\vartheta_t | \log s_t^{(i)}, \tilde{\vartheta}_{t-1}^{(i)}, Y_{1:t}\right)$

3) NEW WEIGHTS

   Compute $w_t^{(i)} \propto \dfrac{p\left(Y_t | \tilde{\vartheta}_{t-1}^{(i)}, \log s_t^{(i)}, Y_{1:t-1}\right)}{p\left(Y_t | \tilde{\vartheta}_{t-1}^{(i)}, g\left(\log \tilde{s}_{t-1}^{(i)}\right), Y_{1:t-1}\right)}$

</div>

A.4   The model for the long run

A.4.1   The dynamics

The vector $\bar{X}_t$ contains three variables: potential output $\bar{y}_t$, trend inflation $\bar{\pi}_t$ and the long-run nominal interest rate $\bar{i}_t$. We assume that potential output is the sum of a trend component and a function of trend inflation:

$$\bar{y}_t = y_t^* + \delta\left(\bar{\pi}_t\right) \tag{A29}$$

where the dynamics of the trend are:

$$y_t^* = y_{t-1}^* + g_{t-1} + \eta_t^y \tag{A30}$$

$$g_t = g_{t-1} + \eta_t^g \tag{A31}$$

and the function $\delta(\bar{\pi})$ is:

$$\delta(\bar{\pi}_t) = \begin{cases} k_1 \bar{\pi}_t & \text{if } \bar{\pi}_t \leq \tau \\ k_2 \bar{\pi}_t + c_k & \text{if } \bar{\pi}_t > \tau. \end{cases} \tag{A32}$$

Imposing continuity in the piecewise linear function, we have $c_k = (k_1 - k_2)\tau$. We assume that trend inflation follows a random walk:

$$\bar{\pi}_t = \bar{\pi}_{t-1} + \eta_t^\pi \tag{A33}$$

and the nominal interest rate in the long run is described by a Fisher equation, so it is equal to the sum of trend inflation and the long-run real interest rate:

$$\bar{i}_t = \bar{\pi}_t + cg_t + z_t. \tag{A34}$$

Following [Laubach and Williams (2003)](#), the long-run real interest rate is assumed to be a linear function of the growth rate of potential output and a random walk component z_t that captures all the slow moving trends that are potentially relevant but not directly present in the model:

$$z_t = z_{t-1} + \eta_t^z \tag{A35}$$

A.4.2   The state space form

Take the first difference of potential output in equation ([A29](#)):

$$\bar{y}_t = \bar{y}_{t-1} + g_{t-1} + \delta(\bar{\pi}_t) - \delta(\bar{\pi}_{t-1}) + \eta_t^y \tag{A36}$$

and note that taking into account the dynamics of trend inflation, $\delta(\bar{\pi}_t)$ is:

$$\delta(\bar{\pi}_t) = \begin{cases} k_1 \bar{\pi}_{t-1} + k_1 \eta_t^\pi & \text{if } \bar{\pi}_t \leq \tau \\ k_2 \bar{\pi}_{t-1} + k_2 \eta_t^\pi + c_k & \text{if } \bar{\pi}_t > \tau. \end{cases} \tag{A37}$$

Then, in computing the difference $\delta(\bar{\pi}_t) - \delta(\bar{\pi}_{t-1})$ we need to distinguish four possible cases:

$$\delta(\bar{\pi}_t) - \delta(\bar{\pi}_{t-1}) = \begin{cases} k_1 \eta_t^\pi & \text{if } \bar{\pi}_{t-1} \leq \tau \text{ and } \bar{\pi}_t \leq \tau \\ (k_2 - k_1)\pi_{t-1} + c_k + k_2 \eta_t^\pi & \text{if } \bar{\pi}_{t-1} \leq \tau \text{ and } \bar{\pi}_t > \tau \\ (k_1 - k_2)\pi_{t-1} - c_k + k_1 \eta_t^\pi & \text{if } \bar{\pi}_{t-1} > \tau \text{ and } \bar{\pi}_t \leq \tau \\ k_2 \eta_t^\pi & \text{if } \bar{\pi}_{t-1} > \tau \text{ and } \bar{\pi}_t > \tau \end{cases} \tag{A38}$$

We are now ready to write our state space in matrix form as in equations ([A6](#)) and ([A7](#)). Define the vector $\theta_t = \begin{pmatrix} \bar{\pi}_t & \bar{y}_t & g_t & z_t \end{pmatrix}'$, and write equation ([A6](#)) as:

$$\begin{pmatrix} \bar{y}_t \\ \bar{\pi}_t \\ \bar{i}_t \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & c & 1 \end{pmatrix} \begin{pmatrix} \bar{\pi}_t \\ \bar{y}_t \\ g_t \\ z_t \end{pmatrix} \tag{A39}$$

where $\bar{D}_t$ is equal to zero and $H_t = H$ is a constant matrix.

In our case the matrices in equation ([A7](#)) depend on trend inflation at time $t-1$ and time t: we have to distinguish the four cases highlighted above in equation ([A38](#)).

Case when $\bar{\pi}_{t-1} \leq \tau$ and $\bar{\pi}_t \leq \tau$.   The dynamics of potential output are described by the following equation:

$$\bar{y}_t = \bar{y}_{t-1} + g_{t-1} + k_1 \eta_t^\pi + \eta_t^y \tag{A40}$$

so the system in matrix form is:

$$
\begin{pmatrix} \bar{\pi}_t \\ \bar{y}_t \\ g_t \\ z_t \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \bar{\pi}_{t-1} \\ \bar{y}_{t-1} \\ g_{t-1} \\ z_{t-1} \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & 0 \\ k_1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \eta_t^\pi \\ \eta_t^y \\ \eta_t^g \\ \eta_t^z \end{pmatrix} . \tag{A41}
$$

Case when $\bar{\pi}_{t-1} \leq \tau$ and $\bar{\pi}_t > \tau$. Potential output in this case follows:

$$
\bar{y}_t = \bar{y}_{t-1} + g_{t-1} + (k_2 - k_1)\pi_{t-1} + c_k + k_2\eta_t^\pi + \eta_t^y \tag{A42}
$$

and equation (A7) becomes:

$$
\begin{pmatrix} \bar{\pi}_t \\ \bar{y}_t \\ g_t \\ z_t \end{pmatrix} = \begin{pmatrix} 0 \\ c_k \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & 0 \\ (k_2 - k_1) & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \bar{\pi}_{t-1} \\ \bar{y}_{t-1} \\ g_{t-1} \\ z_{t-1} \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & 0 \\ k_2 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \eta_t^\pi \\ \eta_t^y \\ \eta_t^g \\ \eta_t^z \end{pmatrix} . \tag{A43}
$$

Case when $\bar{\pi}_{t-1} > \tau$ and $\bar{\pi}_t \leq \tau$. The equation for the dynamics of potential output is:

$$
\bar{y}_t = \bar{y}_{t-1} + g_{t-1} + (k_1 - k_2)\pi_{t-1} - c_k + k_1\eta_t^\pi + \eta_t^y \tag{A44}
$$

and the system becomes:

$$
\begin{pmatrix} \bar{\pi}_t \\ \bar{y}_t \\ g_t \\ z_t \end{pmatrix} = \begin{pmatrix} 0 \\ -c_k \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & 0 \\ (k_1 - k_2) & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \bar{\pi}_{t-1} \\ \bar{y}_{t-1} \\ g_{t-1} \\ z_{t-1} \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & 0 \\ k_1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \eta_t^\pi \\ \eta_t^y \\ \eta_t^g \\ \eta_t^z \end{pmatrix} . \tag{A45}
$$

Case when $\bar{\pi}_{t-1} > \tau$ and $\bar{\pi}_t > \tau$. For this last case the dynamics of potential output follow:

$$
\bar{y}_t = \bar{y}_{t-1} + g_{t-1} + k_2\eta_t^\pi + \eta_t^y \tag{A46}
$$

and the system is:

$$
\begin{pmatrix} \bar{\pi}_t \\ \bar{y}_t \\ g_t \\ z_t \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \bar{\pi}_{t-1} \\ \bar{y}_{t-1} \\ g_{t-1} \\ z_{t-1} \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & 0 \\ k_2 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \eta_t^\pi \\ \eta_t^y \\ \eta_t^g \\ \eta_t^z \end{pmatrix} . \tag{A47}
$$

A.5 Inference on the parameters

In our particle filtering strategy we estimate the parameters combining two approaches: the posterior distribution of τ is approximated through a mixture of Normal densities as in Liu and West (2001); for all the other parameters we use the Particle Learning scheme by Carvalho et al. (2010), which is based on the analytical availability of sufficient statistics that characterize the posterior distributions. For most of the parameters the derivation of these sufficient statistics

is standard. Then, we only describe the inference on the parameters in the model for the long run, except for the variances of the shocks which are standard conjugate Inverse Gamma.

A.5.1 The posterior distribution of parameters in the matrix H

In our case $\bar{D}_t$ is equal to zero and the matrix H is constant. Then, equation (A6) is:

$$\bar{X}_t = H\theta_t. \tag{A48}$$

Rearranging equation (A1) and substituting the definition of $\bar{X}$ we have:

$$Y_t = H\theta_t - A_1 H\theta_{t-1} - A_2 H\theta_{t-2} - ...A_p H\theta_{t-p} + \epsilon_t. \tag{A49}$$

where Y_t has been defined in Section A.2.1.

Then, indicating with $\vec{H}$ the vectorized matrix H we obtain the regression:

$$Y_t = X_t^H \vec{H} + \epsilon_t \tag{A50}$$

where

$$X_t^H = \left[\theta_t' \otimes I_n - \theta_{t-1}' \otimes A_1 - \theta_{t-2}' \otimes A_2 - ...\theta_{t-p}' \otimes A_p \right]. \tag{A51}$$

Note that usually $\vec{H}$ contains some known coefficients and some unknown coefficients that we want to estimate. In our specific case we only have c as unknown: the other coefficients are all ones or zeros. As general practice, collect all the known coefficients in $\vec{H}_K$ and the unknown coefficients in $\vec{H}_U$. With a similar notation indicate with $X_{t,K}^H$ and $X_{t,U}^H$ the corresponding columns of X_t^H. We can write:

$$Y_t - X_{t,K}^H \vec{H}_K = X_{t,U}^H \vec{H}_U + \epsilon_t \tag{A52}$$

$$Y_t^H = X_{t,U}^H \vec{H}_U + \epsilon_t \tag{A53}$$

Equation (A53) is our regression: using Gaussian priors for the coefficients in $\vec{H}_U$, the inference is obtained as in standard Bayesian regression models.

A.5.2 The posterior distribution of slopes k_1 and k_2

First, define $Y_t^G = \bar{y}_t - \bar{y}_{t-1} - g_{t-1}$ and $\vec{k} = \begin{pmatrix} k_1 & k_2 \end{pmatrix}'$. Then, we define the vector X_t^G distinguishing the usual four possible cases:

$$X_t^G = \begin{cases} \begin{pmatrix} \bar{\pi}_t - \bar{\pi}_{t-1} & 0 \end{pmatrix} & \text{if } \bar{\pi}_{t-1} \leq \tau \text{ and } \bar{\pi}_t \leq \tau \\ \begin{pmatrix} \tau - \bar{\pi}_{t-1} & \bar{\pi}_t - \tau \end{pmatrix} & \text{if } \bar{\pi}_{t-1} \leq \tau \text{ and } \bar{\pi}_t > \tau \\ \begin{pmatrix} \bar{\pi}_t - \tau & \tau - \bar{\pi}_{t-1} \end{pmatrix} & \text{if } \bar{\pi}_{t-1} > \tau \text{ and } \bar{\pi}_t \leq \tau \\ \begin{pmatrix} 0 & \bar{\pi}_t - \bar{\pi}_{t-1} \end{pmatrix} & \text{if } \bar{\pi}_{t-1} > \tau \text{ and } \bar{\pi}_t > \tau \end{cases} \tag{A54}$$

Finally, we can write the dynamics of potential output as:

$$Y_t^G = X_t^G \vec{k} + \eta_t^y. \tag{A55}$$

The equation above is a regression with coefficients $\vec{k}$: using a Gaussian prior for $\vec{k}$ we easily obtain a conjugate posterior distribution.

B   Appendix: The GNK model

Households. The first-order conditions with respect to consumption, labor supply and bond holdings are:

$$\lambda_t = \frac{d_t}{C_t - hC_{t-1}}, \tag{B56}$$

$$\frac{W_t}{P_t} = \frac{d_n d_t N_t^\varphi}{\lambda_t},$$

$$1 = E_t \frac{\beta \lambda_{t+1}}{\lambda_t} \frac{R_t}{\pi_{t+1}}, \tag{B57}$$

where λ_t is the marginal utility of consumption, and $\pi_t = \frac{P_t}{P_{t-1}}$ is the gross inflation rate.

Firms. In each period t, a final good, Y_t, is produced by a perfectly competitive representative final-good firm, by combining a continuum of intermediate inputs, $Y_{i,t}$, $i \in [0,1]$, via the technology

$$Y_t = \left[\int_0^1 Y_{i,t}^{\frac{\varepsilon-1}{\varepsilon}} \, di \right]^{\frac{\varepsilon}{\varepsilon-1}}, \tag{B58}$$

where $\varepsilon > 1$ is the elasticity of substitution among intermediate inputs. The first-order condition for profit maximization yields the final-good firm's demand for intermediate good i

$$Y_{i,t} = \left(\frac{P_{i,t}}{P_t} \right)^{-\varepsilon} Y_t. \tag{B59}$$

The final-good market clearing condition is given by $Y_t = C_t$.

We assume that the price of a firm that cannot change the price is automatically indexed to past-inflation with a degree χ, that is $P_{i,t} = P_{i,t-1} \left(\frac{P_{t-1}}{P_{t-2}} \right)^\chi = P_{i,t-1} (\pi_{t-1})^\chi$, where hence $\pi_t \equiv \frac{P_t}{P_{t-1}}$. Hence if a firm fix $P_{i,t}^*$ today and will not be able to change it in the future then the price evolves accordingly to $P_{i,t+1} = P_{i,t}^* (\pi_t)^\chi$, $P_{i,t+2} = P_{i,t}^* (\pi_t)^\chi (\pi_{t+1})^\chi$, $P_{i,t+j} = P_{i,t}^* (\pi_t)^\chi (\pi_{t+1})^\chi ... (\pi_{t+j-1})^\chi = P_{i,t}^* \pi_{t|t-1+j}$, where

$$
\begin{aligned}
\pi_{t|t+j-1} &= \frac{P_t}{P_{t-1}} \times \frac{P_{t+1}}{P_t} \times ... \times \frac{P_{t+j-1}}{P_{t+j-2}} \quad &&\text{for } j \geqslant 1 \quad \text{and} \quad \pi_{t|t} = \frac{P_t}{P_{t-1}} = \pi_t \\
&= 1 &&\text{for } j = 0.
\end{aligned}
\tag{B60}
$$

The intermediate goods producers face a constant probability, $0 < (1-\theta) < 1$, of being able to adjust prices to a new optimal one, $P_{i,t}^*$. Thus, to maximize expected discounted profit they solve the following problem

$$E_t \sum_{j=0}^\infty \theta^j \beta^j \frac{\lambda_{t+j}}{\lambda_t} \left[\frac{P_{i,t}^* \pi_{t|t+j-1}^\chi}{P_{t+j}} Y_{i,t+j} - \frac{W_{t+j}}{P_{t+j}} \left[\frac{Y_{i,t+j}}{A_{t+j}} \right]^{\frac{1}{1-\alpha}} \right]$$

$$s.t. \qquad Y_{i,t+j} = \left[\frac{P_{i,t}^* \pi_{t|t+j-1}^\chi}{P_{t+j}} \right]^{-\varepsilon} Y_{t+j},$$

Defining "average" marginal cost as $MC_t = \frac{A_t^{\frac{1}{1-\alpha}}}{1-\alpha} \frac{W_t}{P_t} Y_t^{\frac{\alpha}{1-\alpha}}$, first order condition for the opti-

mized relative price $p_t^*(= \frac{P_{i,t}^*}{P_t})$ can be written as

$$(p_t^*)^{1+\frac{\varepsilon\alpha}{1-\alpha}} = \frac{\varepsilon}{\varepsilon-1} \frac{E_t \sum_{j=0}^{\infty}(\theta\beta)^j \lambda_{t+j} Y_{t+j} \left[\frac{\pi_{t|t+j-1}^{\chi}}{\pi_{t+1|t+j}}\right]^{-\frac{\varepsilon}{1-\alpha}} MC_{t+j}}{E_t \sum_{j=0}^{\infty}(\theta\beta)^j \lambda_{t+j} Y_{t+j} \left[\frac{\pi_{t|t+j-1}^{\chi}}{\pi_{t+1|t+j}}\right]^{1-\varepsilon}}. \tag{B61}$$

The aggregate price level, $P_t = \left[\int_0^1 P_{i,t}^{1-\varepsilon}di\right]^{\frac{1}{1-\varepsilon}}$, evolves according to

$$p_t^* = \left[\frac{1 - \theta\pi_{t-1}^{(1-\varepsilon)\chi}\pi_t^{\varepsilon-1}}{1-\theta}\right]^{\frac{1}{1-\varepsilon}}. \tag{B62}$$

Lastly, define price dispersion $s_t \equiv \int_0^1 (\frac{P_{i,t}}{P_t})^{-\varepsilon}di$. Under the Calvo price mechanism, the above expression can be written recursively as

$$s_t = (1-\theta)(p_t^*)^{-\varepsilon} + \theta\pi_{t-1}^{-\varepsilon\chi}\pi_t^{\varepsilon}s_{t-1}. \tag{B63}$$

Recursive formulation of the optimal price-setting equation. The joint dynamics of the optimal reset price and inflation can be compactly described by rewriting the first-order condition for the optimal price in a recursive formulation as follows:

$$(p_t^*)^{1+\frac{\varepsilon\alpha}{1-\alpha}} = \frac{\varepsilon}{(\varepsilon-1)}\frac{\psi_t}{\phi_t}, \tag{B64}$$

where ψ_t and ϕ_t are auxiliary variables that allow one to rewrite the infinite sums that appear in the numerator and denominator of the above equation in recursive formulation:

$$\psi_t = MC_t Y_t \lambda_t + \theta\beta\pi_t^{-\frac{\varepsilon\chi}{1-\alpha}} E_t \left[\pi_{t+1}^{\frac{\varepsilon}{1-\alpha}}\psi_{t+1}\right], \tag{B65}$$

and

$$\phi_t = Y_t\lambda_t + \theta\beta\pi_t^{\chi(1-\varepsilon)} E_t \left[\pi_{t+1}^{\varepsilon-1}\phi_{t+1}\right]. \tag{B66}$$

Note that in defining these two auxiliary variables, we used the definition $\lambda_t = \frac{d_t}{C_t - hC_{t-1}} = \frac{d_t}{Y_t - hY_{t-1}}$.

Monetary Policy. The central bank's policy is described by the following Taylor rule

$$\ln R_t = \rho \ln R_{t-1} + (1-\rho)\ln \bar{R}_t + (1-\rho)\psi_{\pi}\ln\left(\frac{\pi_t}{\bar{\pi}_t}\right) + (1-\rho)\psi_x \ln\left(\frac{X_t}{\bar{X}_t}\right) + (1-\rho)\psi_{\Delta y}\ln\left(\frac{g_t^y}{g^y}\right) + \sigma_{r,t}\epsilon_{r,t}, \tag{B67}$$

where X_t is the output gap defined as the deviation of output from its natural (or flexible price) level, $\bar{X}_t$ is the steady state output gap, g_t^y is the growth rate of output, $g^y = \bar{g}$ is the steady state growth rate of output, and $\epsilon_{r,t}$ is an i.i.d. $N(0,1)$ monetary policy shock with time-varying standard deviation $\sigma_{r,t}$. The parameters ψ_{π}, ψ_x and $\psi_{\Delta y}$ govern the central bank's responses to the inflation gap, output gap and output growth, respectively. Here $\bar{\pi}_t$ denotes trend inflation, which is the central's banks (time-varying) inflation target and follows a unit root process

$$\ln \bar{\pi}_t = \ln \bar{\pi}_{t-1} + \sigma_{\bar{\pi},t}\epsilon_{\bar{\pi},t}, \tag{B68}$$

where $\epsilon_{\bar{\pi},t}$ is i.i.d. $N(0,1)$ and $\sigma_{\bar{\pi},t}$ denotes time-varying standard deviation of the inflation target shock.

By considering flexible prices, the law of motion for Y_t^n is given by

$$\left(\frac{Y_t^n}{A_t}\right)^{\frac{1+\varphi}{1-\alpha}} = \frac{(\varepsilon-1)(1-\alpha)}{\varepsilon d_n} + h\left(\frac{Y_t^n}{A_t}\right)^{\frac{\varphi+\alpha}{1-\alpha}}\frac{Y_{t-1}^n}{A_t}. \tag{B69}$$

B.1 Final equations of the log-linearized GNK model

We log-linearize the equilibrium conditions arising from the DSGE model described above around a steady state characterized by a shifting trend inflation and derive the following set of log-linearized equations. A separate online supplement provides further details including the full set of non-linear equilibrium conditions of the model, the steady state and the details on the log-linearization around the time-varying steady state. As in Cogley and Sbordone (2008), the steady state of the model is time-varying because of drifts in trend inflation. As such, care must be taken when log-linearizing the model. As in Cogley and Sbordone (2008) and following Kreps (1998), we assume that agents treat drifting parameters as if they would remain constant at the current level going forward in time. In what follows, hatted variables denote log-deviations of stationary variables from their steady state values, which are, in turn, denoted with a bar.[25]

$$\widehat{\lambda}_t = -\left(\frac{h}{\overline{g}-h}\right)\widehat{g}_t - \left(\frac{\overline{g}}{\overline{g}-h}\right)\widehat{Y}_t + \left(\frac{h}{\overline{g}-h}\right)\widehat{Y}_{t-1} - \left(\frac{h}{\overline{g}-h}\right)\widehat{g}_t^{\overline{Y}} + \widehat{d}_t \tag{1L}$$

$$\widehat{w}_t = \widehat{d}_t + \varphi\widehat{N}_t - \widehat{\lambda}_t \tag{2L}$$

$$\widehat{\lambda}_t = \widehat{R}_t - E_t\widehat{\pi}_{t+1} + E_t\widehat{\lambda}_{t+1}, \tag{3L}$$

$$0 = \left[1 - \theta\overline{\pi}_t^{(1-\chi)(\varepsilon-1)}\right]\widehat{p}_t^* - \theta\overline{\pi}_t^{(1-\chi)(\varepsilon-1)}\left[\widehat{\pi}_t - \chi\widehat{\pi}_{t-1} + \chi\widehat{g}_t^{\overline{\pi}}\right] \tag{4L}$$

$$\left(1 + \frac{\varepsilon\alpha}{1-\alpha}\right)\widehat{p}_t^* = \widehat{\psi}_t - \widehat{\phi}_t \tag{5L}$$

$$\widehat{\psi}_t = \left[1 - \theta\beta\overline{\pi}_t^{\frac{\varepsilon(1-\chi)}{1-\alpha}}\right]\left(\widehat{mc}_t + \widehat{Y}_t + \widehat{\lambda}_t\right) + \theta\beta\overline{\pi}_t^{\frac{\varepsilon(1-\chi)}{1-\alpha}}E_t\left(\widehat{\psi}_{t+1} + \frac{\varepsilon}{1-\alpha}\widehat{\pi}_{t+1} - \frac{\varepsilon\chi}{1-\alpha}\widehat{\pi}_t\right) \tag{6L}$$

$$\widehat{\phi}_t = \left[1 - \theta\beta\overline{\pi}_t^{(1-\chi)(\varepsilon-1)}\right]\left(\widehat{Y}_t + \widehat{\lambda}_t\right) + \theta\beta\overline{\pi}_t^{(1-\chi)(\varepsilon-1)}E_t\left(\widehat{\phi}_{t+1} + (\varepsilon-1)\widehat{\pi}_{t+1} + \chi(1-\varepsilon)\widehat{\pi}_t\right) \tag{7L}$$

$$\widehat{N}_t = \widehat{s}_t + \left(\frac{1}{1-\alpha}\right)\widehat{Y}_t \tag{8L}$$

$$\widehat{s}_t = \left[-\frac{\varepsilon}{1-\alpha}\left(1 - \theta\overline{\pi}_t^{\frac{\varepsilon(1-\chi)}{1-\alpha}}\right)\right]\widehat{p}_t^* + \theta\overline{\pi}_t^{\frac{\varepsilon(1-\chi)}{1-\alpha}}\left[\frac{\varepsilon}{1-\alpha}\widehat{\pi}_t - \frac{\varepsilon\chi}{1-\alpha}\widehat{\pi}_{t-1} + \widehat{s}_{t-1} + \frac{\varepsilon\chi}{1-\alpha}\widehat{g}_t^{\overline{\pi}} - \widehat{g}_t^{\overline{s}}\right] \tag{9L}$$

$$\widehat{mc}_t = \widehat{w}_t + \left(\frac{\alpha}{1-\alpha}\right)\widehat{Y}_t \tag{10L}$$

---

[25]Specifically, for any variable Z_t, $\widehat{Z}_t = \ln Z_t - \ln\overline{Z}_t$. In addition, $\widehat{g}_t^{\overline{\pi}} = \ln(\overline{\pi}_t/\overline{\pi}_{t-1})$, $\widehat{g}_t^{\overline{Y}} = \ln(\overline{Y}_t/\overline{Y}_{t-1})$, and $\widehat{g}_t^{\overline{s}} = \ln(\overline{s}_t/\overline{s}_{t-1})$.

$$\hat{R}_t = \rho \hat{R}_{t-1} - \rho \hat{g}_t^{\overline{\pi}} + (1-\rho)\,\psi_\pi \hat{\pi}_t + (1-\rho)\,\psi_x \hat{x}_t + (1-\rho)\,\psi_{\Delta y}\hat{g}_t^y + \epsilon_{r,t} \qquad (11L)$$

$$\hat{g}_t^y = \hat{Y}_t - \hat{Y}_{t-1} + \hat{g}_t + \hat{g}_t^{\overline{Y}} \qquad (12L)$$

$$\hat{g}_t = \rho_g \hat{g}_{t-1} + \epsilon_{g,t} \qquad (13L)$$

$$\hat{d}_t = \rho_d \hat{d}_{t-1} + \epsilon_{d,t} \qquad (14L)$$

$$\left[\frac{\bar{g}(1+\varphi) - h(\varphi+\alpha)}{h(1-\alpha)}\right]\hat{Y}_t^n = \hat{Y}_{t-1}^n - \hat{g}_t \qquad (15L)$$

$$\hat{x}_t = \hat{Y}_t - \hat{Y}_t^n \qquad (16L)$$

We can write the system of log-linearized equations in Sims (2002)'s canonical form, but now the system will have time-varying parameters, because they are functions of $\overline{\pi}_t$:

$$\Gamma_0(\overline{\pi}_t)\hat{Z}_t = \Gamma_1(\overline{\pi}_t)\hat{Z}_{t-1} + \Psi(\overline{\pi}_t)\varepsilon_t + \Pi(\overline{\pi}_t)\eta_t.$$

However, for any given value of $\overline{\pi}_t$ (and the realization of stochastic volatility), the above system is conditionally linear and can be solved using standard methods.