

The Macroeconomic Consequences of Natural Rate Shocks: An Empirical Investigation*

Stephanie Schmitt-Grohé[†] Martín Uribe[‡]

November 5, 2023

Abstract

This paper studies the effects of natural rate shocks on output and inflation. It estimates a semi-structural model inspired by the DSGE literature. A decline in the natural rate is found to have a large negative effect on the trend of output and to be contractionary and deflationary in the short run. When the economy is constrained by the zero lower bound (ZLB), these results are consistent with the secular stagnation hypothesis. However, negative natural rate shocks are found to depress the trend of output even outside of the ZLB, calling for a more general theory of secular stagnation.

JEL classification: E3, E4.

Keywords: Natural rate of interest, permanent real interest rate shocks, secular stagnation, semi-structural model.

*We thank for comments Marco Del Negro and Christian Wolf, participants at the 2022 Paris School of Economics International Macroeconomics conference, the 2023 ASSA meetings, the 2023 Thomas Laubach Research conference, the 2023 Macroeconometric Caribbean Conference, and the 57th Annual Conference of the Canadian Economics Association, and workshop participants at Columbia University, the Bank of Canada, the European Central Bank, the Federal Reserve Bank of Kansas City, the Federal Reserve Bank of New York, and University of California Riverside. Giovanni Bonfanti provided excellent research assistance.

[†]Columbia University, CEPR, and NBER. E-mail: stephanie.schmittgrohe@columbia.edu.

[‡]Columbia University and NBER. E-mail: martin.uribe@columbia.edu.

1 Introduction

The real rate of interest has displayed a persistent decline over the past decades. A substantial amount of empirical research has been devoted to ascertaining whether this phenomenon corresponds to a fall in the natural rate of interest, understood as the permanent component of the short-term real interest rate. Less research exists on the empirical question of how movements in the natural rate affect macroeconomic indicators in the short and long runs. Does a fall in the natural rate increase or decrease the trend of output? Is a fall in the natural rate contractionary or expansionary in the short run? Is it inflationary or deflationary? This paper aims to address these questions.

We contribute to the related empirical literature by allowing the permanent component of the real interest rate to be a source of movements in real activity and prices. This approach is motivated by the dynamic optimizing general equilibrium literature in which low frequency movements in the real interest rate can occur due to exogenous permanent movements in subjective discount rates, demographic trends, or other exogenous factors determining secular shifts in the propensities to save of domestic or foreign agents.

With this motivation in mind, we develop an empirical model that combines elements of a structural vector autoregression (SVAR) model and a dynamic stochastic general equilibrium (DSGE) model. Like standard SVAR formulations, the model contains a relatively small number of cross-equation restrictions. And like DSGE structures, the model allows for identified trend shocks—including permanent disturbances to the natural rate—to affect not only the trends of endogenous variables but also their cyclical components. These features make it possible to estimate jointly the long-run and business-cycle effects of movements in the permanent component of the real interest rate. A further similarity with DSGE models is that the proposed formulation accommodates more identified shocks than observables. This is important because it creates a competition for the data between the permanent natural rate shock and a rich set of alternative disturbances with a precise economic interpretation.

The model incorporates three permanent shocks: a permanent real interest rate shock,

a permanent productivity shock, and a permanent monetary shock. In addition, it includes a transitory monetary shock and a transitory real shock. We estimate the model using Bayesian techniques on annual U.S. data for output, inflation, and the short-term nominal interest rate over the period 1900 to 2021.

The main finding of the paper is that a shock that permanently lowers the real interest rate permanently lowers the trend path of output and is contractionary and deflationary already in the short run. Specifically, a one-percentage-point fall in the permanent component of the real interest rate reduces the trend level of output by about 8 percentage points. Of this 8-percentage-point fall in the trend of output two-thirds are estimated to take place in the short run (within a year). In addition, inflation falls on impact by 25 basis points and remains below target for about 4 years. These findings suggest that the identified natural rate shock is a demand side shock as it moves output and inflation in the same direction.

Importantly, we estimate that negative natural rate shocks have a negative effect on output trend even when the economy is not constrained by the zero lower bound (ZLB) on policy rates. This finding suggests that the mechanism that triggers the effect may be more general than the one invoked by the secular stagnation hypothesis in its different varieties. This is so because in this theory being near the ZLB is key for a fall in the natural rate to push the economy onto a lower trend trajectory.

In this paper, we adopt the conventional view in applied macroeconomic analysis that output is integrated of order 1. An implication of this assumption is that the growth rate of trend output is stationary. By contrast, in the empirical related literature on the natural rate it is often assumed that output is integrated of order 2 (see, for example, Laubach and Williams, 2003). The assumption that output is integrated of order 2 is made to allow for permanent movements in the real interest rate to be driven by permanent variations in the growth rate of output. In a robustness analysis we adopt an agnostic prior for the serial correlation of output trend shocks. The results of this exercise indicate that the growth rate of trend output is strongly mean reverting. The economic relevance of this exercise

is that the main result of this paper, namely, that falls in the natural rate of interest are contractionary, is unlikely to be accounted for by permanent declines in the growth rate of trend output. The sensitivity analysis also shows that the baseline results of the paper are robust to allowing for heteroskedasticity in measurement error, alternative identification schemes for monetary shocks, and the possibility that permanent monetary disturbances affect the natural rate.

The present investigation is related to two strands of literature. One strand is concerned with the estimation of the time path of the natural rate of interest. A seminal contribution in this body of work is Laubach and Williams (2003), who model the natural rate as an unobserved permanent component in the real interest rate. Laubach and Williams (2016) apply this model and estimate that the natural rate in the United States has experienced a persistent decline since the 1980s with a particularly large fall around the financial crisis of 2008. Holston et al. (2017), Del Negro et al. (2019), Ferreira and Shousha (2021), Cesa-Bianchi et al. (2022), and Hamilton et al. (2016) provide international evidence on the persistent decline of the natural rate over the past few decades. The latter paper estimates a larger degree of uncertainty around the level of the natural rate relative to other papers in the literature. Del Negro et al. (2017) find that the decline in the natural rate in the United States was primarily caused by changes in the convenience yield of safe assets. The contribution of the present paper to this strand of the literature is to estimate the effects of innovations in the natural rate of interest (i.e., innovations to the permanent component of the real interest rate) on output, its trend, and inflation. To our knowledge this represents the first attempt to characterize empirically the macroeconomic effects of permanent real interest rate shocks.

The second strand of literature to which this paper is related is theoretical and goes back to the contributions of Hansen (1939) who linked secular declines in the real interest rate to permanent declines in potential output when the economy is near or at the zero lower bound on interest rates. Modern formulations of this hypothesis have been advanced by

Summers (2014), Eggertsson et al. (2019), and Benigno and Fornaro (2018) among others. In all of these formulations, the zero lower bound is a key ingredient in the formation of a secular stagnation. The contribution of the present paper in this regard is the finding that the secular stagnation phenomenon appears to be more general, in the sense that negative natural rate shocks appear to depress the trend path of output even when the economy is reasonably far away from the ZLB.

There also exists a class of models in which population aging can have negative effects on output and real interest rates through declines in the rate of innovation (Aksoy et al., 2019; Gagnon et al. 2021). These models have been used to shed light on whether the low real rates observed in the past two decades represent a new normal. However, explaining movements in the natural rate of interest from the perspective of this class of models over a longer period of time is not straightforward because the percentage of Americans age 65 and older has been increasing steadily over the period 1900 to 2021 whereas, as we document in the paper, the natural rate displays supercycles—for example, values of the permanent component of the real interest rate as low as those prevailing in the 2000s have also been observed in the 1930s.

The remainder of the paper is organized as follows. Section 2 presents the empirical model and the identification strategy. It also describes the data and the estimation procedure. Section 3 presents the main results of the paper, and section 4 conducts a sensitivity analysis. Section 5 concludes.

2 Empirical Model, Identification, and Estimation

In this section, we present the semi-structural empirical model and discuss the identification scheme, the data, and the estimation strategy.

2.1 Empirical Model

The model is cast in terms of the logarithm of real output per capita, denoted y_t , the inflation rate, denoted π_t , and the short-term nominal interest rate, denoted i_t . For simplicity, the exposition of the model omits intercepts and deterministic trends. We suppose that output is cointegrated with a nonstationary real shock X_t and a permanent real-interest-rate shock X_t^r . Specifically, we define the unobservable stationary variable \hat{y}_t as

$$\hat{y}_t = y_t - X_t - \delta X_t^r.$$

The endogenous latent variable \hat{y}_t is interpreted as the cyclical component of output, the exogenous latent variable X_t captures permanent movements in the state of technology, and the exogenous latent variable X_t^r captures permanent changes in the natural rate of interest. Permanent movements in the real rate of interest could stem from, for example, secular variations in demographic variables, exogenous changes in subjective discount rates, or in other factors determining the domestic or external willingness to save. The output trend is then given by

$$\text{output trend} = X_t + \delta X_t^r. \tag{1}$$

The parameter δ governs the long-run effect of permanent real-interest-rate shocks on the trend of output. As will become clear shortly, a fall in X_t^r corresponds to a decline in the natural rate of interest. Thus, a positive value of δ implies that a permanent decline in the natural rate of interest lowers the trend of output.

The model assumes that inflation is cointegrated with an exogenous stochastic nonstationary nominal disturbance denoted X_t^m . Specifically, we define the cyclical component of inflation, $\hat{\pi}_t$, as

$$\hat{\pi}_t = \pi_t - X_t^m.$$

The variable $\hat{\pi}_t$ is stationary. The permanent monetary shock X_t^m can be interpreted as the

permanent component of a stochastic (de facto) inflation target.

The model assumes that the nominal interest rate is cointegrated with the inflation target X_t^m and with the permanent real-interest-rate shock X_t^r , so that its cyclical component, denoted \hat{i}_t , is given by

$$\hat{i}_t = i_t - X_t^m - X_t^r.$$

Defining the natural rate of interest as the permanent component of the real interest rate, we have that¹

$$\text{natural rate of interest} = X_t^r. \quad (2)$$

Because the permanent nominal shock X_t^m enters with the same coefficient in the trend components of inflation and the interest rate, the model assumes that permanent changes in the inflation rate have no permanent effects on the real interest rate (i.e., that the Fisher effect holds in the long run). Section 4.4 shows that the results of the paper are robust to assuming that X_t^m enters with a coefficient different from one in the trend of i_t . It further shows that this coefficient is not significantly different from one.

In addition to the three permanent shocks (X_t , X_t^m , and X_t^r), the model includes a stationary monetary shock, denoted z_t^m , and a stationary real shock, denoted z_t . The dynamics of the cyclical components of the three endogenous variables of the model are assumed to be given by the following first-order autoregressive system

$$\begin{bmatrix} \hat{y}_t \\ \hat{\pi}_t \\ \hat{i}_t \end{bmatrix} = B \begin{bmatrix} \hat{y}_{t-1} \\ \hat{\pi}_{t-1} \\ \hat{i}_{t-1} \end{bmatrix} + C \begin{bmatrix} \Delta X_t^m \\ z_t^m \\ \Delta X_t \\ z_t \\ \Delta X_t^r \end{bmatrix}. \quad (3)$$

This structure is inspired by that of optimizing dynamic models, in which the equilibrium

¹In the related empirical literature, X_t^r is often denoted r_t^* (e.g., Laubach and Williams, 2003; Del Negro et al., 2017).

values of the cyclical components of endogenous variables depend not only on realizations of stationary shocks but also on disturbances to trend components. In particular, this feature of the model allows us to estimate the impulse responses of output, inflation, and the interest rate to an innovation in X_t^r . It is also an important difference with related papers using VAR models in which trend shocks are not allowed to affect the cycle. (See, for example, Del Negro et al., 2017). In terms of the present notation, that paper assumes that the first, third, and fifth columns of C are nil.

The stochastic driving vector $\begin{bmatrix} \Delta X_t^m & z_t^m & \Delta X_t & z_t & \Delta X_t^r \end{bmatrix}'$ follows the first-order autoregressive law of motion

$$\begin{bmatrix} \Delta X_{t+1}^m \\ z_{t+1}^m \\ \Delta X_{t+1} \\ z_{t+1} \\ \Delta X_{t+1}^r \end{bmatrix} = \rho \begin{bmatrix} \Delta X_t^m \\ z_t^m \\ \Delta X_t \\ z_t \\ \Delta X_t^r \end{bmatrix} + \Psi \begin{bmatrix} \epsilon_{t+1}^{X^m} \\ \epsilon_{t+1}^{z^m} \\ \epsilon_{t+1}^X \\ \epsilon_{t+1}^z \\ \epsilon_{t+1}^{X^r} \end{bmatrix}, \quad (4)$$

where ρ and Ψ are diagonal matrices and ϵ_t^s , for $s = X^m, z^m, X, z, X^r$, are i.i.d. disturbances distributed $N(0, 1)$. This component of the model is also inspired by DSGE formulations in which all exogenous shocks are identified and follow distinct laws of motion.

The model is cast in terms of latent variables. To be able to estimate it—that is, to estimate the matrices B , C , ρ , Ψ and the parameter δ —we introduce a vector of observable variables, for which the model has precise predictions. Specifically, we assume that the econometrician observes with measurement error the growth rate of real GDP per capita, denoted Δy_t , the change in the inflation rate, denoted $\Delta \pi_t$, and the change in the short-term nominal interest rate, denoted Δi_t . The observation equations are given by the following three identities:

$$\Delta y_t = \hat{y}_t - \hat{y}_{t-1} + \Delta X_t + \delta \Delta X_t^r + \mu_t^y, \quad (5)$$

$$\Delta\pi_t = \hat{\pi}_t - \hat{\pi}_{t-1} + \Delta X_t^m + \mu_t^\pi, \quad (6)$$

and

$$\Delta i_t = \hat{i}_t - \hat{i}_{t-1} + \Delta X_t^m + \Delta X_t^r + \mu_t^i, \quad (7)$$

where the vector $[\mu_t^y \mu_t^\pi \mu_t^i]'$ is a normally distributed mean-zero i.i.d. measurement error with a diagonal variance-covariance matrix denoted R . Observation equations (5)–(7) say that, up to a measurement error, the growth rate of an observed variable can be expressed as the sum of the growth rate of its cyclical component and the growth rate of its trend component.

2.2 Identification and Priors

The cointegration restrictions discussed thus far are themselves identification restrictions for the nonstationary shocks. In particular, they imply that X_t is the only shock that has a long-run effect on output but not on inflation or the nominal interest rate; X_t^m is the only shock that can have a long-run effect on inflation; and X_t^r is the only shock that can have a long-run effect on output and the nominal interest rate but not on inflation.

To identify the stationary shocks z_t^m and z_t , we proceed as follows. We identify the transitory monetary shock, z_t^m , by two alternative methods, sign restrictions and zero restrictions. Both methods deliver similar results. The sign restriction identification scheme allows for stationary monetary contractions to have a nonpositive impact effect on output and inflation. Formally, it imposes

$$C_{12}, C_{22} \leq 0$$

in equation (3). The zero restriction scheme assumes that z_t^m does not affect output or inflation on impact, that is, it imposes

$$C_{12} = C_{22} = 0.$$

Here we focus on the zero restriction scheme and in section 4.3 we discuss the sign restriction

identification scheme.

Without loss of generality, we normalize the impact effects of the stationary monetary shock on the interest rate and of the stationary real shock (z_t) on output to unity, that is, we set

$$C_{14} = C_{32} = 1.$$

The assumption that z_t^m has a zero impact effect on output and inflation may sound too restrictive, given the annual frequency of the data. However, in quarterly data, the effects of monetary shocks often manifest with a delay of several quarters (see, for example, Christiano et al., 2005). This issue is more naturally handled by the sign restriction approach presented in section 4.3.

We impose a zero prior mean for the parameter δ in observation equation (5), which implies a prior belief that the real interest rate shock X_t^r has no effect on the trend of output. Specifically, we assume that δ has a normal prior distribution with mean 0. In addition, we assume that this distribution is quite diffuse by setting a standard deviation of 5, which implies that if δ is one standard deviation above its prior mean, then an X_t^r shock that lowers the permanent component of the real interest rate by 1 percentage point reduces the trend of output by 5 percent. The rationale behind adopting a diffuse prior for δ is a body of theoretical work, going back to Hansen (1939), that attributes large variations in potential output to movements in the natural rate of interest. The rationale for the assumed symmetry around zero is that, while models of secular stagnation are designed to deliver a negative effect of a drop in the natural rate on trend output, one can show that in the neoclassical growth model this effect is positive and, under plausible calibrations, also large.

The elements of the matrix B in equation (3) are assumed to have normal prior distributions. In the spirit of the Minnesota prior, we impose a prior mean of 0.5 on the diagonal elements of B and a prior mean of 0 on the off diagonal elements. All elements of B are assumed to have a prior standard deviation of 0.3.

Next, we present the assumed prior distributions of the estimated elements of the matrix

C in equation (3). All estimated elements of C are assumed to have normal prior distributions with standard deviation equal to 1. Under the mean of the prior distribution, an innovation to X_t^m is assumed to have a zero impact effect on inflation and the interest rate. This requires setting the prior means of C_{21} and C_{31} to -1. All other estimated elements of C are assumed to have a prior mean of 0.

Consider now the prior distributions of the estimated parameters of the laws of motion of the exogenous driving forces given in equation (4). The diagonal elements of the matrix ρ are assumed to have a beta prior distribution with mean 0.3 and a standard deviation of 0.2, with the exception of element (4,4), which is assumed to have a mean of 0.5 and a standard deviation of 0.2. The rationale behind the assumed mean values of ρ_{11} , ρ_{33} , and ρ_{55} is that the first differences of output, inflation, and the nominal interest rate have relatively low serial correlations. The rationale for the prior mean of ρ_{22} is that transitory monetary disturbances are typically assumed to be i.i.d. or to have a low persistence. (For example, Smets and Wouters, 2007, assign this parameter a prior mean of 0.5 for a model estimated on quarterly data, which at an annual frequency corresponds to a value of about 0.2.) The justification for the higher prior mean for the serial correlation of z_t (element (4,4) of ρ) is that in business cycle analysis stationary productivity shocks are typically estimated to have a relatively high persistence. The standard deviations of all exogenous shocks, given by the diagonal of the matrix Ψ , are assumed to have gamma prior distributions with means and standard deviations equal to 1.

The trend of output is assumed to have a deterministic component. The growth rate of this component, given by the intercept of equation (5) (not shown), is assumed to have a normal prior distribution with a mean and a standard deviation equal to the sample mean and standard deviation of output growth. Changes in inflation and the interest rate are assumed to have a zero mean. That is, observation equations (6) and (7) are assumed to have no intercept. Finally, the diagonal elements of the variance-covariance matrix R of measurement errors in observation equations (5)–(7) are assumed to have uniform prior

Table 1: Prior Distributions

| Parameter | Distribution | Mean | Std. Dev. |
|-------------------------------------|--------------|---------------------------------------|---|
| Diagonal elements of B | Normal | 0.5 | 0.3 |
| Off diagonal elements of B | Normal | 0 | 0.3 |
| C_{21}, C_{31} | Normal | -1 | 1 |
| All other estimated elements of C | Normal | 0 | 1 |
| $\rho_{ii}, i = 1, 2, 3, 5$ | Beta | 0.3 | 0.2 |
| ρ_{44} | Beta | 0.5 | 0.2 |
| Diagonal elements of ψ | Gamma | 1 | 1 |
| δ | Normal | 0 | 5 |
| Diagonal elements of R | Uniform | $\frac{\text{var}(o_t)}{10 \times 2}$ | $\frac{\text{var}(o_t)}{10 \times \sqrt{12}}$ |
| Estimated element of A | Normal | $E(\Delta y_t)$ | $\sqrt{\frac{\text{var}(\Delta y_t)}{T}}$ |

Notes. T denotes the sample length. The vector o_t contains the observables, $o_t = [\Delta y_t \ \Delta \pi_t \ \Delta i_t]'$. The vector A denotes the mean of the vector of observables, $A = E(o_t)$.

distributions. The lower bounds of these distributions are set to zero, and the upper bounds are set to ensure that under the prior distribution measurement errors account for no more than 10 percent of the variance of the data.

Table 1 provides a summary of the assumed prior distributions.

2.3 Data and Estimation

The model is estimated on annual U.S. data from 1900 to 2021. For the period 1900 to 2017 the data is taken from the Jordá-Schularick-Taylor Macrohistory Database, see Jordá et al. (2017). Specifically, we downloaded from that database the series rgdppc, cpi, and stir corresponding, respectively, to real GDP per capita, the consumer price index, and the nominal short-term interest rate. Since 1955 the short-term nominal interest rate measure corresponds to the effective federal funds rate. For the period 2018 to 2021, the data source for real GDP per capita is the Bureau of Economic Analysis, for the consumer price index the Bureau of Labor Statistics, and for the federal funds rate the Federal Reserve Board. The annual value of the consumer price index is computed as the arithmetic average over the

corresponding monthly observations. Output growth, inflation, and the short-term interest rate are expressed in percent per year.

We estimate the model using Bayesian methods and use the random walk Metropolis-Hastings algorithm to obtain 50 million draws from the posterior distribution of the estimated parameter vector. We target an acceptance rate between 0.25 and 0.35. The results presented in the remainder of the paper are based on random subsamples from the 50-million draws of length either 1 million or 100 thousand, depending on the computation time involved in the quantitative analysis.

3 Results

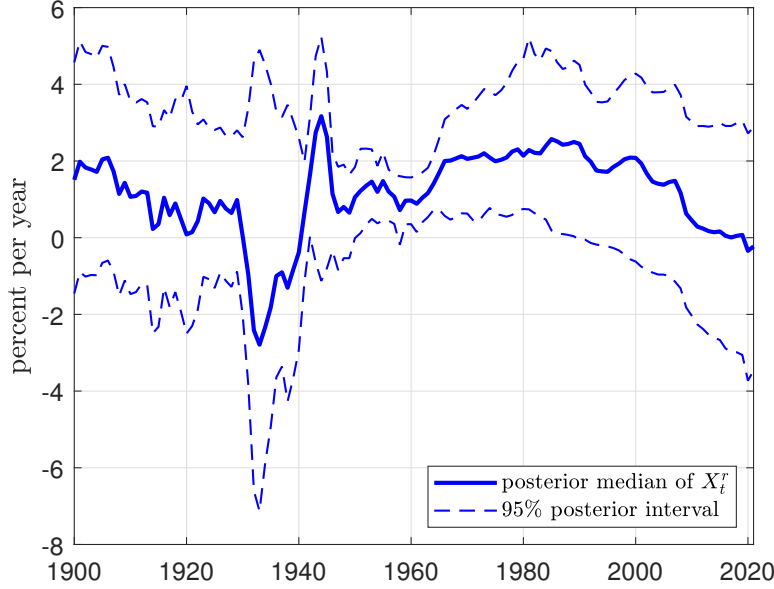
This section addresses the question of how innovations in the permanent component of the real interest rate affect output and inflation in the long and short runs.

3.1 Natural Rate Supercycles

A large number of papers devoted to the estimation of the natural rate of interest have documented a persistent decline in this variable over the past decades. Figure 1 offers a long run perspective on the behavior of the natural rate, X_t^r . It displays its inferred path over the period 1900-2021. The variable X_t^r is computed by two-sided Kalman smoothing using 100,000 draws from the posterior chain of parameter vectors of length 50 million. This method delivers an estimate of X_t^r up to a constant. To facilitate interpretation, we set this constant to ensure that X_t^r and the observed ex-post real rate, $i_t - \pi_{t+1}$, have the same sample mean, 1.15 percent per year. The solid line plots the posterior median of X_t^r and the broken lines indicate the 2.5th and 97.5th percentiles of its posterior estimate.

The figure shows that the natural rate has experienced a significant fall of about 3 percentage points since the late 1980s. This finding is in line with those reported in the related literature. However, the figure also shows that the natural rate of interest displays

Figure 1: The Natural Rate of Interest: 1900-2021



Notes. The variable X_t^r is computed by two-sided Kalman smoothing. It is normalized by adding a constant to match the observed sample mean of $i_t - \pi_{t+1}$ (1.15 percent per year). The solid line is the posterior median of X_t^r and the broken lines indicate the 2.5th and 97.5th posterior percentile of X_t^r , respectively. These statistics are computed using 100,000 randomly picked draws from an MCMC chain of length 50 million.

long cycles. We refer to those as natural rate supercycles. For the estimation period, the first supercycle begins in 1900 and ends in the mid 1980s. The trough of this supercycle coincides with the trough of the Great Depression in the 1930s. The figure suggests that the fall in the natural rate that has been taking place over the past three decades is part of a second supercycle that began in the mid 1980s.

The long swings in the natural rate of interest predicted by the model deliver three insights. First, there were significant movements in the natural rate of interest pre 1980. This finding departs from that stressed in Del Negro et al. (2019). These authors estimate, also using long data but a different empirical strategy, that the natural rate of interest displayed little variation prior to the Great Moderation. As these authors show, their results are sensitive to the fact that they impose a small mean on the prior distribution of the standard deviation of the innovation to the trend component of the real interest rate. Specifically, under their assumed prior mean the expected change in the natural rate from one year to

the next has a standard deviation of 1 basis point. They also show that if they increase this prior mean to 100 basis points, then their estimation strategy, like ours, delivers significant movements in the natural rate of interest pre 1980 (see Figure A4 of the online appendix to Del Negro et al., 2019).

Second, the persistent fall in the natural rate of interest over the past three decades is not unprecedented. Between 1900 and 1933, the natural rate fell by 4 percentage points. This figure is larger than the estimated 3 percentage points decline in the natural rate since the beginning of the most recent downward swing of the supercycle in 1985.

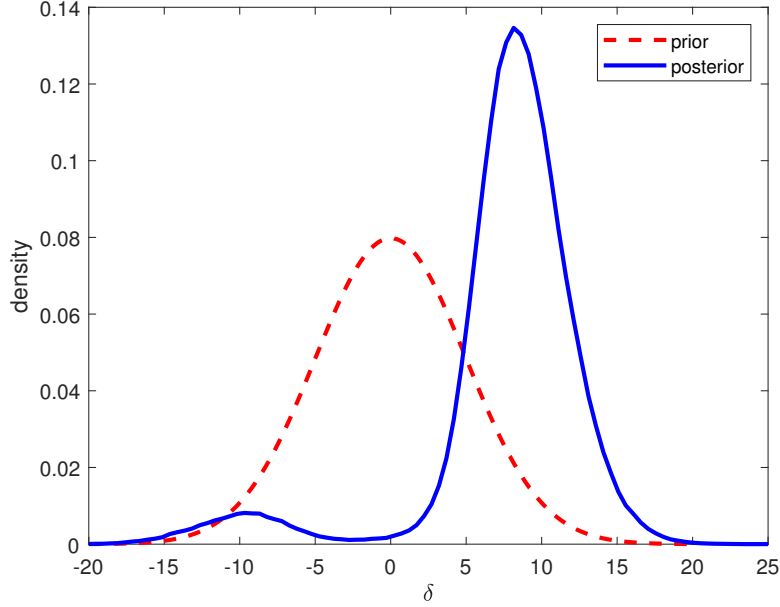
Third, the nonmonotonic behavior of the natural rate over the past 120 years sheds light on the role of aging as a driver of the permanent component of the real interest rate. Specifically, it has been argued that aging could have played a significant role in the observed decline in the natural rate of interest and output growth since the 1980s. However, aging is not a recent phenomenon; the U.S. population has been aging steadily since 1900. The fact that the permanent component of the real interest rate was as low in the trough of the first supercycle, which took place in the first half of the 20th century, as it was in 2021 suggests the presence of additional significant drivers of the natural rate.

We now turn to the main results of the paper, namely, the long- and short-run effects of innovations in the natural rate of interest on output and inflation.

3.2 Effects of Natural Rate Shocks on the Trend of Output

According to equation (1), the effect of changes in the permanent component of the real interest rate, X_t^r , on the trend of output, $X_t + \delta X_t^r$, is governed by the parameter δ . Figure 2 displays the prior and posterior distributions of this parameter. The data appears to favor positive values of δ . Under the posterior distribution, the probability that δ is positive is 93 percent compared to 50 percent under the prior distribution. The posterior median of δ is 8.4 and the posterior mean is 7.6, which means that on average a one percentage point fall in the natural rate of interest X_t^r lowers the trend level of output by about 8 percent.

Figure 2: Prior and Posterior Densities of δ



Notes. The parameter δ measures the effect of a change in the natural rate of interest, X_t^r , on the trend of output. A positive value of δ means that a decline in the natural rate of interest (a fall in X_t^r) lowers the trend of output.

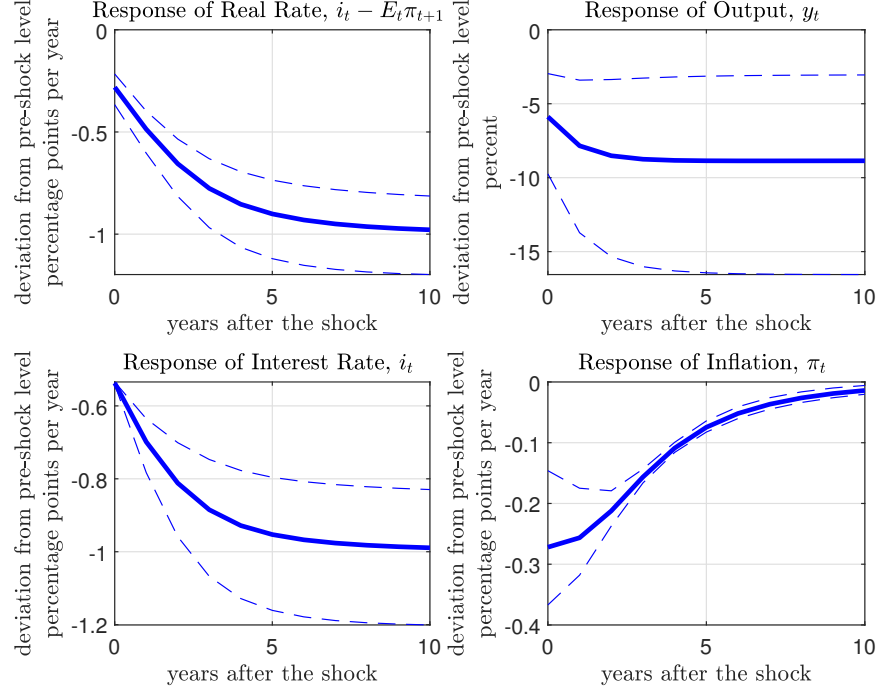
This finding suggests a sizable degree of confidence that a fall in the natural rate has a large negative effect on the permanent component of output per capita.

We note, however, that the posterior distribution of δ does have some mass (about 7 percent) to the left of zero, and that in this range it displays a second mode. For negative values of δ , a fall in the natural rate elevates trend output. The second mode can give rise to wide confidence bands of objects such as impulse response functions. For this reason, in the dynamic analysis that follows we will focus on model predictions conditional on positive values of δ .

3.3 Short-Run Effects of Natural Rate Shocks

What does the transition to a lower natural rate look like? Figure 3 addresses this question. It displays the posterior mean response to a negative natural rate shock (a fall in X_t^r) that lowers the real interest rate by 1 percentage point in the long run. The figure also includes

Figure 3: Impulse Responses to a Decline in the Natural Rate of Interest



Notes. Solid lines display the posterior mean response to a negative natural rate shock (a decrease in X_t^r) that lowers the real interest rate by 1 annual percentage point in the long run. Broken lines are asymmetric 95-percent confidence bands computed using the Sims-Zha (1999) method from 100,000 randomly picked draws from an MCMC chain of length 50 million. Impulse responses and confidence bands are conditional on $\delta > 0$.

asymmetric 95-percent confidence bands computed using the Sims-Zha (1999) method. Per the discussion in subsection 3.2, the plotted impulse responses and their associated error bands are conditional on $\delta > 0$. The figure is largely unchanged when one does not condition on positive δ values except that the error band around the output impulse response is significantly wider.

The top right panel of the figure shows that a fall in the permanent component of the real interest rate is contractionary in the short run. On impact output falls by 5.9 percent, which represents two thirds of its long run decline of 8.9 percent (the posterior mean of δ conditional on it being positive). The fall in the natural rate is also deflationary in the short run (bottom right panel of Figure 3). A one percentage point fall in the natural rate lowers inflation by about 25 basis points on impact. The decline in inflation is persistent

with a half life of about 4 years. Throughout the transition, the nominal rate falls by more than inflation, implying that a permanent decline in the natural rate of interest leads to a reduction in the real interest rate, $i_t - E_t\pi_{t+1}$, not only in the long run but also in the short run (left panels of Figure 3).

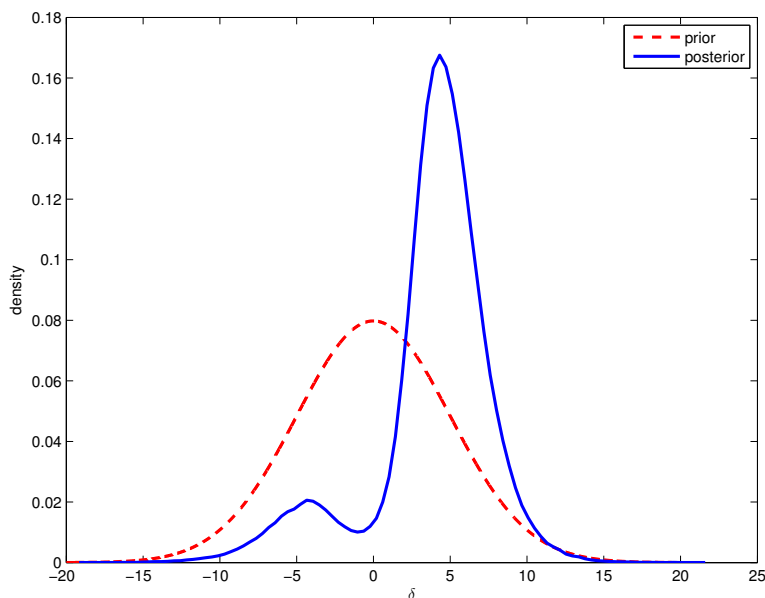
Taken together, the findings presented in this section suggest that a fall in the natural rate of interest causes a downward parallel shift in the trend of the logarithm of output and a contraction and deflation in the short run.

3.4 Excluding ZLB Episodes

The secular stagnation hypothesis argues that when the policy rate is near its effective lower bound, in response to negative natural rate shocks the real interest rate may not be able to fall to the level required for aggregate demand to meet aggregate supply. As a consequence, this theory predicts that in such circumstances, the path of output may lie chronically below potential. A natural question is whether the ZLB is a necessary condition for negative natural rate shocks to lower the trend path of output. Put differently, can negative natural rate shocks push the economy into a secular stagnation even when the policy rate is not near its effective lower bound?

In the United States, an example of a continuous period in which the policy rate was away from its effective lower bound runs from the end of World War II until the eve of the global financial crisis. Accordingly, we reestimate the model using data from 1946 to 2007. Figure 4 displays the resulting posterior density of δ . The shape of this distribution is bimodal with a peak at a positive value of 4.3 and another at a negative value of -4.3. However, as in the baseline case, the bulk of the posterior probability of δ , 87 percent, lies in the positive range, and the posterior median of δ coincides with the positive peak of 4.3. This finding suggests that even when the economy is away from the zero lower bound, there is a substantial probability that a permanent fall in the real interest rate will cause a fall in the trend of output. We note that the estimated fall is smaller than the one that results

Figure 4: Estimation Sample 1946 to 2007: Prior and Posterior Densities of δ



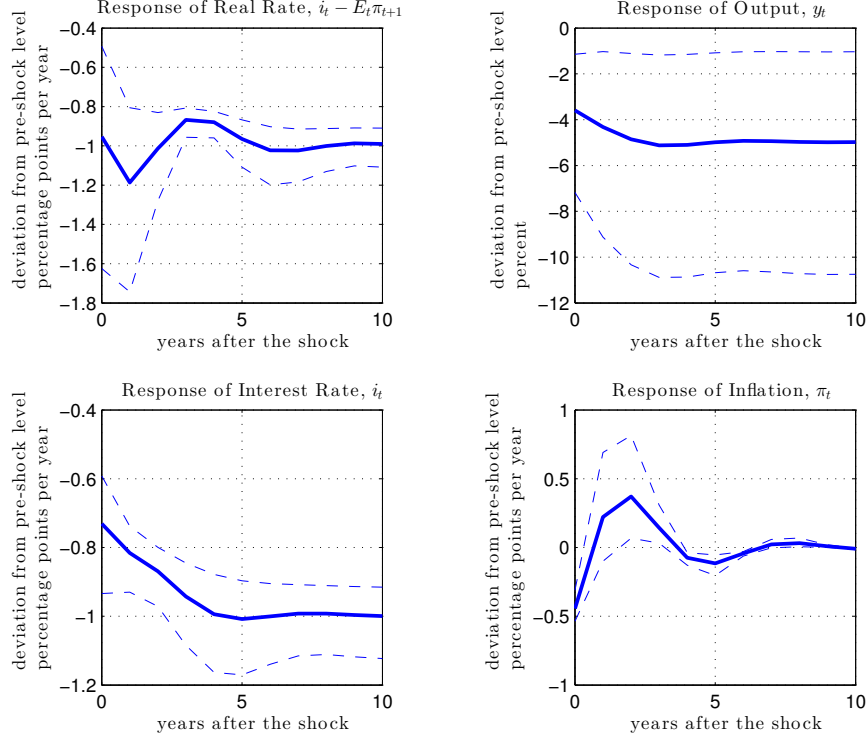
Note. See notes to Figure 2.

from an estimation on the full sample (medians of 4.4 versus 8.4 percentage points).

The short-run effects of natural rate shocks are also robust to estimating the model on a sample in which the policy rate is not constrained by the zero lower bound. Figure 5 presents the impulse responses of output, inflation, and nominal and real interest rates to an X_t^r shock that lowers the real interest rate by 1 percentage point in the long run conditional on δ being positive. The impulse responses suggest that negative natural rate shocks continue to be contractionary in the short run.

Overall, the message conveyed by this robustness check is that the ZLB, which is a key element of the secular stagnation hypothesis, does not appear to be a necessary condition for negative natural rate shocks to cause a downward shift in the trend path of output or a recession. This finding suggests that more theory is required to understand the macroeconomic effect of shocks to the natural rate of interest.

Figure 5: Estimation Sample 1946 to 2007: Impulse Responses to a Decline in the Natural Rate of Interest



Notes. See notes to Figure 3.

4 Robustness Analysis

This section presents a number of exercises aimed at ascertaining the robustness of the main results of the paper. It begins by estimating the stability of the degree of mean reversion in the growth rate of trend output to imposing agnostic priors on the serial correlations of ΔX_t and ΔX_t^r . This exercise speaks to whether the data favors a specification in which output is integrated of order 2, as assumed in some related studies. It then gauges the sensitivity of the results to allowing for heteroskedasticity in measurement errors. Next, it considers an alternative identification schemes of monetary shocks, and finally it allows for monetary shocks to have a permanent effect on the real interest rate—i.e., it allows for the Fisher effect not to hold in the long run.

4.1 Agnostic Prior on the Serial Correlation of Trend Output Growth

One possible explanation for the persistent fall in the natural rate observed over the past decades is that it was driven by permanent declines in the growth rate of productivity. This explanation requires that productivity growth be nonstationary, which in turn implies that the level of output is integrated of order 2. The assumption that output is integrated of order 2 is maintained by a number of empirical papers on the natural rate, most notably Laubach and Williams (2003), who assume that the growth rate of the trend component of output follows a pure random walk.

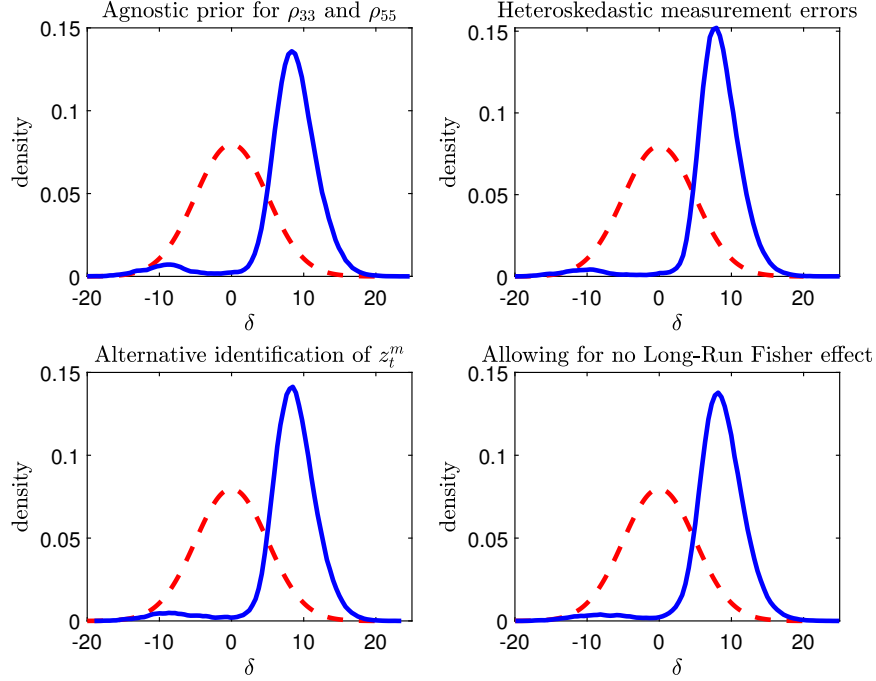
Though common in empirical papers on the natural rate, the assumption that output is integrated of order 2 is somewhat unusual in the analysis of macroeconomic time series, where it is standard to model output as being integrated of order 1 (Stock and Watson, 1998). Nonetheless, it is of interest to ascertain the degree to which in the present model the data favors parameters implying that the growth rate of trend output is close to a pure random walk.

In the model, the growth rate of trend output is given by $\Delta X_t + \delta \Delta X_t^r$. The model assumes that ΔX_t and ΔX_t^r follow AR(1) processes with serial correlations ρ_{33} and ρ_{55} , respectively. A random walk in the growth rate of trend output occurs when either $\rho_{33} = 1$ or $\rho_{55} = 1$.

The baseline estimation of the model assumes that ρ_{33} and ρ_{55} have beta prior distributions with mean 0.3 and standard deviation 0.2. The resulting posterior means are 0.26 and 0.31, respectively, and the posterior standard deviations are 0.15 and 0.11. These estimates suggest that the growth rates of the trend components of output are strongly mean reverting.

To ascertain whether this result is driven by the assumed priors, here we consider the case in which ρ_{33} and ρ_{55} have uniform prior distributions over the interval $[0, 1]$. The prior distributions of all other estimated parameters are the same as those shown in Table 1. We find that the posterior estimates are little changed relative to their baseline counterparts.

Figure 6: Prior and Posterior Densities of δ : Sensitivity Analysis



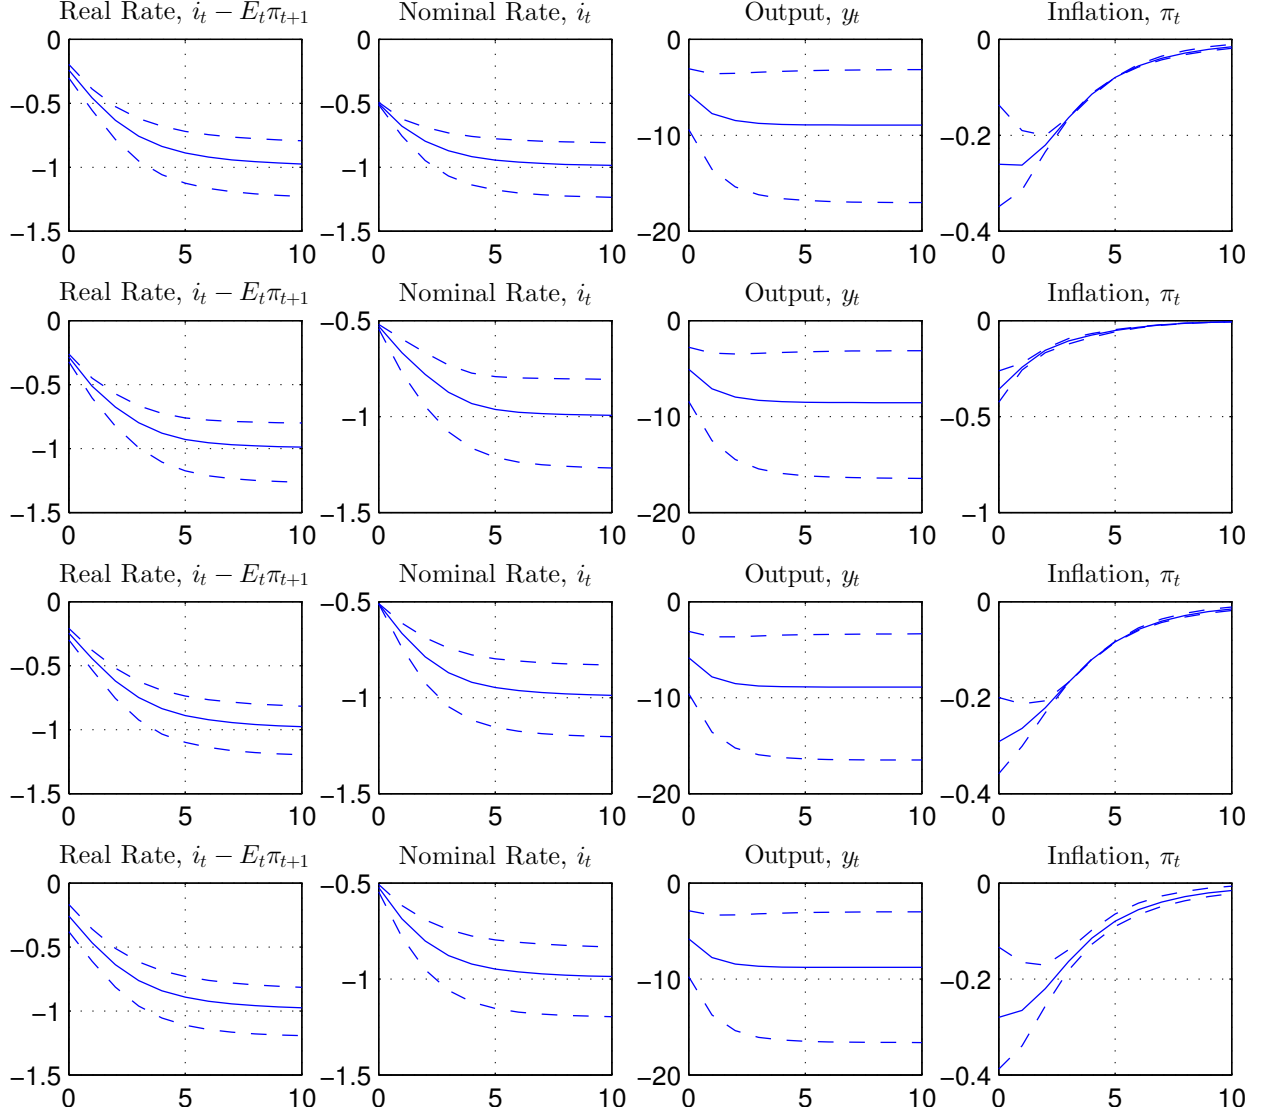
Notes. Solid lines display the posterior density of the parameter δ and broken lines the prior density.

Specifically, in the model with uniform priors, the posterior means of ρ_{33} and ρ_{55} are both 0.33 and their posterior standard deviations are 0.21 and 0.12, respectively. This suggests that the strong mean reversion in the growth rates of the trend components of output obtained in the baseline estimation is not driven by the assumed priors. The top left panel of Figure 6 and the first row of Figure 7 show that the main predictions of the baseline model are robust to assuming uniform priors for ρ_{33} and ρ_{55} . From an economic point of view this result is relevant because it suggests that the estimated contraction in the level of output in response to a fall in the natural rate of interest estimated in this paper is unlikely to be driven by a permanent decline in productivity growth.

4.2 Heteroskedasticity in Measurement Errors

The second robustness check we perform is to allow for heteroskedasticity in measurement errors pre and post 1955. The rationale for this exercise is that arguably, the systematic

Figure 7: Impulse Responses to a Decline in the Natural Rate: Sensitivity Analysis



Notes. Row 1, uniform prior distributions over the interval $[0, 1]$ for ρ_{33} and ρ_{55} ; row 2, heteroskedastic measurement errors; row 3, alternative identification of the stationary monetary shock z_t^m ; and row 4, allowing for no long-run Fisher effect, by letting the trend component of the nominal interest rate be $(1 + \alpha)X_t^m + X_t^r$ and giving α a uniform prior distribution over the interval $[-0.5, 0.5]$. The fall in X_t^r is such that the real rate falls by 1 percentage point in the long run. Solid lines are posterior means, and broken lines are asymmetric 95-percent confidence bands computed with the Sims-Zha (1999) method and using 100,000 draws from the posterior distribution conditional on $\delta > 0$.

compilation of NIPA data and aggregate price indices began in earnest after World War II, which conceivably could have resulted in larger measurement errors in the earlier period. Accordingly, we estimate the model allowing measurement errors to capture up to 20 percent of the variance of each observable prior to 1955. This is twice the amount allowed in the baseline estimation. In the post 1955 subperiod, we continue to limit the fraction of the variance of the observables explained by measurement error at 10 percent.

We find that measurement errors are indeed much larger in the pre-postwar period. The fraction of the variance of the observables explained by measurement error is 9 times larger for output growth, 13 times larger for the change in inflation, and 2 times larger for the change in the interest rate. Nonetheless, the structural matrices of the model are little changed relative to the baseline case. Consistent with this result, Figures 6 and 7 show that the main insights derived from the baseline specification are robust to allowing for heteroskedasticity in measurement errors.

4.3 Alternative Identification of the Temporary Monetary Shock

In the baseline formulation, the transitory monetary shock z_t^m is identified by assuming that on impact it can change the policy rate but not output or inflation ($C_{12} = C_{22} = 0$). One potential criticism of this approach is that it may not be too reasonable at an annual frequency, as one year may be enough time for z_t^m to affect output and inflation. To address this concern, we allow z_t^m to have a nonpositive impact effect on these two variables. Specifically we impose the restrictions

$$C_{12}, C_{22} \leq 0.$$

If one is to allow for only one stationary monetary shock with these characteristics, it is necessary to impose the additional restriction that the impact effect on the policy rate of an increase in z_t —the other temporary shock in the model—be nonnegative, that is, one must

impose

$$C_{34} \geq 0.$$

We implement this identification scheme by imposing appropriate restrictions on the prior distributions of the three parameters involved. Specifically, we assume that $-C_{12}$ and $-C_{22}$ have Gamma prior distributions with means 0.7 and 0.2, respectively, and standard deviations 0.2 and 0.1, respectively. We chose these values to match the impact effect on output and inflation of a monetary innovation after one year in the VAR model estimated in Christiano et al. (2005).² For the parameter C_{34} , we assume a uniform prior distribution over the interval $[0, 2]$.

As shown in the bottom left panel of Figure 6 and the third row of Figure 7, the main results obtained from the baseline model—in particular, the short- and long-run effects of a fall in the natural rate on output, inflation, and interest rates—also hold under this alternative identification of the stationary monetary shock.

As noted in Uribe (2022), one need not restrict C_{34} to be nonnegative to achieve an economically meaningful identification of both stationary shocks. Allowing C_{34} to take negative and positive values opens the door for the existence of two monetary shocks. In this case, identification is achieved by the assumed differences in the prior distributions of the parameters of the second and fourth columns of C and those of the second and fourth diagonal elements of the matrix ρ . For example, as explained in section 2.2, under the prior distribution z_t^m is on average less persistent than z_t . We find (not shown) that the results of the paper also go through under this identification scheme when C_{34} is assigned a standard normal prior distribution. As it turns out, the posterior means of C_{34} and C_{24} are positive, implying that z_t is most likely real demand shock.

²Specifically, Figure 1 in that paper shows that a 0.75 percentage point fall in the policy rate causes 4 quarters later an increase in output of about 0.5 percent and an elevation in inflation of about 0.15 percentage points, with centered 95% error bands of approximate widths of 0.25 and 0.15 on each side (which we take to be approximately 2 standard deviations). Thus, we set the prior means of C_{12} and C_{22} at $0.7 \approx 0.5/0.75$ and $0.2 \approx 0.15/0.75$, respectively, and their prior standard deviations at $0.2 \approx 0.25/(0.75 \times 2)$ and $0.1 \approx 0.15/(0.75 \times 2)$, respectively.

4.4 Allowing Deviations from the Long-Run Fisher Effect

The baseline model assumes that X_t^m enters with a unit coefficient in the permanent components of both inflation (π_t) and the nominal interest rate (i_t). This means that a permanent monetary shock leaves the real interest rate unchanged in the long run, or, equivalently, that the long-run Fisher effect holds. We now relax this assumption and allow for permanent monetary shocks to affect the real rate in the long run. Specifically, we redefine the trend component of the nominal interest rate to be $(1 + \alpha)X_t^m + X_t^r$, while maintaining the assumption that the trend component of inflation is X_t^m . Thus, if α is positive, a permanent increase in inflation causes an increase in the real rate in the long run. The rest of the model is unchanged.

We estimate α along with all other parameters of the modified model. We assume that at the mean of the prior distribution the long-run Fisher effect holds, that is, $\alpha = 0$. Specifically, we assume that α has a uniform prior distribution over the interval $[-0.5, 0.5]$. This means that if α is one standard deviation above its prior mean, then a permanent increase in inflation of 1 percentage point leads to a permanent increase in the nominal interest rate of 1.29 percentage points and therefore to a permanent increase in the real interest rate of 0.29 percentage points.

The assumption that the prior mean of α is zero is motivated by the following three observations. First, across time and countries over long periods of time, inflation and the nominal interest rate tend to move one for one (see, for example, Uribe, 2022, Figures 1 and 2). Second, existing cointegration analyses are inconclusive about the sign of the deviation from the long-run Fisher effect, with some finding that a permanent increase in the inflation rate leads in the long run to a permanent increase in the real interest rate (Azevedo et al., forthcoming), to a decrease (King and Watson, 1997), or to no change (Uribe, 2022). Third, a large number of studies on the joint behavior of output, inflation, and the nominal interest rate assumes, on theoretical grounds, that the real interest rate is independent of the rate of inflation in the long run (see, for example, Galí, 1992, among many others).

We find little evidence against the long-run Fisher effect. The posterior mean of α is -0.019. This means that a permanent monetary shock that increases inflation by 1 percentage point in the long run causes in the long run an increase of 0.98 percentage points in the nominal interest rate and a fall of less than 2 basis points in the real rate. Further, α is imprecisely estimated, with a posterior standard deviation of 0.28.

We interpret this result as providing support to our baseline assumption that the long-run Fisher effect holds, that is, that monetary policy does not have permanent effects on the real interest rate. Finally, as can be seen in the bottom right panel of Figure 6 and the last row of Figure 7, the main results obtained from the baseline model (i.e., the short- and long-run effects of a fall in the natural rate on output, inflation, and interest rates) also obtain under the present modification of the model.

5 Conclusion

What theories should we use to inform us about the economic consequences of persistent changes in consumers' subjective discount factors, the world real interest rate, or population growth? One way to make progress on these fundamental questions is to ask a related but simpler question: What happens with actual macroeconomic aggregates in the short and long runs in response to a change in the permanent component of the real interest rate? Having an answer to this question is important because it provides restrictions on the class of theories of the natural rate of interest that are empirically sound.

In this paper, we formulate an empirical model with minimal identification restrictions suitable for estimating the macroeconomic effects of permanent disturbances to the real interest rate. We estimate the model on U.S. data over the period 1900 to 2021. The estimation results suggest that the answer to the question posed above is that a fall in the natural rate of interest puts the economy on a lower trend trajectory and is contractionary and deflationary in the short run.

Importantly, the estimated model predicts that a fall in the natural rate of interest depresses the trend of output not only when the economy is at the zero lower bound but also when it is away from it. This result is intriguing because the most prominent theory suggesting that negative natural rate shocks can put the economy on a lower growth trajectory is the secular stagnation hypothesis, which relies on the economy being at the zero lower bound. A challenge for future research is therefore to develop a theoretical framework in which falls in the natural rate lower potential output even when the policy rate is unconstrained.

The results of this paper suggest that such theory must clear two additional hurdles. First, while population has been aging more or less monotonically since the beginning of 1900, the natural rate displays supercycles. In particular, values as low or lower than those seen in the 2000s were also seen in the 1930s. Second, postulating that the natural rate is driven by permanent movements in productivity growth requires assuming that output is integrated of order two, which lacks empirical support.

References

- Aksoy, Yunus, Henrique S. Basso, Ron P. Smith, and Tobias Gras, “Demographic Structure and Macroeconomic Trends,” *American Economic Journal: Macroeconomics* 11, January 2019, 193–222.
- Azevedo, João Valle, João Ritto, and Pedro Teles, “The Neutrality of Nominal Rates: How Long is the Long Run?,” *International Economic Review*, forthcoming.
- Benigno, Gianluca, and Luca Fornaro, “Stagnation Traps,” *Review of Economic Studies* 85, July 2018, 1425–1470.
- Cesa-Bianchi, Ambrogio, Richard Harrison, and Rana Sajedi, “Decomposing the drivers of Global R^* ,” Bank of England, Staff Working Paper No. 990, July 2022.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy* 113, 2005, 1–45.
- Del Negro, Marco, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti, “Safety, Liquidity, and the Natural Rate of Interest,” *Brookings Papers on Economic Activity* 48, Spring 2017, 235–316.
- Del Negro, Marco, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti, “Global Trends in Interest Rates,” *Journal of International Economics* 118, May 2019, 248–262.
- Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins, “A Model of Secular Stagnation: Theory and Quantitative Evaluation,” *American Economic Journal: Macroeconomics* 11, January 2019, 1–48.
- Ferreira, Thiago, and Samer Shousha, “Supply of Sovereign Safe Assets and Global Interest Rates,” mimeo, Board of Governors of the Federal Reserve System, March 2021.
- Fisher, Irving, *Appreciation and Interest*, New York: Macmillan, 1896.
- Galí, Jordi, “How Well Does the IS-LM Model Fit Postwar U.S. Data?,” *Quarterly Journal of Economics* 107, May, 1992, 709–738.
- Gagnon, Etienne, Benjamin K. Johannsen, and David López-Salido, “Understanding the

- New Normal: The Role of Demographics,” *IMF Economic Review* 69, June 2021, 357–390.
- Hamilton, James D., Ethan S. Harris, Jan Hatzius, and Kenneth D. West, “The Equilibrium Real Funds Rate: Past, Present, and Future,” *IMF Economic Review* 64, 2016, 660–707.
- Hansen, Alvin H., “Economic Progress and Declining Population Growth,” *American Economic Review* 29, March 1939, 1–15.
- Holston, Kathryn, Thomas Laubach, and John C. Williams, “Measuring the Natural Rate of Interest: International Trends and Determinants,” *Journal of International Economics* 108, 2017, S59–S75.
- Jordá, Óscar, Moritz Schularick, and Alan M. Taylor, “Macrofinancial History and the New Business Cycle Facts,” in Martin Eichenbaum and Jonathan A. Parker (eds.), *NBER Macroeconomics Annual 2016, Volume 31*, Chicago: University of Chicago Press, 2017, 213–263.
- King, Robert G., and Mark W. Watson, “Testing Long-Run Neutrality,” *Federal Reserve Bank of Richmond Economic Quarterly* 83, Summer 1997, 69–101.
- Laubach, Thomas, and John C. Williams, “Measuring the Natural Rate of Interest,” *Review of Economics and Statistics* 85, November 2003, 1063–70.
- Laubach, Thomas, and John C. Williams, “Measuring the Natural Rate of Interest Redux,” *Business Economics* 51, April 2016, 57–67.
- Sims, Christopher, and Tao Zha, “Error Bands for Impulse Responses,” *Econometrica* 67, 1999, 1113–1156.
- Smets, Frank, and Rafael Wouters, “Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach,” *American Economic Review* 97, June 2007, 586–606.
- Stock, James H., and Mark W. Watson, “Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model,” *Journal of the American Statistical Association* 93, 1998, 349–58.
- Summers, Lawrence H., “U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the

Zero Lower Bound,” *Business Economics* 49, April 2014, 65–73.

Uribe, Martín, “The Neo-Fisher Effect: Econometric Evidence from Empirical and Optimizing Models,” *American Economic Journal: Macroeconomics* 14, July 2022, 133–62.