Economic Modelling 132 (2024) 106641

Contents lists available at ScienceDirect

Economic Modelling

journal homepage: www.journals.elsevier.com/economic-modelling

Econometric issues in the estimation of the natural rate of interest ☆

Daniel Buncic *

Stockholm Business School, Stockholm University, Sweden

ARTICLE INFO

JEL classification:
C32
E43
E52
O40

Keywords:
Natural rate of interest
Median Unbiased Estimation
Kalman Filter
Spurious relations
Misspecified econometric models

ABSTRACT

The natural rate of interest is a key benchmark steady-state rate of return that is used by central banks to measure the stance of monetary policy. Asset managers rely on estimates of the natural rate to make long-term investment decisions for their clients. This paper outlines a number of important econometric issues in the estimation of the natural rate of interest from a widely used structural model. Using a simulation experiment as well as empirical data for the U.S., I show that one of the key parameters of the structural model is overestimated and leads to a spuriously amplified downward trend in the natural rate. The paper shows how the overestimation can be remedied and provides alternative estimates of the natural rate. Various other issues are discussed that policy makers should be aware of when utilizing this model's natural rate estimates for policy decision.

1. Introduction

Since the global financial crisis, nominal interest rates have declined substantially to levels last witnessed following the Great Depression. The academic as well as policy literature has attributed this decline in nominal interest rates to a decline in the natural rate of interest; namely, the rate of interest consistent with employment at full capacity and inflation at its target. In this literature, Holston et al.'s (2017, HLW henceforth) estimates of the natural rate have become particularly influential and are widely regarded as a benchmark. In fact, the Federal Reserve Bank of New York (FRBNY) maintains an entire webpage dedicated to providing updates to HLW's estimates of the natural rate, not only for the United States (U.S.), but also for the Euro Area, Canada and the United Kingdom (U.K.) (see https://www.newyorkfed.org/research/policy/rstar).

In HLW's model, the natural rate of interest is defined as the sum of (annualized) trend growth in output and '*other factor*' z_t. This ' *other factor*' z_t is meant to capture various underlying structural factors such as savings/investment imbalances, demographic changes, and fiscal imbalances that influence the natural rate, but which are not captured by trend growth. In Fig. 1, I show filtered estimates of '*other factor*' z_t. The dashed (blue) and solid (red) lines in Fig. 1 show estimates obtained with data ending in 2016:Q3 and 2019:Q2, respectively. The strong and persistent downward trending behaviour of ' *other factor*' z_t is striking from Fig. 1, particularly from 2012:Q1 onwards. The two (black) dashed vertical lines mark the periods 2012:Q1 and 2015:Q4. In 2015:Q4, the Federal Reserve started the tightening cycle and raised nominal interest rates by 25 basis points. In 2012:Q1, real rates began to rise due to a (mild) deterioration in inflation expectations. Both led to an increase in the real rate. Yet, the estimate of '*other factor*' z_t declined by about 50 basis points from 2012:Q1 to 2015:Q4, and then another 50 basis points from 2015:Q4 to 2019:Q2, reaching a value of -1.58 in 2019:Q2. Because z_t evolves as a driftless random walk in

☆ Without implications, I am grateful to Jim Stock, Adrian Pagan, Neil Ericsson, Paolo Giordani, Michael Kiley, Refet Gürkaynak (EEA Session Chair), Fabrizio Venditti, Kirstin Hubrich, Benoit Mojon, Gabriel Perez Quiros, Glenn Rudebusch, Luc Bauwens, Francesco Ravazollo, Simon van Norden, Fabio Canova, Eric Leeper, Lawrence Summers, Georgi Krustev, Bernd Schwaab, Alessandro Galesi, Claus Brand, Wolfgang Lemke, Eric Renault, Robin Lumsdaine, Jean-Paul Renne, Robert Carl Michael Beyer, Volker Wieland, John Paul Broussard, Christopher Gibbs, Andrea Papetti, Per Krusell, Ingvar Strid, Xin Zhang, Emanuel Moench, Elmar Mertens, Stéphane Dupraz, Mikkel Plagborg-Møller, as well as seminar and conference participants at Stockholm University, Sveriges Riksbank, the ECB, Deutsche Bundesbank, Norges Bank, Banque de France, NIER/KI, Banca d'Italia, the 2020 EEA Congress, the 2020 SNDE Symposium, eMAF 2020, IAAE 2021, and SED 2021 for comments that helped to improve the paper. I thank two anonymous referees and the editor for feedback and suggestions that improved the readability of the paper. Parts of this paper were written while I was visiting the Research Division of the Monetary Policy Department at the Riksbank. I thank Jesper Lindé and my former colleagues at FOE for their hospitality and many stimulating discussions. Lorand Abos provided excellent research assistance. I am grateful to the Jan Wallander and Tom Hedelius Foundation and the Tore Browaldh Foundation for research support (Grant Number: P20-0155).

* Correspondence to: Daniel Buncic, Stockholm Business School, Stockholm University, SE-103 37, Stockholm, Sweden.
E-mail address: daniel.buncic@sbs.su.se.
URL: https://www.danielbuncic.com.

https://doi.org/10.1016/j.econmod.2023.106641
Received 31 January 2023; Received in revised form 28 December 2023; Accepted 28 December 2023
Available online 5 January 2024

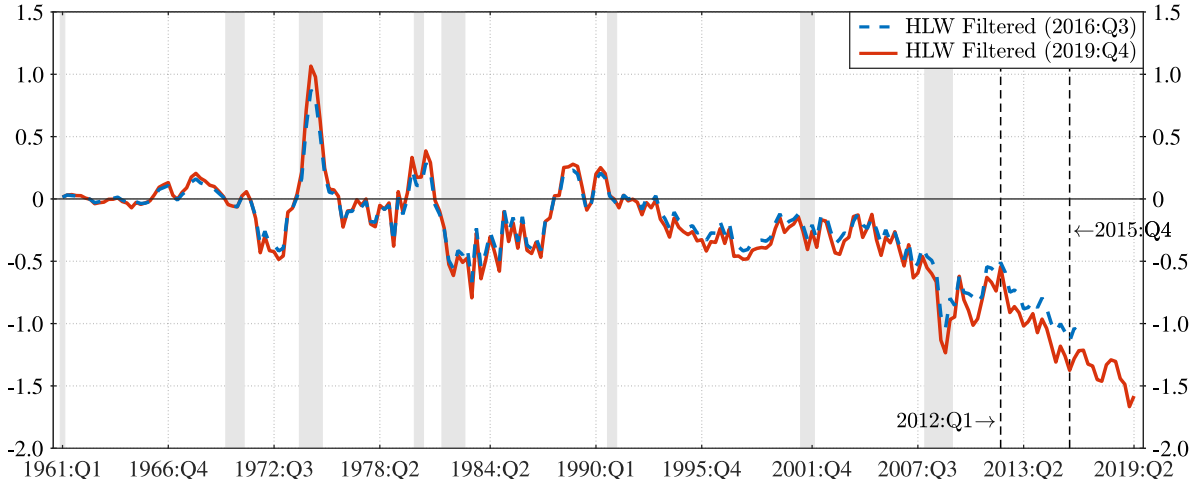

Fig. 1. Filtered estimates of HLW's *'other factor'* z_t.

the model, the only parameter that affects the *'trending behaviour'* of z_t is the ' *signal-to-noise*' ratio parameter λ_z. The size of λ_z thus directly affects the natural rate estimate, with changes in the trend in z_t directly transferred to changes in the trend in the natural rate.

The goal of this paper is to raise a number of econometric issues with HLW's model of the natural rate. More specifically, I show that the implementation of Stock and Watson's (1998) Median Unbiased Estimation (MUE) of the signal-to-noise ratio in Stage 2 is based on a misspecified model, which yields an excessively large estimate of λ_z. Since the magnitude of λ_z determines the severity of the downward trend in '*other factor*' z_t, this misspecification affects the estimate of the trend in the natural rate. Using a simulation experiment, I show that HLW's MUE procedure in Stage 2 leads to spuriously large estimates of λ_z even when the true value in the data generating process is zero. The strong and persistent downward trend in the empirical estimate of the natural rate for the U.S. is thus spurious.

The paper also outlines five other issues with the model that are relevant for policy analysis. These are related to: (1) the circular relationship between the exogenously included federal funds policy rate and the resulting estimate of the natural rate, (2) the excessive sensitivity of the estimates to the chosen starting date used in the estimation, (3) the specification and/or existence of *'other factor'* z_t in the model, (4) the use of filtered as opposed to smoothed states in the visual presentation of that natural rate for policy analysis, and (5) the ability to recover the shocks driving the natural rate, and thus the natural rate itself.

This paper is related to, but distinct from, a broader literature on natural rates of interest and the estimation thereof. Fiorentini et al. (2018) examine the role of Kalman filter uncertainty on the precision of the estimates of the components of the natural rate and show that when the IS and/or Phillips curve parameters b_y and a_r (γ and κ in their notation) go to zero, filter uncertainty increases. Berger and Kempa (2019) use a Bayesian estimation approach and modify HLW's baseline model by allowing for time varying (stochastic) volatility. Lewis and Vazquez-Grande (2018) assess the validity of the random walk specification of '*other factor*' z_t, while Martínez-García (2021) uses a two-country New Keynesian model to estimate the natural rate.

The rest of the paper is organized as follows. In Section 2, the structural model of the natural rate of interest is described, together with Stock and Watson's (1998) MUE, the three Stage procedure that is implemented to estimate the model, as well as some simulation results. In Section 3, estimates from a corrected implementation are provided. Section 4 outlines other issues with the model that policy makers should be aware of. Section 5 concludes the study.

2. Estimating the natural rate

This section describes HLW's 3 Stage procedure to estimate the natural rate. It also discusses some implications of Stock and Watson's (1998) Median Unbiased Estimator (MUE) and relates this back to the implementation of MUE in Stage 2 in HLW.

2.1. HLW's structural model

Holston, Laubach, and Williams (2017) use the following *'structural'* model to estimate the natural rate of interest:

$$\text{Output:} \quad y_t = y_t^* + \tilde{y}_t \tag{1a}$$

$$\text{Inflation:} \quad \pi_t = b_\pi \pi_{t-1} + (1-b_\pi)\pi_{t-2,4} + b_y \tilde{y}_{t-1} + \varepsilon_t^\pi \tag{1b}$$

$$\text{Output gap:} \quad \tilde{y}_t = a_{y,1}\tilde{y}_{t-1} + a_{y,2}\tilde{y}_{t-2} + \frac{a_r}{2}\left[\left(r_{t-1} - r_{t-1}^*\right) + \left(r_{t-2} - r_{t-2}^*\right)\right] + \varepsilon_t^{\tilde{y}} \tag{1c}$$

$$\text{Output trend:} \quad y_t^* = y_{t-1}^* + g_{t-1} + \varepsilon_t^{y^*} \tag{1d}$$

$$\text{Trend growth:} \quad g_t = g_{t-1} + \varepsilon_t^g \tag{1e}$$

$$\text{Other factor:} \quad z_t = z_{t-1} + \varepsilon_t^z, \tag{1f}$$

where y_t is 100 times the (natural) log of real GDP, y_t^* is the permanent or trend component of GDP, $\tilde{y}_t$ is its cyclical component, π_t is annualized quarter-on-quarter PCE inflation, and $\pi_{t-2,4} = \left(\pi_{t-2} + \pi_{t-3} + \pi_{t-4}\right)/3$. The real interest rate r_t is computed as:

$$r_t = i_t - \pi_t^e, \tag{2}$$

where expected inflation is constructed as:

$$\pi_t^e = (\pi_t + \pi_{t-1} + \pi_{t-2} + \pi_{t-3})/4, \tag{3}$$

and i_t is the *exogenously* determined nominal interest rate, the federal funds rate.

The natural rate of interest r_t^* is computed as the sum of (annualized) trend growth g_t and *'other factor'* z_t, both of which are $I(1)$ processes.[1] The real interest rate gap is defined as $\tilde{r}_t = (r_t - r_t^*)$. The error terms $\varepsilon_t^\ell, \forall \ell = \{\pi, \tilde{y}, y^*, g, z\}$ are assumed to be *i.i.d* normal distributed,

[1] From the description of the data, we can see that the nominal interest rate i_t as well as inflation π_t are defined in annual or annualized terms, while output, and hence the output gap, trend and trend growth in output are defined at a quarterly rate. Due to this measurement mismatch, trend growth g_t is scaled by 4 to be expressed at an annualized rate whenever it enters equations that relate it to annualized variables. The natural rate is thus factually computed as $r_t^* = 4g_t + z_t$.

mutually uncorrelated, with zero means and time-invariant variances denoted by σ^2_ℓ. Inflation is restricted to follow an integrated AR(4) process.

Holston et al. (2017) argue that due to *'pile-up'* at zero problems with Maximum Likelihood (ML) estimation of the variances of the innovation terms ϵ^g_t and ϵ^z_t in (1), estimates of σ^2_g and σ^2_z are "*likely to be biased towards zero* " (page S64). To avoid such *'pile-up'* at zero problems, they employ MUE of Stock and Watson (1998) in two preliminary steps — Stage 1 and Stage 2 — to get estimates of what they refer to as signal-to-noise ratios defined as $\lambda_g = \sigma_g/\sigma_{y^*}$ and $\lambda_z = a_r\sigma_z/\sigma_{\tilde{y}}$. In Stage 3, the remaining parameters of the full model in (1) are estimated, conditional on the median unbiased estimates $\hat{\lambda}_g$ and $\hat{\lambda}_z$ obtained in Stages 1 and 2, respectively. Before I describe how MUE in Stage 2 is implemented in HLW, I summarize the main findings on MUE from Stock and Watson (1998) in the section below.

2.2. Median unbiased estimation

Stock and Watson (1998) proposed MUE in the general setting of Time Varying Parameter (TVP) models. TVP models are commonly specified in a way that allows their parameters to change gradually or smoothly over time. This is achieved by defining the parameters to evolve as driftless random walks (RWs), with the variances of the innovation terms in the RW equations assumed to be small. One issue with Kalman Filter based ML estimation of such models is that estimates of these variances *can* frequently *'pile-up'* at zero when the true error variances are *'very'* small, but nevertheless, non-zero in population.

Stock and Watson (1998) show simulation evidence of *'pile-up'* at zero problems with Kalman Filter based ML estimation in Table 1 on page 353 of their paper. In their simulation set-up, they consider the following data generating process for the series GY_t:

$$GY_t = \beta_t + \varepsilon_t \tag{4a}$$

$$\beta_t = \beta_{t-1} + (\lambda/T)\eta_t, \tag{4b}$$

where ε_t and η_t are drawn from $i.i.d.$ standard normal distributions, β_{00} is initialized at 0, and the sample size is held fixed at $T = 500$ observations, using 5000 replications. The λ values that determine the size of the variance of $\Delta\beta_t$ are generated over a grid from 0 to 30, with unit increments. Four median unbiased estimators relying on four different structural break test statistics are compared to two ML estimators. The first ML estimator, referred to as the maximum profile likelihood estimator (MPLE), treats the initial state vector as an unknown parameter to be estimated. The second estimator, referred to as maximum marginal likelihood estimator (MMLE), treats the initial state vector as a Gaussian random variable with a given mean and variance. When the variance of the integrated part of the initial state vector goes to infinity, MMLE produces a likelihood with a diffuse prior.

How one treats the initial condition in the Kalman Filter recursions matters substantially for the *'pile-up'* at zero problem with MLE. This fact has been known, at least, since the work of Shephard and Harvey (1990).[2] The simulation results reported in Table 1 on page 353 in Stock and Watson (1998) show that *'pile-up'* at zero frequencies are *considerably* lower for MMLE than for MPLE, which estimates the initial state vector. For instance, for the smallest considered non-zero population value of $\lambda = 1$, which implies a standard deviation of $\Delta\beta_t$ ($\sigma_{\Delta\beta}$ henceforth) of $\lambda/T = 1/500 = 0.002$, MMLE produces an at most 14 percentage points higher *'pile-up'* at zero frequency than MUE

[2] On page 340, Shephard and Harvey (1990) write to this: *"…we show that the results for the fixed and known start-up and the diffuse prior are not too different. However, in Section 4 we demonstrate that the sampling distribution of the ML estimator will change dramatically when we specify a fixed but unknown start-up procedure."* Their Tables II and III quantify how much worse the ML estimator that attempts to estimate the initial condition in the local level model performs compared to MLE with a diffuse prior.

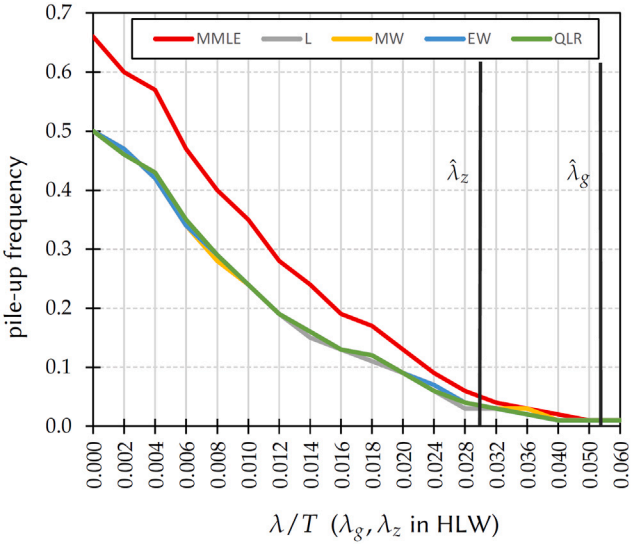

Fig. 2. Pile-up at zero frequencies of MMLE and 4 MUEs. The x-axis values show λ/T, which is the representation used in HLW (i.e., λ_z and λ_g are normalized by the sample size T).

(i.e., 0.60 or 60% for MMLE versus 0.46 or 46% for MUE based on the Quandt (1960) Likelihood Ratio, henceforth QLR, structural break test statistic). For MPLE, this frequency is 45 percentage points higher at 0.91 (91%). At $\lambda = 5$ ($\sigma_{\Delta\beta} = 0.01$) and $\lambda = 10$ ($\sigma_{\Delta\beta} = 0.02$), these differences in the *'pile-up'* at zero frequencies reduce to 11 and 4 percentage points, respectively, for MMLE, but remain still sizeable for MPLE. At $\lambda = 20$ ($\sigma_{\Delta\beta} = 0.04$), the *'pile-up'* at zero problem disappears nearly entirely for MMLE and MUE, with *'pile-up'* frequencies dropping to 2 and 1 percentage points, respectively, for these two estimators, staying somewhat higher at 7 percentage points for MPLE.

Using MUE instead of MLE to mitigate *'pile-up'* at zero problems comes, nevertheless, at a cost; that is, a loss in estimator efficiency whenever λ (or $\sigma_{\Delta\beta}$) is not *very* small. From Table 2 on page 353 in Stock and Watson (1998), which shows the asymptotic relative efficiency of MUE (and MPLE) relative to MMLE, it is evident that for true λ values of 10 or greater ($\sigma_{\Delta\beta} \geq 0.02$), the 4 different MUEs yield asymptotic relative efficiencies (AREs) as low as 0.65 (see the results under the L and MW columns in Table 2). This means that MMLE only needs 65% of MUE's sample size to achieve the same efficiency. Only for very small values of $\lambda \leq 4$ ($\sigma_{\Delta\beta} \leq 0.008$) are the AREs of MUE and MMLE of a similar magnitude, i.e., close to 1, suggesting that both estimators achieve approximately the same precision.

The following main points are to be taken away from the above review of the simulation results in Stock and Watson (1998). Since, HLW do not estimate the initial condition of the state-vector and instead use very tightly specified priors, the potential for a significant reduction in the *'pile-up'* at zero frequencies is rather limited, and depends on the true size of λ in the data. To get an intuitive visual impression of the *'pile-up'* at zero probabilities, we can plot the information from Table 1 in Stock and Watson (1998), with the λ values deflated by the sample size T to make them comparable to the estimates in HLW. Fig. 2 shows these *'pile-up'* at zero frequencies (y-axis), versus various λ/T values (x-axis) for MMLE and the 4 different MUE estimators. The two empirical values that HLW estimate for the U.S. in their first two stages are 0.053 and 0.030 for λ_g and λ_z, respectively (see Table 6 on page S73 in Holston et al. (2017)). Evidently, the increase in the *'pile-up'* at zero probability of MMLE relative to MUE is rather small for λ_z, and effectively 0 for λ_g. Thus, given these sizeable MUE estimates of λ_g and λ_z in HLW, we can expect — a priori — the likelihood of *pile-up* at zero problems materializing with MMLE to be small, therefore precluding the need to employ MUE in the first place.

2.3. HLW's three stage estimation procedure

Holston et al. (2017) employ MUE in two preliminary stages that are based on restricted versions of the full model in (1) to obtain estimates of the *'signal-to-noise'* ratios $\lambda_g = \sigma_g/\sigma_{y^*}$ and $\lambda_z = a_r\sigma_z/\sigma_{\tilde{y}}$. These ratios are then held fixed in Stage 3 of their procedure, which produces estimates of the remaining parameters of the model in (1).

HLW adopt the general state–space model (SSM) notation of Hamilton (1994) in their three stage procedure. The SSM is formulated as follows:[3]

$$\begin{aligned}\mathbf{y}_t &= \mathbf{A}\mathbf{x}_t + \mathbf{H}\boldsymbol{\xi}_t + \boldsymbol{\nu}_t \\ \boldsymbol{\xi}_t &= \mathbf{F}\boldsymbol{\xi}_{t-1} + \mathbf{S}\boldsymbol{\epsilon}_t\end{aligned}, \quad \text{where } \begin{bmatrix}\boldsymbol{\nu}_t \\ \boldsymbol{\epsilon}_t\end{bmatrix} \sim \text{MNorm}\left(\begin{bmatrix}\mathbf{0} \\ \mathbf{0}\end{bmatrix}, \begin{bmatrix}\mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}\end{bmatrix}\right), \tag{5}$$

where I define $\boldsymbol{\epsilon}_t = \mathbf{S}\boldsymbol{\epsilon}_t$, so that $\text{Var}(\boldsymbol{\epsilon}_t) = \text{Var}(\mathbf{S}\boldsymbol{\epsilon}_t) = \mathbf{SWS}' = \mathbf{Q}$ to make it consistent with the notation used in Holston et al. (2017). The (observed) measurement vector is denoted by $\mathbf{y}_t$ in (5), $\mathbf{x}_t$ is a vector of exogenous variables, $\boldsymbol{\xi}_t$ is the latent state vector, $\mathbf{S}$ is a selection matrix. The disturbance terms $\boldsymbol{\nu}_t$ and $\boldsymbol{\epsilon}_t$ are serially uncorrelated, and the (individual) covariance matrices $\mathbf{R}$ and $\mathbf{W}$ are assumed to be diagonal matrices. The measurement vector $\mathbf{y}_t$ in (5) is the same for all three stages and is defined as $\mathbf{y}_t = [y_t,\ \pi_t]'$, where y_t and π_t are the log of real GDP and annualized PCE inflation, respectively, as defined in Section 2.1.

As emphasized in the description of MUE in Section 2.2, the simulation results of Stock and Watson (1998) show that *'pile-up'* at zero frequencies for MLE are not only a function of the size of the variance of $\Delta\beta_t = (\lambda/T)\eta_t$ (or alternatively λ), but also depend critically on whether the initial condition of the state vector is estimated or not. Now HLW *do not* estimate the initial condition of the state vector in any of the three stages that are implemented. Instead, they apply the HP filter to log GDP data with the smoothing parameter set to 36000 to get a preliminary estimate of y_t^* and trend growth g_t (computed as the first difference of the HP filter estimate of y_t^*) using data from 1960:Q1 onwards. *'Other factor'* z_t is initialized at 0. This means that $\boldsymbol{\xi}_{00}$ has known and fixed quantities in all three estimation stages. Also, Holston et al. (2017) determined the covariance matrix of the initial state vector in a somewhat unorthodox way. Even though every element of the state vector $\boldsymbol{\xi}_t$ in all three estimation stages is (at least) an $I(1)$ variable, they do not use a diffuse prior on the state vector. Instead, the covariance matrix is determined effectively as 0.2 times an identity matrix (for more details on the exact computations that are carried out by HLW to find the initial state vector and covariance matrix, see Section 4 in Buncic (2021)).

2.4. Stage 1 model

The first stage model takes the following restricted form of the full model presented in Eq. (1):

$$y_t = y_t^* + \tilde{y}_t \tag{6a}$$

$$\pi_t = b_\pi\pi_{t-1} + \left(1 - b_\pi\right)\pi_{t-2,4} + b_y\tilde{y}_{t-1} + \epsilon_t^\pi \tag{6b}$$

$$\tilde{y}_t = a_{y,1}\tilde{y}_{t-1} + a_{y,2}\tilde{y}_{t-2} + \mathring{\varepsilon}_t^{\tilde{y}} \tag{6c}$$

$$y_t^* = g + y_{t-1}^* + \mathring{\varepsilon}_t^{y^*}, \tag{6d}$$

where the vector of Stage 1 parameters to be estimated is:

$$\boldsymbol{\theta}_1 = [a_{y,1},\ a_{y,2},\ b_\pi,\ b_y,\ g,\ \sigma_{\tilde{y}},\ \sigma_\pi,\ \sigma_{y^*}]'. \tag{7}$$

Details on the exact matrix SSM form are given in Appendix A.1 in Buncic (2021).

[3] The state–space form that they use is described on pages 9 to 11 of their online appendix that is included with the R-Code HLW_Code.zip file from Williams' website at the New York Fed. For convenience of notation, I do not transpose the system matrices $\mathbf{A}$ and $\mathbf{H}$ in (5), as HLW do.

2.5. Stage 2 model

The second stage model in HLW consists of the following system of equations, which are again a restricted version of the full model in (1) (See Appendix A.2 in Buncic (2021) for the matrix SSM form of the Stage 2 model and how it is re-arranged to arrive at Eq. (8) below):[4]

$$y_t = y_t^* + \tilde{y}_t \tag{8a}$$

$$\pi_t = b_\pi\pi_{t-1} + \left(1 - b_\pi\right)\pi_{t-2,4} + b_y\tilde{y}_{t-1} + \epsilon_t^\pi \tag{8b}$$

$$a_y(L)\tilde{y}_t = a_0 + \tfrac{a_r}{2}(r_{t-1} + r_{t-2}) + a_g g_{t-1} + \mathring{\varepsilon}_t^{\tilde{y}} \tag{8c}$$

$$y_t^* = y_{t-1}^* + g_{t-2} + \mathring{\varepsilon}_t^{y^*} \tag{8d}$$

$$g_{t-1} = g_{t-2} + \epsilon_{t-1}^g. \tag{8e}$$

Given an estimate of λ_g, the vector of Stage 2 parameters to be estimated by MLE is:

$$\boldsymbol{\theta}_2 = [a_{y,1},\ a_{y,2},\ a_r,\ a_0,\ a_g,\ b_\pi,\ b_y,\ \sigma_{\tilde{y}},\ \sigma_\pi,\ \sigma_{y^*}]'. \tag{9}$$

Examining the formulation of the Stage 2 model in (8) and comparing it to the full model in (1), it is evident that HLW make two 'misspecification' choices that are important to highlight. First, they include g_{t-2} instead of g_{t-1} in the trend equation in (8d), so that the misspecified error term $\mathring{\varepsilon}_t^{y^*}$ is in fact:

$$\begin{aligned}\mathring{\varepsilon}_t^{y^*} &= \epsilon_t^{y^*} + \overbrace{g_{t-1} - g_{t-2}}^{\epsilon_{t-1}^g \text{ from (8e)}} \\ &= \epsilon_t^{y^*} + \epsilon_{t-1}^g.\end{aligned} \tag{10}$$

Due to the ϵ_{t-1}^g term in (10), the covariance between the two error terms in (8d) and (8e) is no longer zero, but rather σ_g^2. Thus, treating W in (5) as a diagonal variance–covariance matrix in the estimation of the second stage model is incorrect.[5]

Second, HLW account for only one lag in trend growth g_t and do not impose the $a_g = -4a_r$ restriction in the estimation of a_g.[6] Due to this, the error term $\mathring{\varepsilon}_t^{\tilde{y}}$ in (8c) can now be seen to consist of the following two components:

$$\begin{aligned}\mathring{\varepsilon}_t^{\tilde{y}} &= \overbrace{-a_r(L)4g_t - a_r(L)z_t + \epsilon_t^{\tilde{y}}}^{\text{missing true model part}} - \overbrace{(a_0 + a_g g_{t-1})}^{\text{added Stage 2 part}} \\ &= \underbrace{-a_r(L)z_t + \epsilon_t^{\tilde{y}}}_{\text{desired terms}} - \underbrace{\left[a_0 + a_g g_{t-1} + a_r(L)4g_t\right]}_{\text{unnecessary terms}},\end{aligned} \tag{11}$$

where the 'desired terms' on the right-hand side of (11) are needed for HLW's implementation of MUE in the second stage, while the 'unnecessary terms' are purely due to the ad hoc changing of the lag structure on g_t and failure to impose the $a_g = -4a_r$ restriction, as well as the addition of an intercept term.

To be consistent with the full model specification in (1), the relations in (8c) and (8d) should have been formulated as:

$$a_y(L)\tilde{y}_t = a_r(L)[r_t - 4g_t] + \mathring{\varepsilon}_t^{\tilde{y}} \tag{12a}$$

$$y_t^* = y_{t-1}^* + g_{t-1} + \mathring{\varepsilon}_t^{y^*}, \tag{12b}$$

so that only the two missing lags of z_t from (12a) appear in the error term $\mathring{\varepsilon}_t^{\tilde{y}}$, specifically:

$$\mathring{\varepsilon}_t^{\tilde{y}} = -a_r(L)z_t + \epsilon_t^{\tilde{y}}. \tag{13}$$

[4] I use the ring symbol ($\mathring{}$) on the disturbance terms in (8c) and (8d) to distinguish them from the *i.i.d.* error terms of the full model in (1).

[5] In the post COVID-19 updated version of the model in Holston et al. (2023), they correct this part of the Stage 2 model misspecification (See Appendix A1: State–Space Models on pages 41–42 in Holston et al. (2023)).

[6] They also add an intercept term a_0 to the output gap equation in (8c), which, according to the full model specification, is superfluous.

Table 1
Summary statistics of the λ_z estimates obtained from applying Holston et al.'s (2017) Stage 2 MUE procedure to simulated data.

Summary statistics	(1) DGPs when θ_2 held fixed at $\hat{\theta}_2$	(2) DGPs when θ_2 held fixed at $\hat{\theta}_2$	(3) DGPs when θ_2 is re-estimated	(4) DGPs when θ_2 is re-estimated
	$r_t^* = 4g_t$	$r_t^* = 4g_t + z_t$	$r_t^* = 4g_t$	$r_t^* = 4g_t + z_t$
Minimum	0.000000	0.000000	0.000000	0.000000
Maximum	0.101220	0.096427	0.116886	0.116445
Standard deviation	0.016245	0.016582	0.018512	0.019647
Mean	0.028842	0.030726	0.025103	0.027462
Median	0.028394	0.029609	0.022215	0.025115
$\Pr(\hat{\lambda}_z^s > 0.030217)$	0.457000	0.490000	0.339000	0.393000

Notes: This table reports summary statistics of the λ_z estimates that one obtains from implementing Holston et al.'s (2017) Stage 2 MUE procedure on artificial data that was simulated from two different data generating processes (DGPs). The first DGP simulates data from the full structural model in (1) under the parameter estimates of Holston et al. (2017), but where the natural rate is determined solely by trend growth. That is, in the output gap equation in (1c), $r_t^* = 4g_t$. The second DGP simulates data from the full model of Holston et al. (2017) where $r_t^* = 4g_t + z_t$. The summary statistics that are reported are the minimum, maximum, standard deviation, mean, median, as well as the empirical frequency of observing a value larger than the estimate of 0.030217 obtained by Holston et al. (2017), denoted by $\Pr(\hat{\lambda}_z^s > 0.030217)$. The table shows four different estimates, grouped in 2 block pairs. The left block under the heading 'DGPs when θ_2 is held fixed' shows the simulation results for the two DGPs when the Stage 2 parameter vector θ_2 is held fixed at the Stage 2 estimates and is not re-estimated on the simulated data. The right block under the heading 'DGPs when θ_2 is re-estimated' shows the simulation results when θ_2 is re-estimated for each simulated series. Simulations are performed on a sample size equivalent to the empirical data, with 1000 repetitions.

That is, the correctly specified Stage 2 model should have been formulated as:

$$y_t = y_t^* + \tilde{y}_t \tag{14a}$$

$$\pi_t = b_\pi \pi_{t-1} + \left(1 - b_\pi\right) \pi_{t-2,4} + b_y \tilde{y}_{t-1} + \varepsilon_t^\pi \tag{14b}$$

$$a_y(L)\tilde{y}_t = a_r(L)[r_t - 4g_t] + \mathring{\varepsilon}_t^{\tilde{y}} \tag{14c}$$

$$y_t^* = y_{t-1}^* + g_{t-1} + \varepsilon_t^{y^*} \tag{14d}$$

$$g_{t-1} = g_{t-2} + \varepsilon_{t-1}^g. \tag{14e}$$

2.6. Simulation evidence

In Buncic (2022) (see Section 2.3) it is shown algebraically that MUE applied to the misspecified model yields to the signal-to-noise ratio parameter $\lambda_z = \frac{a_r \sigma_z}{\bar{\sigma}(\tilde{y}_t^{\tilde{y}})}$, where $\bar{\sigma}(\cdot)$ denotes the long-run standard deviation and the misspecified error term of the output-gap equation is $\mathring{v}_t^{\tilde{y}} = \varepsilon_t^{\tilde{y}} - [a_0 + \frac{(a_g + 4a_r)}{2}(g_{t-1} + g_{t-2}) + \frac{a_g}{2}\varepsilon_{t-1}^g]$ in the local-level model for Stage 2 MUE. Note that, even under the assumption of $(a_g + 4a_r) = 0$, the relation for $\mathring{v}_t^{\tilde{y}}$ yields the extra term $\frac{a_g}{2}\sigma_g$ in λ_z, giving $\lambda_z = \frac{a_r \sigma_z}{(\sigma_{\tilde{y}} + a_g \sigma_g/2)}$. This shows that MUE applied to the misspecified Stage 2 model does not recover the signal-to-noise ratio of interest $\frac{\lambda_z \sigma_{\tilde{y}}}{a_r}$ as used in Stage 3 of the estimation in HLW.

In this section, I provide simulation evidence to show that MUE as applied in HLW's (misspecified) Stage 2 model leads to spuriously large λ_z estimates when computed from simulated data generated from a model that sets $\lambda_z = \sigma_z = 0$. To illustrate this, I implement two simulation experiments. In the first experiment, I simulate data from the full structural model in (1) using the Stage 3 parameter estimates from HLW as the true values, but with '*other factor*' z_t set to zero for all t (or alternatively, $\lambda_z = \sigma_z = 0$). The natural rate r_t^* in the output gap equation in (1c) is thus solely determined by (annualized) trend growth, that is, $r_t^* = 4g_t$. I then implement HLW's Stage 2 MUE procedure on the simulated data to yield a sequence of $S = 1000$ estimates of λ_z $\left(\{\hat{\lambda}_z^s\}_{s=1}^S\right)$.

I use two different scenarios for the Stage 2 parameters θ_2 in the Kalman Smoother recursions used to extract the latent cycle as well as trend growth series needed for the construction of the dummy variable regressions used in MUE. The first scenario simply takes HLW's empirical Stage 2 estimate of θ_2 and keeps those values fixed for all 1000 generated data sequences when applying the Kalman Smoother. In the second scenario, I re-estimate the Stage 2 model parameters for each simulated sequence to obtain new estimates $\hat{\theta}_2^s, \forall s = 1, \ldots, S$. I then apply the Kalman Smoother using these estimates in the dummy variable regressions.

Finally, I repeat the above computations on data that were generated from the full model in (1) with the natural rate of interest determined by both factors, namely, $r_t^* = 4g_t + z_t$, where z_t was simulated as a pure random walk, with its standard deviation set at the Stage 3 estimate of $\sigma_{\tilde{y}}$ and a_r, i.e., at $\sigma_z = \lambda_z \sigma_{\tilde{y}}/a_r \approx 0.15$ (see row σ_z (implied) of column one in Table 3). The goal here is to provide a comparison of the range of λ_z estimates that can be obtained when implementing HLW's Stage 2 MUE procedure on data that were generate with and without '*other factor*' z_t in the data generating processes (DGPs) of the natural rate.

In Table 1, summary statistics of the $\{\hat{\lambda}_z^s\}_{s=1}^S$ from the two different DGPs are reported. The left column block shows results for the two different DGPs when the Stage 2 parameter vector θ_2 is held fixed at HLW's estimates. The right column block shows corresponding results when θ_2 is re-estimated for each simulated data series. The summary statistics are the minimum, maximum, standard deviation, mean, and median of $\{\hat{\lambda}_z^s\}_{s=1}^S$, as well as the relative frequency of obtaining a value larger than the empirical point estimate in HLW. This point estimate and the corresponding relative frequency are denoted by $\hat{\lambda}_z^{\text{HLW}}$ and $\Pr(\hat{\lambda}_z^s > \hat{\lambda}_z^{\text{HLW}})$, respectively. To complement the summary statistics in Table 1, histograms of $\hat{\lambda}_z^s$ are shown in Fig. 3 to provide visual information about its sampling distribution.

From the summary statistics in Table 1 as well as the histograms in Fig. 3 it can be seen how similar the $\hat{\lambda}_z^s$ estimates from these two different DGPs are. For instance, when the data were simulated without '*other factor*' z_t (i.e., $\lambda_z = 0$), the sample mean of $\hat{\lambda}_z^s$ is 0.028842. When the data were generated from the full model with $r_t^* = 4g_t + z_t$, the sample mean of $\hat{\lambda}_z^s$ is only 6.53% higher at 0.030726. Similarly, the relative frequencies $\Pr(\hat{\lambda}_z^s > \hat{\lambda}_z^{\text{HLW}})$ for these two DGPs are 45.70% and 49%, respectively. The inclusion of '*other factor*' z_t in the DGP of the natural rate thus results in only a 3.3 percentage points higher $\Pr(\hat{\lambda}_z^s > \hat{\lambda}_z^{\text{HLW}})$.[7] The histograms in Fig. 3 paint the same overall picture. The Stage 2 MUE implementation has difficulties to discriminate between these two DGPs.

In a second experiment, I simulate DGPs from entirely unrelated univariate ARMA processes to make the individual components needed

[7] When the Stage 2 parameter vector θ_2 is re-estimated for each simulated sequence shown in the right column block in Table 1, the sample means as well as the relative frequency $\Pr(\hat{\lambda}_z^s > \hat{\lambda}_z^{\text{HLW}})$ are somewhat lower at 0.025103 and 0.027462, and 33.90% and 39.30%, respectively.

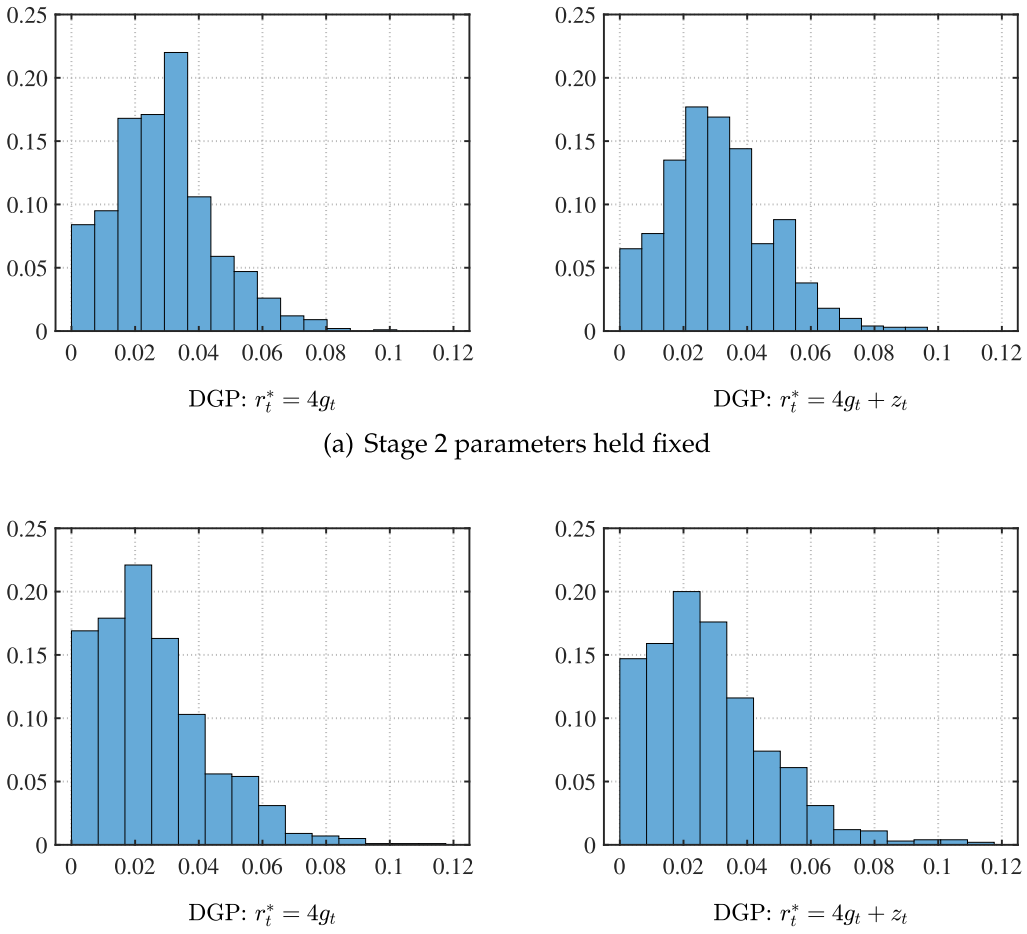

Fig. 3. Histograms of the estimated $\{\hat{\lambda}_z^s\}_{s=1}^S$ sequence corresponding to the summary statistics shown in Table 1. On the left and right sides of the figure, histograms for the two different DGPs are shown. To top two histograms show the results when θ_2 is held fixed in the simulations and is not re-estimated, while the bottom plots show the results when θ_2 is re-estimated on each simulated series that is generated.

Table 2
Summary statistics of λ_z estimates of the Stage 2 MUE procedure applied to data simulated from unrelated univariate ARMA processes.

Summary statistic	$g_{t-1} = \hat{g}_{t-1\|T}$	$g_{t-1} \sim$ RW	$g_{t-1} \sim$ WN	$\Delta g_{t-1} \sim$ ARMA
Minimum	0.000000	0.000000	0.000000	0.000000
Maximum	0.097019	0.095914	0.096789	0.093340
Standard deviation	0.015240	0.015858	0.016803	0.016335
Mean	0.031798	0.029708	0.026117	0.030449
Median	0.030165	0.028647	0.024254	0.029435
$\Pr(\lambda_z^s > 0.030217)$	0.498000	0.456000	0.384000	0.482000

Notes: This table reports summary statistics of the Stage 2 estimates of λ_z that one obtains when applying Holston et al.'s (2017) MUE procedure to simulated data without the z_t process. The summary statistics that are reported are the minimum, maximum, standard deviation, mean, median, as well as the empirical frequency of observing a value larger than the estimate of 0.030217 obtained by Holston et al. (2017), denoted by $\Pr(\hat{\lambda}_z^s > 0.030217)$. The columns show the estimates for the four different data generating processes for trend growth g_t. The first column reports results when the Kalman Smoothed estimate $\hat{g}_{t-1|T}$ is used for g_{t-1}. The second and third columns show estimates when g_{t-1} is generated as a pure random walk (RW) or (Gaussian) white noise (WN) process. The last column reports results when g_{t-1} is computed as the cumulative sum of Δg_{t-1}, which is simulated from the coefficients obtained from a low order ARMA process fitted to $\Delta \hat{g}_{t-1|T}$. The cycle and real rate series are also constructed by first finding the best fitting low order ARMA processes to the individual series and then simulating from fitted coefficients.

for the structural break regressions. To match the time series properties of the time series that enter these regressions, I fit simple low-order ARMA models to $\hat{\tilde{y}}_{t|T}$, r_t and $\hat{g}_{t|T}$, and then use these ARMA estimates to simulate artificial data. I then again apply HLW's Stage 2 MUE procedure to the simulated data as before. The full results from the second experiment are reported in Table 2 and Fig. 4, respectively.

As can be seen from the results in Table 2 and histograms in Fig. 4, the magnitudes of $\hat{\lambda}_z^s$ are similar to those from the first simulation experiment, with mean estimates being between 0.026117 and 0.031798,

Table 3
Stage 3 parameter estimates.

θ_3	(1) HLW.R-File	(2) Replicated	(3) MLE($\sigma_g \mid \lambda_z^{\text{HLW}}$)	(4) MLE($\sigma_g \mid \lambda_z^{\mathcal{M}_0}$)	(5) MLE(σ_g, σ_z)
$a_{y,1}$	1.52957249	1.52957247	1.49442462	1.49566712	1.49566147
$a_{y,2}$	−0.58756415	−0.58756414	−0.55370268	−0.55448942	−0.55448212
a_r	−0.07119569	−0.07119569	−0.07941598	−0.07525496	−0.07525240
b_π	0.66820705	0.66820705	0.67128197	0.66919468	0.66919993
b_y	0.07895778	0.07895778	0.07593604	0.08054901	0.08054716
$\sigma_{\tilde{y}}$	0.35346845	0.35346847	0.36043114	0.37381376	0.37382935
σ_π	0.78919487	0.78919487	0.79029982	0.78948921	0.78949094
σ_{y^*}	0.57241925	0.57241924	0.55915743	0.55293818	0.55293018
σ_g (implied)	(0.03083567)	(0.03083567)	0.04583852	0.04497450	0.04497414
σ_z (implied)	(0.15002080)	(0.15002080)	(0.13714150)	(0.00374682)	0.00000001
λ_g (implied)	0.05386904	0.05386904	(0.08197784)	(0.08133730)	(0.08133782)
λ_z (implied)	0.03021722	0.03021722	0.03021722	0.00075430	(0.00000000)
Log-likelihood	−515.14470528	−515.14470599	−514.83070544	−514.28987426	−514.28958969

Notes: This table reports replication results for the Stage 3 model parameter vector θ_3 of Holston et al. (2017). The first column (HLW.R-File) reports estimates obtained by running Holston et al.'s (2017) R-Code for the Stage 3 model. The second column (Replicated) shows the replicated results using the same set-up as in Holston et al. (2017). The third column (MLE($\sigma_g \mid \lambda_z^{\text{HLW}}$)) reports estimates when σ_g is directly estimated by MLE together with the other parameters of the Stage 3 model, while λ_z is held fixed at $\lambda_z^{\text{HLW}} = 0.030217$ obtained from Holston et al.'s (2017) *"misspecified"* Stage 2 procedure. In the fourth column (MLE($\sigma_g \mid \lambda_z^{\mathcal{M}_0}$)), σ_g is again estimated directly by MLE together with the other parameters of the Stage 3 model, but with λ_z now fixed at $\lambda_z^{\mathcal{M}_0} = 0.000754$ obtained from the *"correctly specified"* Stage 2 model in (14). The last column (MLE(σ_g, σ_g)) shows estimates when all parameters are computed by MLE. Values in round brackets give the implied $\{\sigma_g, \sigma_z\}$ or $\{\lambda_g, \lambda_z\}$ values when either is fixed or estimated. The last row (Log-likelihood) reports the value of the log-likelihood function at these parameter estimates.

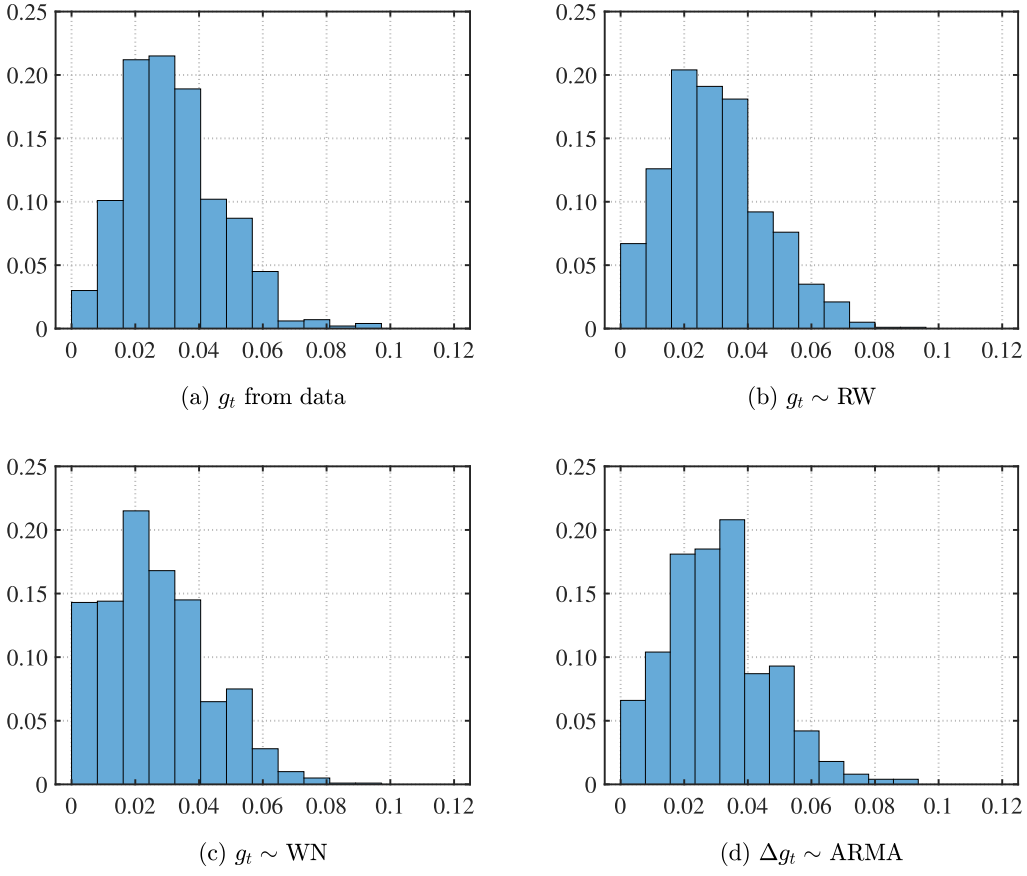

Fig. 4. Histograms of the estimated $\{\hat{\lambda}_z^s\}_{s=1}^S$ sequence corresponding to the summary statistics shown in Table 2.

and relative frequencies corresponding to $\Pr(\hat{\lambda}_z^s > \hat{\lambda}_z^{\text{HLW}})$ being between 38.40% and 49.80%. Moreover, there are many instances where the estimates of λ_z from the simulated data are not only non-zero, but rather sizeable, being larger than the estimate of $\lambda_z = 0.030217$ that HLW compute from the empirical data. Note that there is no z_t process simulated, yet with HLW's Stage 2 MUE procedure one can recover an

estimate that is at least as large as the empirical one around 40 to 50 percent of the time, depending on how g_t is simulated. This simulation exercise highlights that HLW's MUE procedure can lead to spuriously large estimates of λ_z when the true value is 0. As the downward trend in the z_t process drives the movement in the natural rate, where the severity of the downward trend is related to the magnitude of σ_z (through λ_z), HLW's estimates of the natural rate are likely to be downward biased.

3. Natural rate estimates from the correct Stage 2 model

The analysis so far has shown that, due to the misspecification in the Stage 2 model, HLW's MUE procedure overestimates λ_z when it is in fact equal to zero in the DGP. Moreover, from the results in Buncic (2022) and Buncic (2021) it is clear that firstly, λ_z from HLW's misspecified model cannot recover the signal-to-noise ratio of interest $\lambda_z = a_r\sigma_z/\sigma_{\tilde{y}}$. And secondly, the Stage 1 model is not needed for σ_g as its MMLE based estimate never shrinks to zero (see Figure 1 in Buncic (2022) and the discussion in Section 4.1 in Buncic (2021)).

This section provides corrected estimates of the natural rate as well as trend growth, '*other factor*' z_t and the output gap from HLW's structural model.[8] I initially report the relevant parameter estimates in Table 3, and then plot the Kalman filtered and smoothed estimates in Fig. 5 and Fig. 6, respectively. Table 3 is structured as follows. The column with the heading $\boldsymbol{\theta}_3$ lists the full model's parameters $\boldsymbol{\theta}_3 = [a_{y,1},\ a_{y,2},\ a_r,\ b_\pi,\ b_y,\ \sigma_{\tilde{y}},\ \sigma_\pi,\ \sigma_{y^*}]'$. Estimates of the ratios $\lambda_g = \sigma_g/\sigma_{y^*}$ and $\lambda_z = a_r\sigma_z/\sigma_{\tilde{y}}$ from the previous two stages are reported in either direct estimate form (without parenthesis) or in implied from (with parenthesis) when estimated indirectly via MUE or backed out from the MLE estimates using the relevant ratio of interest. In columns (1) and (2) estimates obtained from HLW's R-files and my replicated results are reported. The signal-to-noise ratios λ_g and λ_z are computed following the exact format of HLW. Values in parenthesis for σ_g and σ_z are the implied values obtained from the signal-to-noise ratios, solved for σ_g and σ_z. In column (3) I report parameter estimates following the format of HLW in Stage 2, but where the Stage 1 step is skipped and σ_g is estimated directly by MLE. The notation 'MLE$(\sigma_g|\hat{\lambda}_z^{\text{HLW}})$' signifies this in the tables and in the plots. In column (4) the estimates from the correctly specified Stage 2 model are shown, where σ_g is again estimated directly by MLE. The notation 'MLE$(\sigma_g|\lambda_z^{\mathcal{M}_0})$' is used for that. Lastly, pure MLE based estimates of the full structural model of HLW are reported in column (5), using the notation MLE(σ_g, σ_z).

The results in Table 3 can be summarized as follows. Firstly, the MLE of σ_g does not '*pile-up*' at zero (as expected) and is approximately 50% larger than the estimate implied by the Stage 1 MUE of λ_g. That is, $\hat{\sigma}_g \approx 0.045$ in the last three columns of Table 3. The MLE of σ_z shrinks numerically to zero, while the estimates of the other parameters remain largely unchanged. Notice that the log-likelihood values of the last three models in Table 3 are very similar, i.e., between −514.8307 and −514.2899. Yet, the corresponding estimates of σ_z are either very small at 0, or comparatively large at 0.1371 when computed (implied) from the 'misspecified' Stage 2 model's $\hat{\lambda}_z^{\text{HLW}}$ estimate. This is a discrepancy that is unexpected from the simulation results reported in Stock and Watson (1998). The implied $\hat{\sigma}_z$ coefficient from the 'correctly specified' Stage 2 model shown in column (4) is 0.0037 and thereby nearly 40 times smaller than from the 'misspecified' Stage 2 model. The MLE of σ_z shown in column (5) shrinks to zero, as expected.

The findings from Table 3 are mirrored in the filtered (and smoothed) estimates of r_t^*, g_t, z_t and $\tilde{y}_t$ plotted in Fig. 5 (and Fig. 6). Note that the 'MLE$(\sigma_g|\lambda_z^{\mathcal{M}_0})$' and 'MLE$(\sigma_g, \sigma_z)$' based estimates of '*other factor*' z_t are visually indistinguishable, that is, the overlap for the entire time period, despite having different sized estimates of σ_z of 0.00374682 and 0.00000001. This is entirely consistent with the findings in Stock and Watson (1998). The MUE from the correct Stage 2 specification implies a much closer to zero estimate of '*other factor*' z_t, mirroring the simulation based findings reported earlier. Also, and unsurprisingly, out of the four latent state estimates, '*other factor*' z_t is overall most strongly affected by the two different λ_z (σ_z) values that are conditioned upon, showing either vary large variability and a pronounced downward trend in z_t, or being close to zero with very little variation (see panel (c) in Fig. 5). The effect on the estimate of the natural rate is largest in the immediate aftermath of the global financial crisis, namely, from 2010 onwards. Interestingly, the output gap estimates shown in panel (d) of Fig. 5 are quite similar, with the largest divergence occurring after 2012. The three trend growth estimates in panel (b) of Fig. 5 where σ_g is estimated directly by MLE are visually indistinguishable. Trend growth estimated from HLW's Stage 1 MUE of λ_g is larger from 2009 to 2014. The pure backward looking nature of the Kalman Filtered g_t series exacerbates the effect of the decline in GDP during the financial crisis on the trend growth estimates after the crisis.

In support of the earlier results on the misspecification in the Stage 2 model of HLW and its effect on λ_z, I re-estimate HLW's model with data ending in 2019:Q2. Estimation results, together with corresponding plots of filtered (and smoothed) estimates are reported in Table 4, Fig. 7 and Fig. 8, respectively. The objective of re-estimating the model on updated data and reporting the results here is to provide additional evidence that λ_z computed from HLW's misspecified Stage 2 model cannot be correct with regards to its magnitude. The ML estimates shown in column (5) of Table 3 give a prime example of the '*pile-up*' at zero problem with MLE. In column (5) of Table 4 we see that $\hat{\sigma}_z$ estimated by MLE does *not* shrink to zero in the extended sample period. In fact, the point estimate is comparable in magnitude to the one obtained from the correctly specified Stage 2 model MUE. That is, they are 0.05250494 and 0.06065982, respectively, while the MUE from HLW's implementation yields the much larger (implied) $\hat{\sigma}_z$ of 0.15642060. This provides extra empirical evidence that the Stage 2 MUE procedure of HLW leads to spuriously large estimates of λ_z.

4. Other issues with HLW's structural model

This section raises five other issues with HLW's structural model when aimed for policy analysis. These are structured around: (1) reverse causality (in the structural model at least), that is, the nominal fed funds target rate i_t is driving the movements of the natural rate, not the other way around, (2) excessive sensitivity of the estimates to the chosen sample starting date, (3) what is an appropriate process for other factor z_t, (4) should the focus be on smoothed or filtered estimates of the latent states, and (5) can we recover the shocks driving the natural rate and therefore the natural rate itself?

4.1. Reverse causality: the policy rate is driving the natural rate

In HLW's model, the policy rate i_t is included as an exogenous variable. With $r_t^* = 4g_t + z_t$, and '*other factor*' z_t the free variable due to g_t being driven by GDP growth, z_t effectively matches the leftover movements in the interest rate to make it compatible with trend growth in the model. Since the central bank has full control over the policy rate, it can set i_t to any desired level, and the model will produce a natural rate through '*other factor*' z_t that will match it. Further, there is nothing in the structural model of (1) that makes the system stable. For the output gap relation in (1c) to be stationary, the real rate cycle $\tilde{r}_t = r_t - r_t^* = (i_t - \pi_t^e) - (4g_t + z_t)$ must be $I(0)$, yet there is no co-integrating relation imposed anywhere in the system to ensure that this holds in the model.[9]

[8] Replication files and the factors themselves are available at https://github.com/4db83/Issues-with-HLWs-natural-rate-Code and https://www.danielbuncic.com/data/correct.HLW.factors.zip, respectively.

[9] This insight is not new and has been discussed in, for instance, Pagan and Wickens (2019) (see pages 21 – 23). When trying to simulate from such

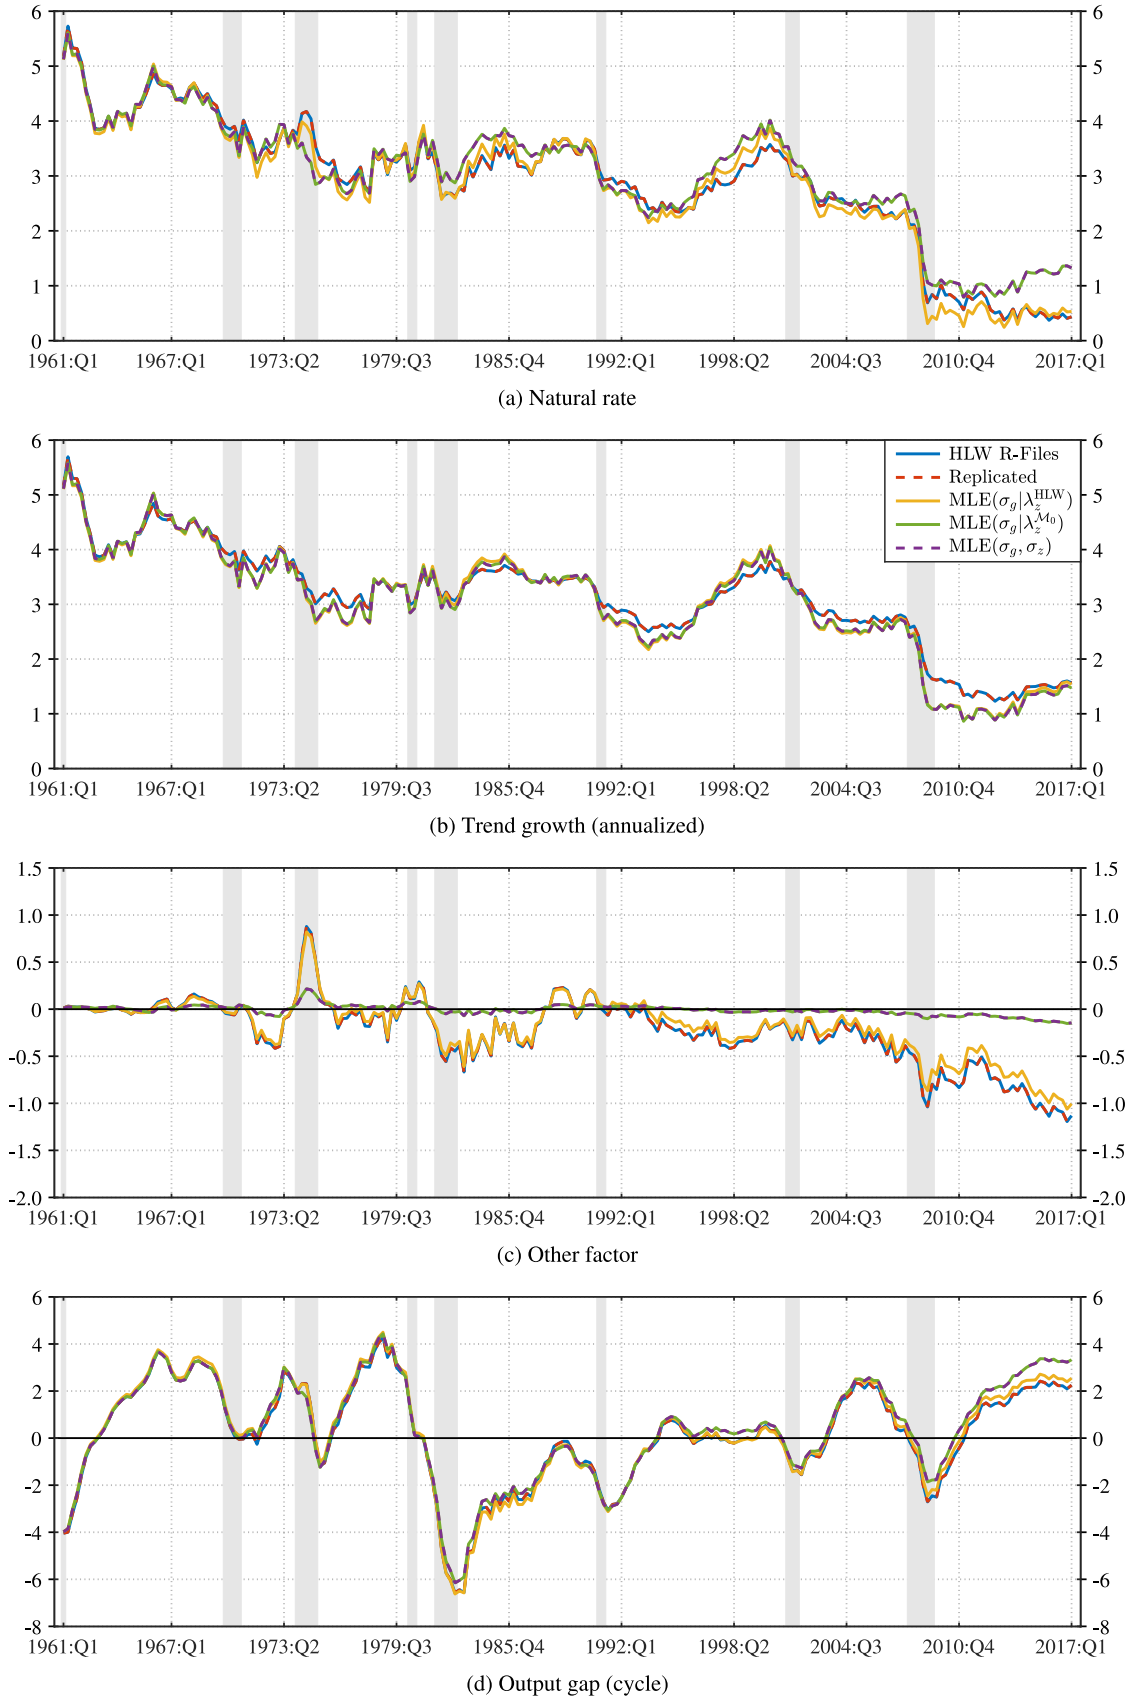

(a) Natural rate

(b) Trend growth (annualized)

(c) Other factor

(d) Output gap (cycle)

Fig. 5. Filtered estimates of the natural rate r_t^*, annualized trend growth g_t, *'other factor'* z_t, and the output gap (cycle) variable $\tilde{y}_t$.

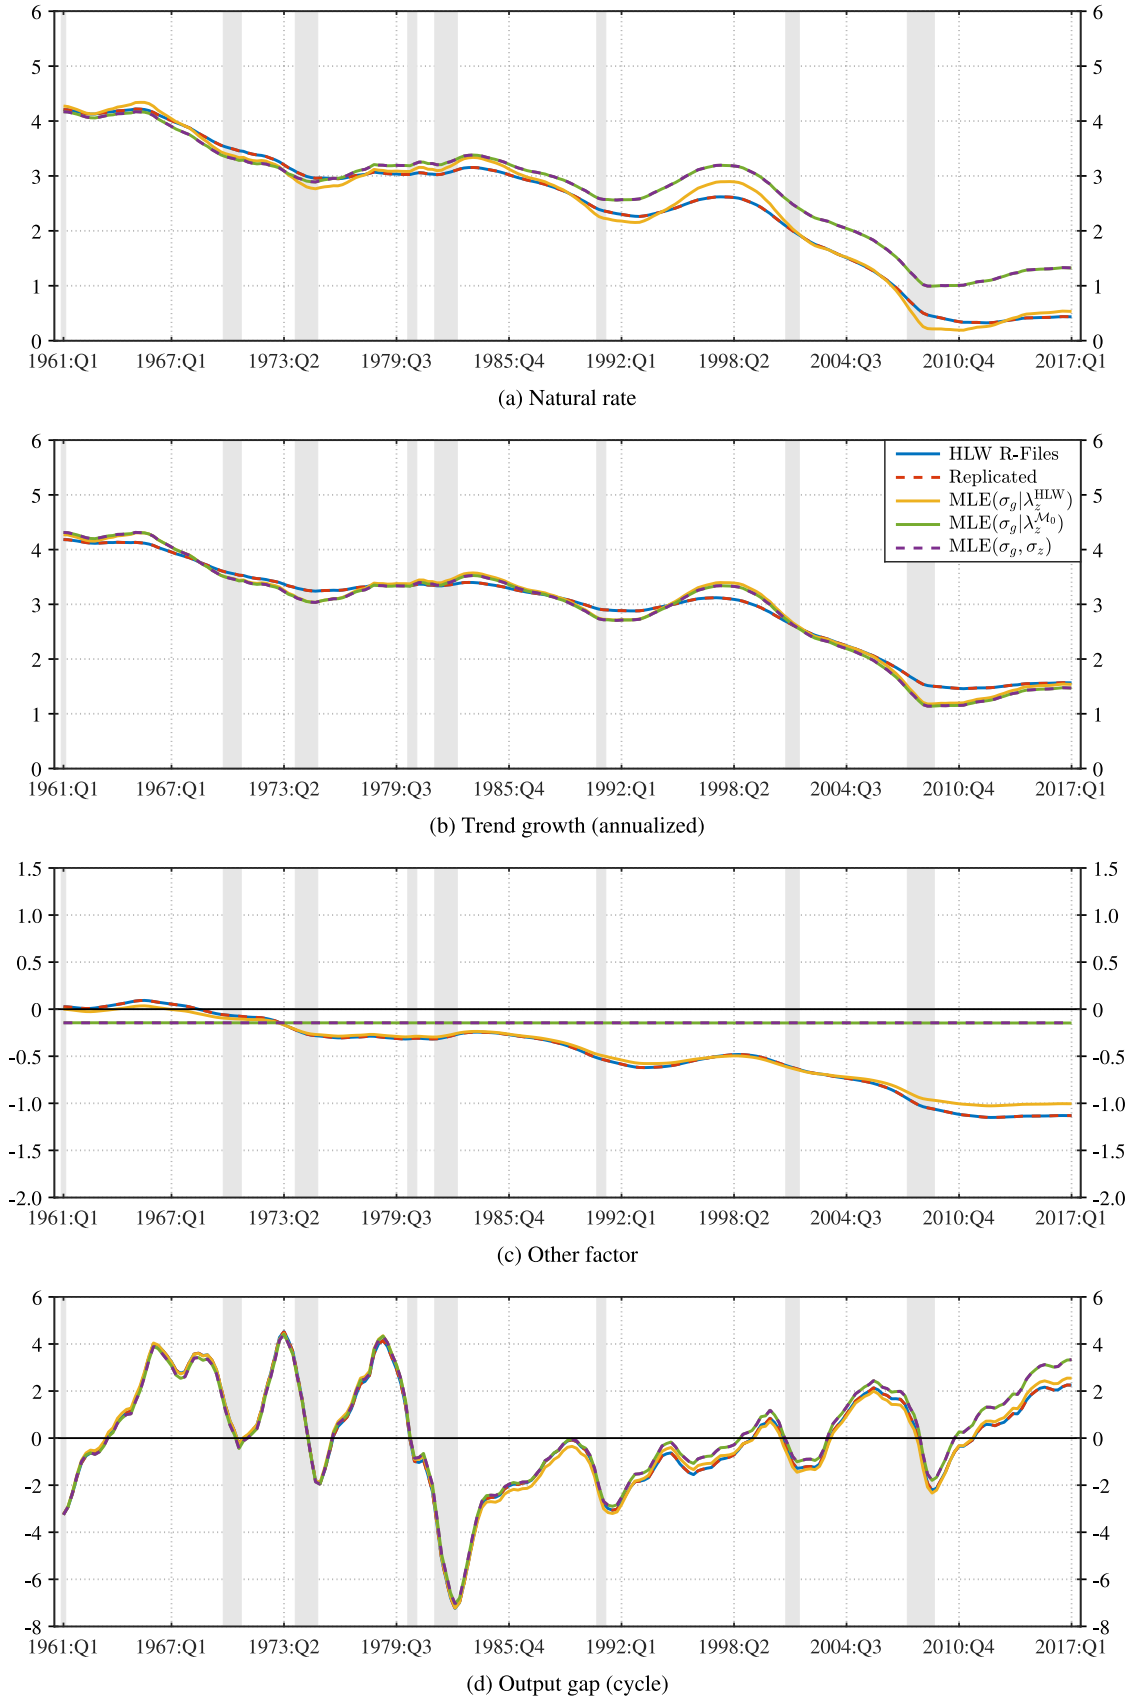

(a) Natural rate

(b) Trend growth (annualized)

(c) Other factor

(d) Output gap (cycle)

Fig. 6. Smoothed estimates of the natural rate r_t^*, annualized trend growth g_t, *'other factor'* z_t, and the output gap (cycle) variable $\tilde{y}_t$.

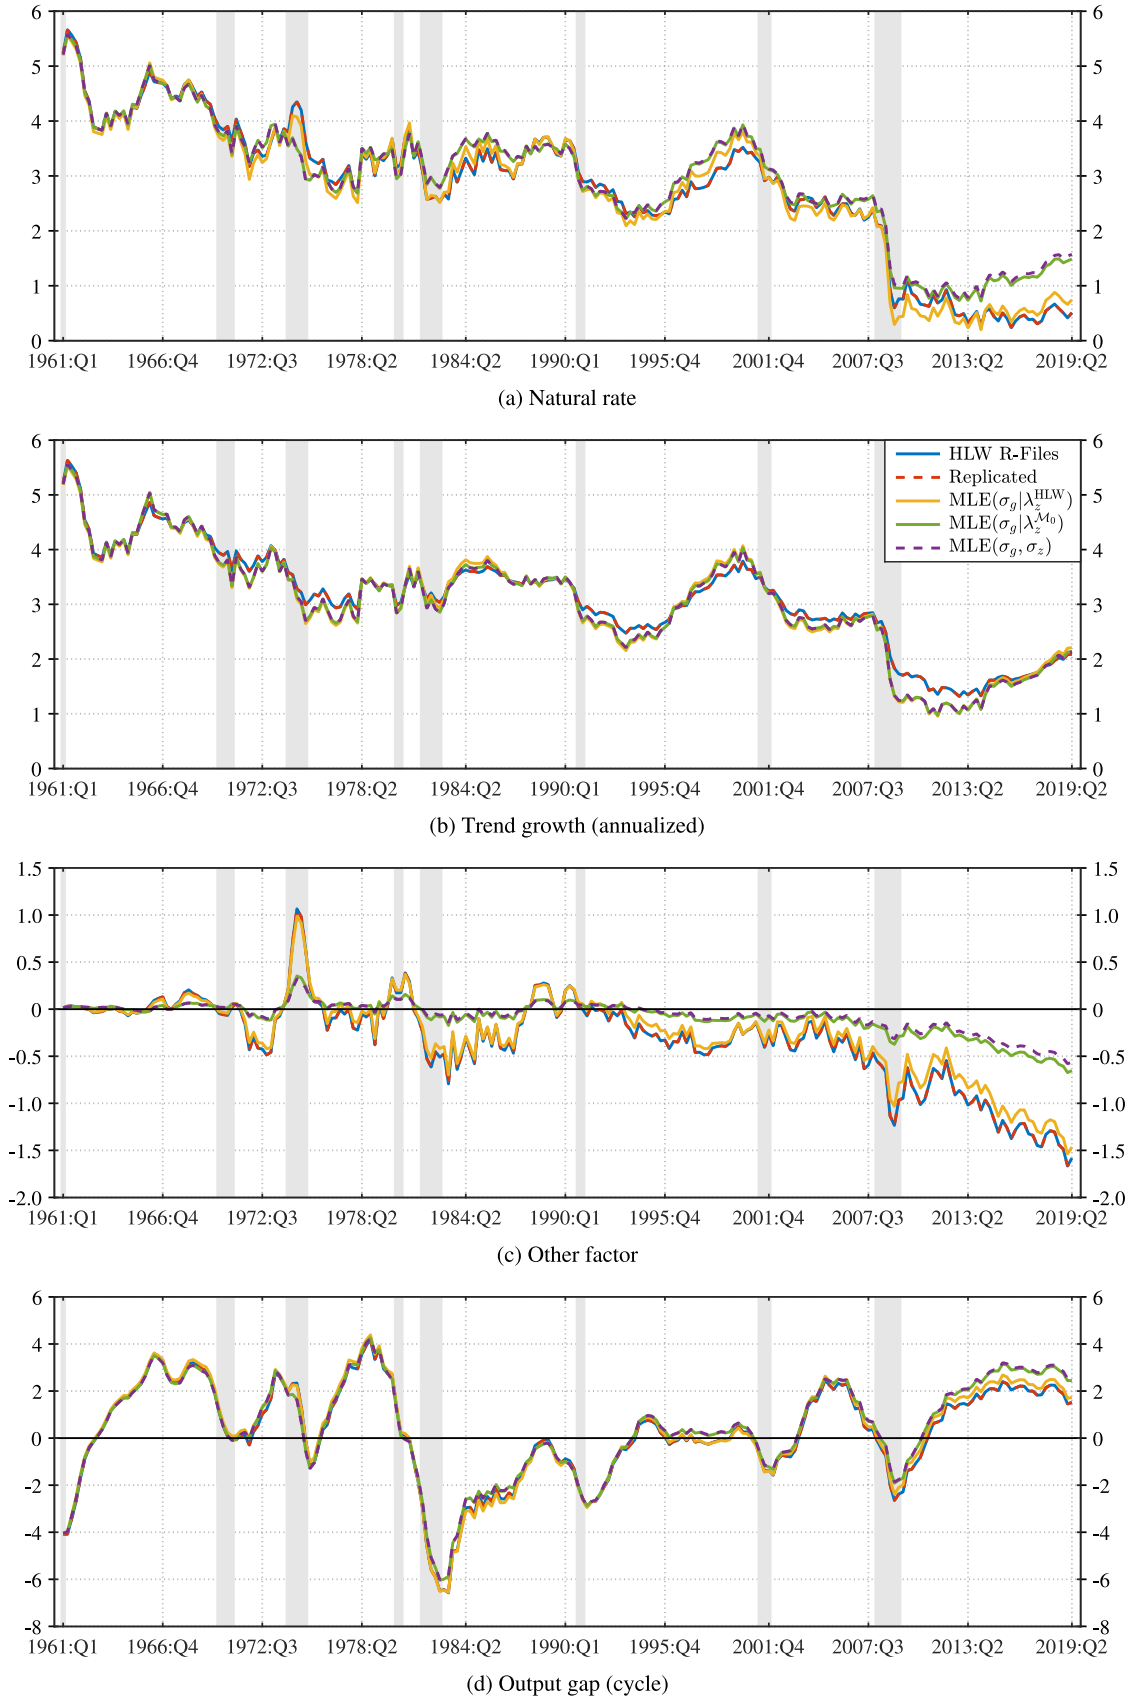

(a) Natural rate

(b) Trend growth (annualized)

(c) Other factor

(d) Output gap (cycle)

Fig. 7. Filtered estimates of the natural rate r_t^*, annualized trend growth g_t, *'other factor'* z_t, and the output gap (cycle) variable $\tilde{y}_t$ up to 2019:Q2.

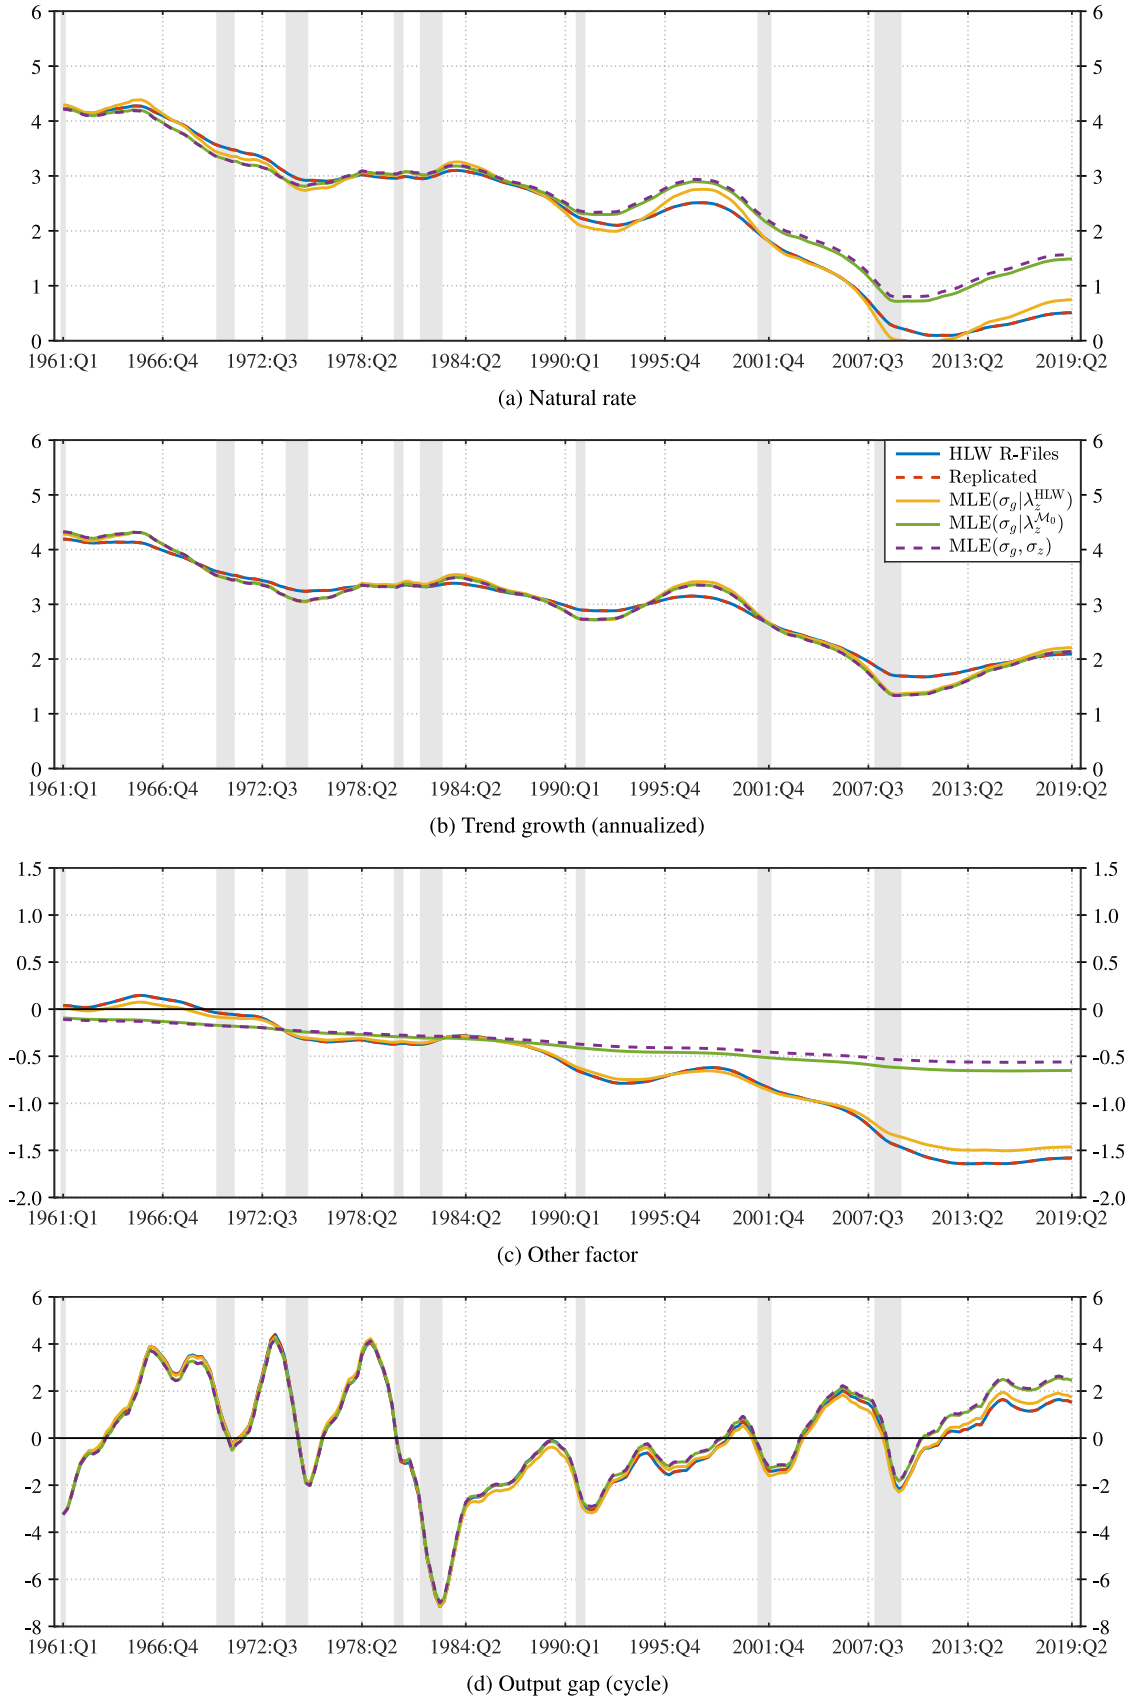

(a) Natural rate

(b) Trend growth (annualized)

(c) Other factor

(d) Output gap (cycle)

Fig. 8. Smoothed estimates of the natural rate r_t^*, annualized trend growth g_t, *'other factor'* z_t, and the output gap (cycle) variable $\tilde{y}_t$ up to 2019:Q2.

Table 4
Stage 3 parameter estimates using data up to 2019:Q2.

θ_3	(1) HLW.R-File	(2) Replicated	(3) MLE($\sigma_g \mid \lambda_z^{\text{HLW}}$)	(4) MLE($\sigma_g \mid \lambda_z^{\mathcal{M}_0}$)	(5) MLE(σ_g, σ_z)
$a_{y,1}$	1.53876458	1.53876459	1.51083223	1.51659115	1.51669620
$a_{y,2}$	−0.59700264	−0.59700265	−0.57053684	−0.57637540	−0.57645931
a_r	−0.06854043	−0.06854043	−0.07561113	−0.07029671	−0.07000675
b_π	0.67331545	0.67331545	0.67638900	0.67463450	0.67483411
b_y	0.07755450	0.07755451	0.07454055	0.07888427	0.07885243
$\sigma_{\tilde{y}}$	0.33590693	0.33590692	0.33598381	0.34747670	0.34826106
σ_π	0.78812554	0.78812554	0.78921814	0.78854952	0.78862036
σ_{y^*}	0.57577319	0.57577320	0.56789520	0.56359236	0.56327803
σ_g (implied)	(0.03082331)	(0.03082331)	0.04517849	0.04386169	0.04378982
σ_z (implied)	(0.17251762)	(0.17251762)	0.15642060	0.06065982	0.05250494
λ_g (implied)	0.05353377	0.05353377	(0.07955428)	(0.07782519)	(0.07774103)
λ_z (implied)	0.03520151	0.03520151	(0.03520151)	(0.01227186)	(0.01055443)
Log-likelihood	−533.36984524	−533.36984550	−533.16547501	−532.82874860	−532.82637541

Notes: This table reports replication results for the Stage 3 model parameter vector θ_3 of Holston et al. (2017). The first column (HLW.R-File) reports estimates obtained by running Holston et al.'s (2017) R-Code for the Stage 3 model. The second column (Replicated) shows the replicated results using the same set-up as in Holston et al. (2017). The third column (MLE($\sigma_g \mid \lambda_z^{\text{HLW}}$)) reports estimates when σ_g is directly estimated by MLE together with the other parameters of the Stage 3 model, while λ_z is held fixed at $\lambda_z^{\text{HLW}} = 0.035202$ obtained from Holston et al.'s (2017) *"misspecified"* Stage 2 procedure. In the fourth column (MLE($\sigma_g \mid \lambda_z^{\mathcal{M}_0}$)), σ_g is again estimated directly by MLE together with the other parameters of the Stage 3 model, but with λ_z now fixed at $\lambda_z^{\mathcal{M}_0} = 0.012272$ obtained from the *"correctly specified"* Stage 2 model in (14). The last column (MLE(σ_g, σ_g)) shows estimates when all parameters are computed by MLE. Values in round brackets give the implied $\{\sigma_g, \sigma_z\}$ or $\{\lambda_g, \lambda_z\}$ values when either is fixed or estimated. The last row (Log-likelihood) reports the value of the log-likelihood function at these parameter estimates.

Kalman filter based estimates of the state vector ξ_t will be (weighted combinations of the) one-sided moving averages of the three observed variables that enter the state–space model; namely, i_t, y_t, and π_t. This can be seen by writing out the Kalman Filtered estimate of the state vector as:

$$\begin{aligned}
\hat{\boldsymbol{\xi}}_{t|t} &= \hat{\boldsymbol{\xi}}_{t|t-1} + \underbrace{\mathbf{P}_{t|t-1}\mathbf{H}'(\mathbf{H}\mathbf{P}'_{t|t-1}\mathbf{H}' + \mathbf{R})^{-1}}_{\mathbf{G}_t}(\mathbf{y}_t - \mathbf{A}\mathbf{x}_t - \mathbf{H}\hat{\boldsymbol{\xi}}_{t|t-1}) \\
&= \hat{\boldsymbol{\xi}}_{t|t-1} + \mathbf{G}_t(\mathbf{y}_t - \mathbf{A}\mathbf{x}_t - \mathbf{H}\hat{\boldsymbol{\xi}}_{t|t-1}) \\
&= (\mathbf{I} - \mathbf{G}_t\mathbf{H})\hat{\boldsymbol{\xi}}_{t|t-1} + \mathbf{G}_t(\mathbf{y}_t - \mathbf{A}\mathbf{x}_t) \\
&= \underbrace{(\mathbf{I} - \mathbf{G}_t\mathbf{H})\mathbf{F}}_{\boldsymbol{\Phi}_t}\hat{\boldsymbol{\xi}}_{t-1|t-1} + \mathbf{G}_t\underbrace{(\mathbf{y}_t - \mathbf{A}\mathbf{x}_t)}_{\bar{\mathbf{y}}_t} \\
&= \boldsymbol{\Phi}_t\hat{\boldsymbol{\xi}}_{t-1|t-1} + \mathbf{G}_t\bar{\mathbf{y}}_t,
\end{aligned}$$

which is a (linear) recursion in $\hat{\boldsymbol{\xi}}_{t|t}$ and can thus be rewritten as:

$$\begin{aligned}
\hat{\boldsymbol{\xi}}_{t|t} &= \boldsymbol{\Psi}_t\boldsymbol{\xi}_{0|0} + \sum_{i=0}^{t-1}\underbrace{\boldsymbol{\Psi}_i\mathbf{G}_{t-i}}_{\boldsymbol{\omega}_{ti}}\bar{\mathbf{y}}_{t-i} \\
&= \boldsymbol{\Psi}_t\boldsymbol{\xi}_{0|0} + \sum_{i=0}^{t-1}\boldsymbol{\omega}_{ti}\bar{\mathbf{y}}_{t-i},
\end{aligned} \tag{15}$$

where $\Psi_i = \prod_{n=0}^{i-1}\boldsymbol{\Phi}_{t-n}, \forall i = 1, 2, \ldots, \Psi_0 = I$, I is the identity matrix, $\hat{\xi}_{t|t-1} = F\hat{\xi}_{t-1|t-1}$ is the predicted state vector, $\xi_{0|0}$ is the prior mean, $P_{t|t-1} = FP_{t-1|t-1}F + Q$ is the predicted state variance, $\omega_{ti} = \Psi_i G_{t-i}$ is a time varying weight matrix, and $\bar{\mathbf{y}}_t$ is a (2×1) vector containing the observed variables y_t, π_t, and i_t.

As the nominal interest rate i_t is contained in $\bar{\mathbf{y}}_t$ in (15) and is directly controlled by the central bank, a circular relationship can be seen to evolve. Any central bank induced change in the policy rate i_t is mechanically transferred to the natural rate r_t^* via the Kalman Filtered estimate of the state vector $\hat{\xi}_{t|t}$ in (15). This confounds the relationship between r_t^* and i_t, making it impossible to address causal policy questions of interest such as: "*Is the natural rate low because i_t is low, or is i_t low because the natural rate is low?*", as one follows as a direct consequence from the other (in this model).

4.2. Excessive sensitivity to the chosen sample starting date

Because of the one-sided moving average nature of the Kalman Filtered estimates of the state vector, any outliers, structural breaks or otherwise extreme observations at the beginning (or end) of the sample period can have a strong impact on these filtered estimates. For the (two-sided) Hodrick and Prescott (1997) filter, such problems (and other ones) are well known and have been discussed extensively in the literature. Kalman Filter based (one-sided) estimates will also be affected. A simple way to appreciate this is by re-estimating the HLW model using four different starting dates, while keeping the end of the sample period the same at 2019:Q2. In Fig. 9 I show filtered estimates of r_t^*, g_t, z_t and $\tilde{y}_t$ for the four starting dates 1967:Q1, 1972:Q1, 1952:Q2 and 1947:Q1 (smoothed estimates are shown in Fig. 10), together with HLW's estimates using 1961:Q1 as the starting date.[10]

Why did I choose these starting dates? The period following the April 1960 to February 1961 recession was marked by temporarily (and unusually) high GDP growth, yielding an annualized mean of 6.07% (median 6.47%), with a low standard deviation of 2.67% from 1961:Q2 to 1966:Q1. Having such excessive growth at the beginning of the sample period has an unduly strong impact on the filtered (less so on the smoothed) estimate of trend growth g_t in the model. Since both g_t and z_t enter the natural rate, this affects the estimate of r_t^*. To illustrate the sensitivity of these estimates to this time period, I estimate the model with data starting 6 years later in 1967:Q1. Also, HLW's Euro Area estimates of r_t^* are negative from around 2013 onwards (see the bottom panel of Figure 3 on page S63 of their paper). To show that one can get the same negative estimates of r_t^* for the U.S., I estimate the model with data starting in 1972:Q1 to match the sample period used for the Euro Area. Lastly, I extend HLW's data back to 1947:Q1 to have estimates from a very long sample, using total PCE inflation prior to 1959:Q2 in place of Core PCE inflation and the Federal Reserve Bank of New York discount rate from 1965:Q1 back to 1947:Q1 as a

a model, with π_t being integrated of order 1, the simulated paths of the real rate $r_t = i_t - \pi_t^e$ can frequently diverge to very large values, even with samples of size $T = 229$ observations (the empirical sample size).

[10] In all computations, I use HLW's R-Code and follow exactly their three stage procedure as before to estimate the factors of interest.

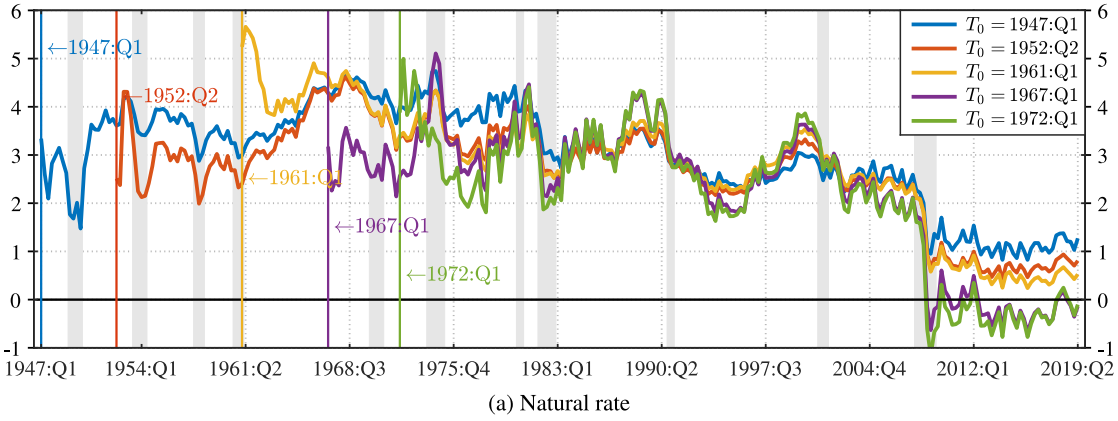

(a) Natural rate

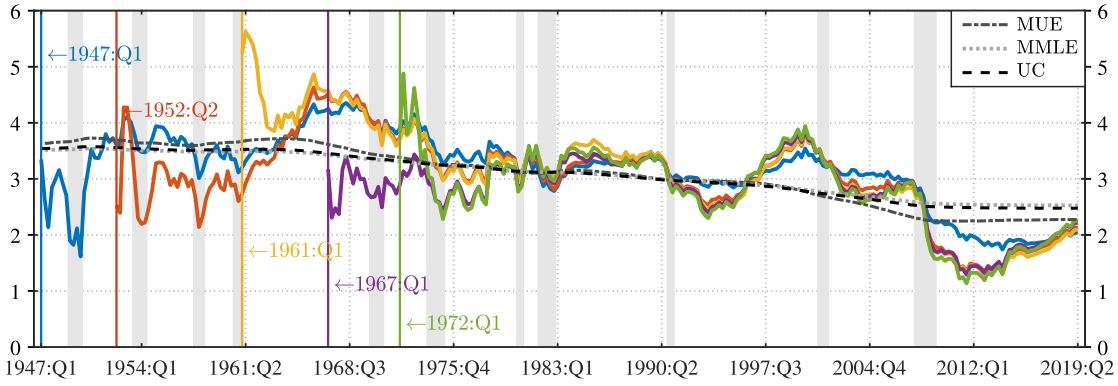

(b) Trend growth (annualized)

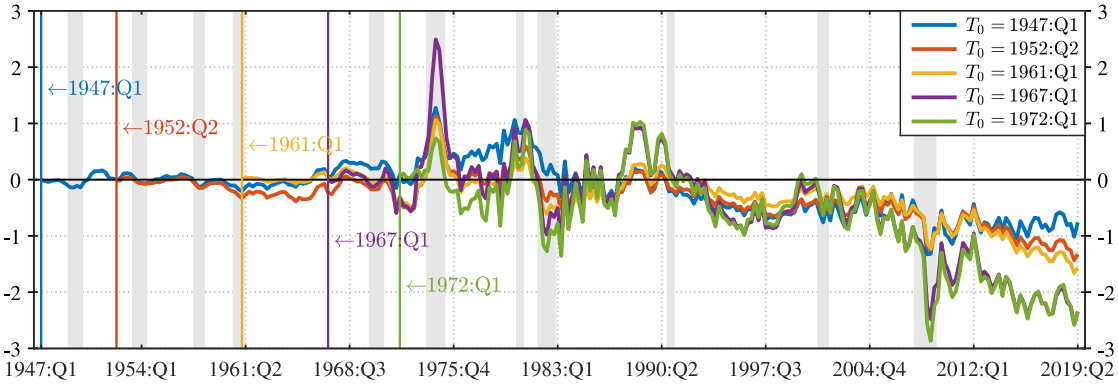

(c) Other Factor

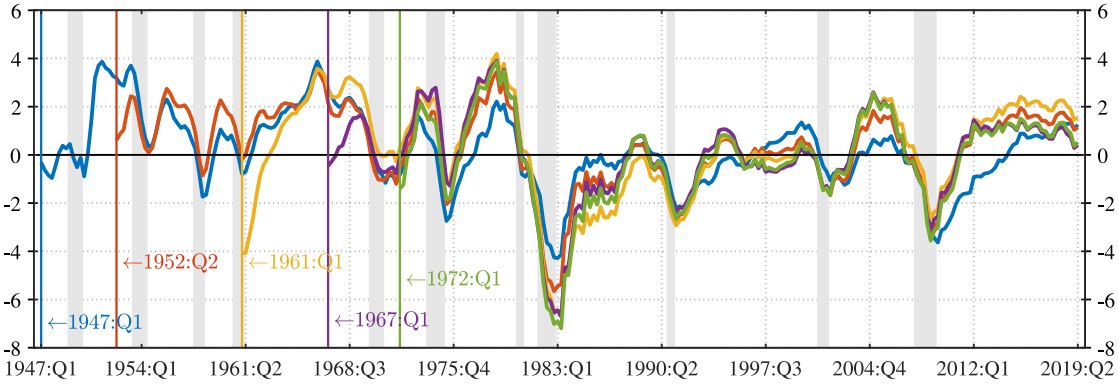

(d) Output gap (cycle)

Fig. 9. Filtered estimates of annualized trend growth g_t, *'other factor'* z_t and the natural rate r_t^* based on different starting dates.

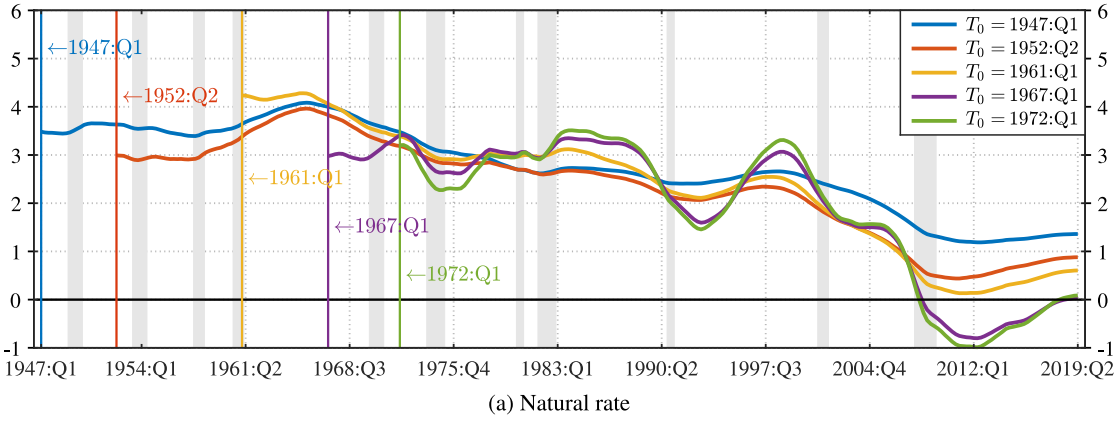

(a) Natural rate

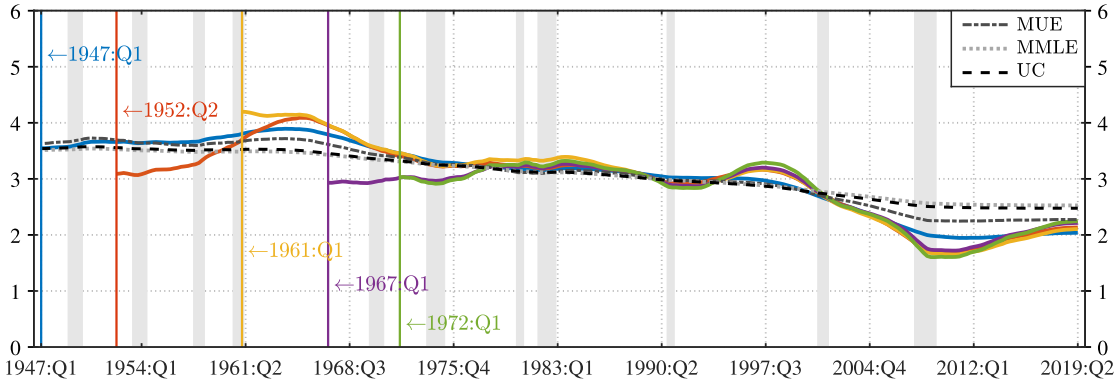

(b) Trend growth (annualized)

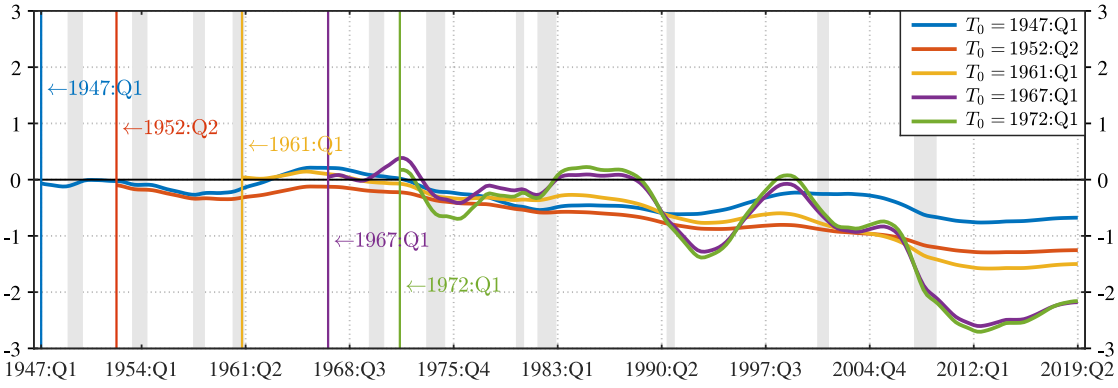

(c) Other Factor

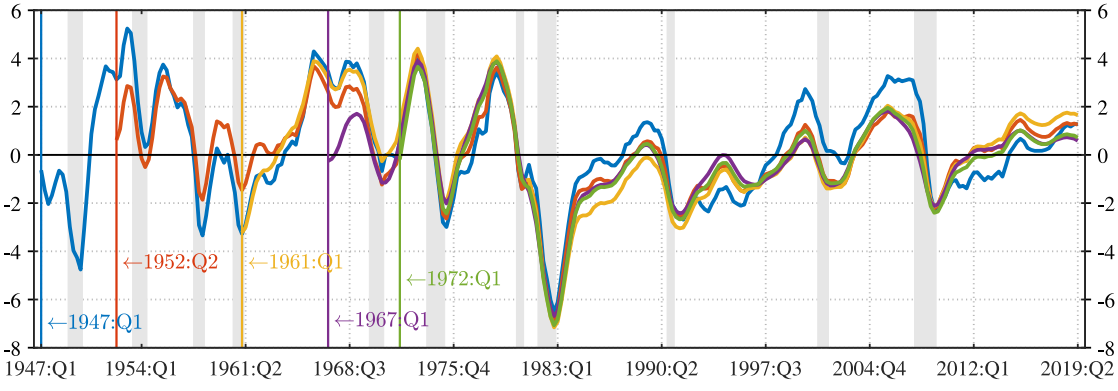

(d) Output gap (cycle)

Fig. 10. Smoothed estimates of annualized trend growth g_t, *'other factor'* z_t and the natural rate r_t^* based on different starting dates.

proxy for the Federal Funds rate, as was done in Laubach and Williams (2003). Since inflation was rather volatile from 1947 to 1952, I also estimate the model with data beginning in 1952:Q2 to exclude this volatile inflation period from the sample.

Panel (a) in Fig. 9 shows how sensitive the natural rate estimates to the different starting dates are, particularly at the beginning of HLW's sample, namely, from 1961 until about 1980, and at the end of the sample from 2009 onwards. Negative natural rate estimates are now also obtained for the U.S. when the sample starts in 1972:Q1 (or 1967:Q1). From panel (b) in Fig. 9 it is evident that the filtered trend growth estimates are the primary driver of the excessive sensitivity in r_t^* over the 1961 to 1980 period. For instance, in 1961:Q1, these estimates can be as high as 6 percent, or as low as 3 percent, depending on the starting date of the sample. Also, the differences in the estimates stay sizeable until 1972:Q1, before converging to more comparable magnitudes from approximately 1981 onwards. Apart from the estimate using the very long sample beginning in 1947:Q1 (see the blue line in panel (b) of Fig. 9), the other four series remain surprisingly similar, even during and after the financial crisis period, that is, from mid 2007 to the end of the sample in 2019:Q2. Thus, the variability of the natural rate estimates at the beginning of the sample are driven by the variability in the estimates of trend growth g_t at the beginning of the sample.

In panel (b) of Fig. 9, I add trend growth estimates computed from Stock and Watson's (1998) model and also an Unobserved Component (UC) model in the spirit of Clark (1987) and Stock and Watson (1988) to provide long-sample benchmarks of trend growth from simple univariate models to HLW's estimates. Comparing the Kalman Filter based estimates from the various starting dates to the (smoothed) MUE, MMLE, and UC ones shows how different these are, particularly, from 2009:Q3 until the end of the sample. In the immediate post-crisis period, the (one-sided) filter based estimates are pulled down excessively by the sharp decline in GDP and converge only slowly at the very end of the sample period towards the three long-sample benchmarks. Trend growth is severely underestimated from 2009:Q3 onwards, and this affects the estimate of r_t^*.

Looking at the estimates of '*other factor*' z_t in panel (c) of Fig. 9, we can see that it is the end of the sample, namely, from 2009:Q1 to 2019:Q2, that is most strongly affected by the different starting dates. In particular the two estimates that are based on the shorter samples starting in 1967:Q1 and 1972:Q1, which exclude the excessive GDP growth period at the beginning of HLW's sample, generate substantially more negative z_t estimates. For instance, in 2009:Q1, the 1972:Q1 based estimate is −2.87 while HLW's is −1.22. Also, the z_t estimates from the shorter samples are well below −2 over nearly the entire 2014:Q4 to 2019:Q2 period.[11] What is particularly surprising to observe here is how stable (and very close to zero) the estimates of z_t from the four earlier sample starts are from 1947:Q1 until about 1971:Q3. Given the change in demographics and population growth, as well as factors related to savings and investment following the end of World War II, one would expect z_t to reflect these changes. It is only from 1990:Q2 onwards that a decisive downward trend in z_t becomes visible.

HLW initialize the state vector for the z_t elements of ξ_t at zero. This choice leads to an anchoring effect and implies that the natural rate is driven solely by trend growth at the beginning of the sample. In the data, it acts as a normalization. Since z_t is specified to evolve as a driftless random walk, an initialization at zero seems sensible from an econometric perspective. Nevertheless, if one is to view '*other factor*' z_t as a factor which is meant to capture underlying structural changes in demographics, saving and investment rates and the likes, in the economy, it needs to be aligned with such trends in the empirical data of the economy that is modelled. That is, the normalization date needs to be justified from an economic perspective. Due to its large impact on the downward trend in the estimates of the natural rate, understanding how the zero initialization affects the estimates and what exactly z_t captures is crucial from a policy perspective.

4.3. What is an appropriate process for 'other factor' z_t?

'*Other factor*' z_t is specified as a driftless random walk in HLW and one might ask whether the estimate of σ_z shrinks towards zero because of the way it has been formulated. My conjecture is that the empirical data reject the $I(1)$ restriction on '*other factor*' z_t.[12] To investigate if the z_t series indeed follows an $I(1)$ process, we can examine the time series properties of the difference between the observed GDP growth series Δy_t and the real rate r_t, that is, the $(\Delta y_t - r_t)$ series. From HLW's structural model in (1) we know that:

$$\begin{aligned} r_t &= r_t^* + \tilde{r}_t \\ &= g_t + z_t + \tilde{r}_t \\ &= g_{t-1} + \varepsilon_t^g + z_t + \tilde{r}_t, \end{aligned} \tag{16}$$

where $\tilde{r}_t$ is the real rate gap defined earlier, which must be an $I(0)$ process for the output gap equation to be stationary. From the relations in (1a) and (1d) we then obtain

$$\Delta y_t = g_{t-1} + \varepsilon_t^{y^*} + \Delta \tilde{y}_t, \tag{17}$$

so that the difference between real GDP growth Δy_t in (17) and the real rate in (16) is:

$$\begin{aligned} (\Delta y_t - r_t) &= (g_{t-1} + \varepsilon_t^{y^*} + \Delta \tilde{y}_t) - (g_{t-1} + \varepsilon_t^g + z_t + \tilde{r}_t) \\ &= \underbrace{\Delta \tilde{y}_t - \tilde{r}_t + \varepsilon_t^{y^*} - \varepsilon_t^g}_{\text{stationary ARMA}} - z_t, \end{aligned} \tag{18}$$

where the variables in the first block on the right hand side of (18) will be the sum of two stationary ARMA terms (the two gaps) and two uncorrelated error terms ($\varepsilon_t^{y^*}$ and ε_t^g), while z_t is $I(1)$. In the model we thus have that $(\Delta y_t - r_t)$ is $I(1)$.

In Fig. 11 I show time series plots of Δy_t with r_t superimposed in Panel (a) in the top of the figure. The GDP growth minus real rate series $(\Delta y_t - r_t)$ is plotted in Panel (b), with the autocorrelation function (ACF) and partial autocorrelation function (PACF) plots of $(\Delta y_t - r_t)$ in Panel (c) in the bottom. Both, the time series plot in Panel (b) as well as the ACF and PACF plots in Panel (c) give the visual impression of a stationary $(\Delta y_t - r_t)$ series, with a fast decaying correlation structure. The first order autocorrelation coefficient is 0.576. More formal unit-root tests confirm that the null hypothesis of a unit-root in $(\Delta y_t - r_t)$ is strongly rejected.[13] Either '*other factor*' z_t is stationary, or it does not appear in $(\Delta y_t - r_t)$ at all, which is only possible if σ_z is zero. This supports the estimation results obtained from MLE and the correct Stage 2 model's MUE. From Panel (a) in Fig. 11 it is also visible that the real rate remained well below real GDP growth for an extended period of time in the aftermath of the financial crisis; arguably, the longest stretch in the sample, with only the period following the Dotcom bubble showing some similarities.

[11] This is even more pronounced in the smoothed estimates of z_t shown in panel (c) of Fig. 10.

[12] Lewis and Vazquez-Grande (2018) look at the likelihood of the z_t process to follow a unit-root process and find overwhelming evidence against this (see in particular Figure 4 on page 433 of their paper).

[13] Augmented Dickey–Fuller and Elliott et al. (1996) DF-GLS P_T and $t-$ tests yield test statistics of −5.33, 0.58, and −5.30, which are all substantially lower than their respective 1% critical values of −3.46, 1.92, and −2.58.

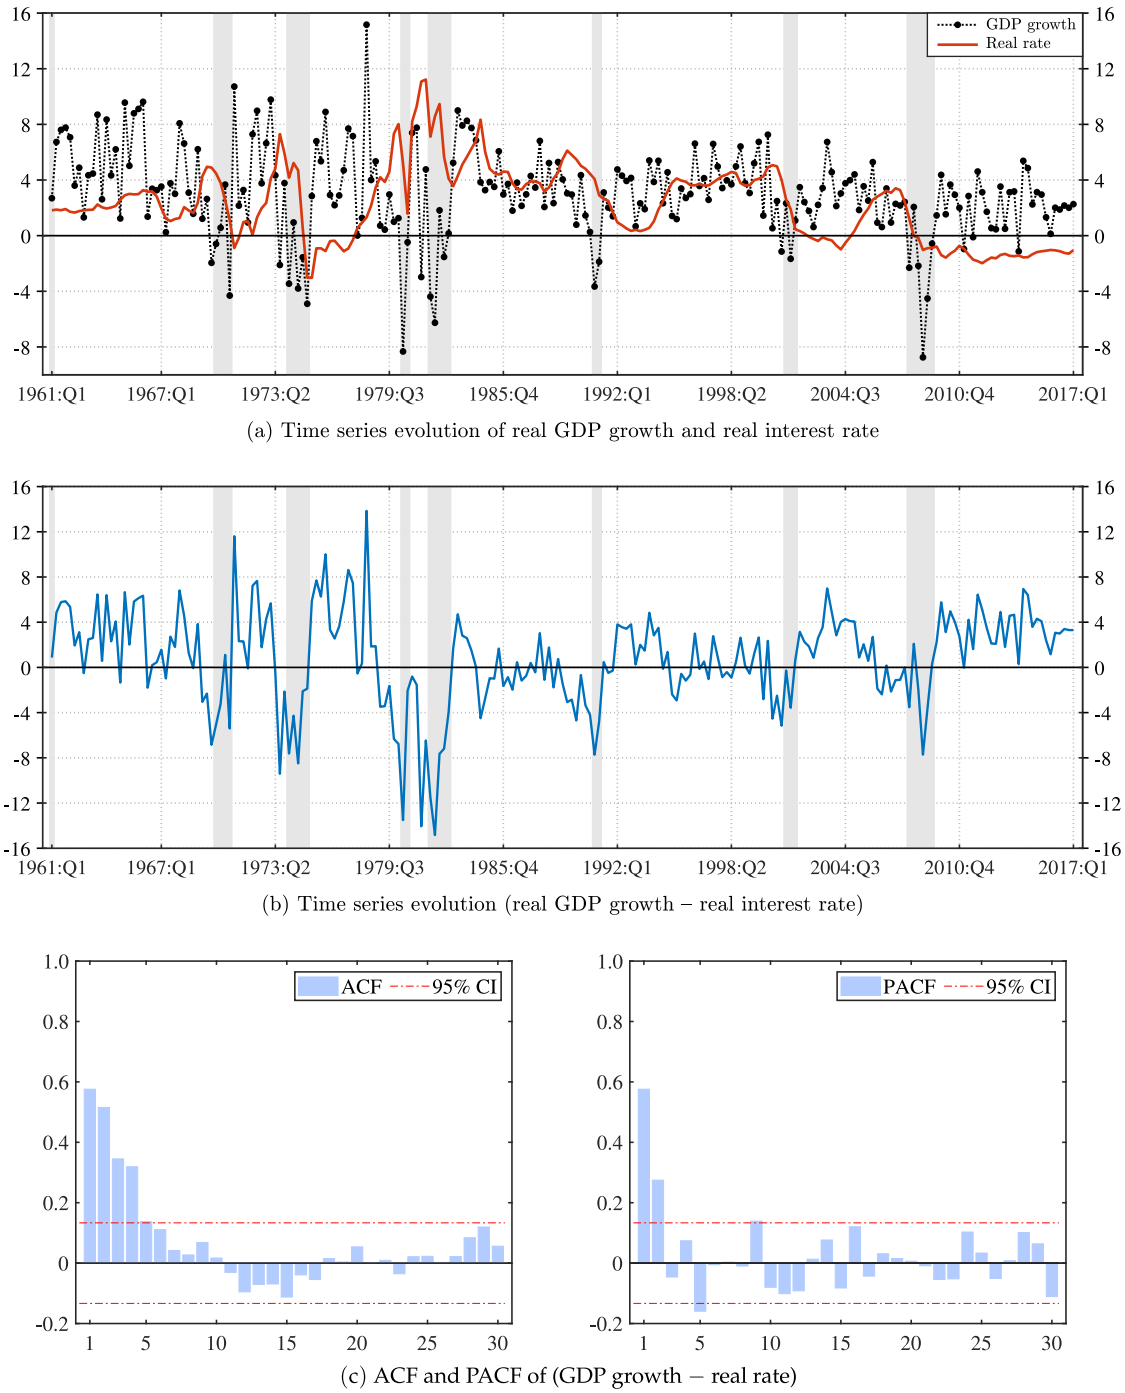

(a) Time series evolution of real GDP growth and real interest rate

(b) Time series evolution (real GDP growth − real interest rate)

(c) ACF and PACF of (GDP growth − real rate)

Fig. 11. Time series and autocorrelation plots of real GDP growth and the real interest rate.

4.4. *To filter or to smooth?*

HLW's seem to prefer to report filtered (as opposed to smoothed) estimates of the latent state variables (both in the paper as well as on the FRBNY website). This is rather surprising, as it is well known that the mean squared error (MSE) of the filtered states is in general larger than from the smoothed ones (see the discussion on page 151 in Harvey (1989)). That is, the smoother is the more accurate estimator of the latent state vector. The greater variability in the filtered states is visible from the estimates of r_t^*, g_t and z_t, and less so from the output gap (cycle) estimates. While it is sometimes stated that filtered states are real-time estimates, and are thus more relevant for policy analysis, one can see that this is a straw-man argument here. Not only are the parameter estimates of the model, i.e., the $\hat{\theta}_3$ in Table 3, based on full sample information, the GDP and PCE inflation data are not real time data, that is, data that were available to policy makers at time $t < T$. Reporting filtered (one-sided) estimates of the states as in HLW (or on the FRBNY website) gives a misleading visual impression of the magnitude of the natural rate and trend growth, which are important variables in the policy decision making process.

4.5. Can we recover the shocks driving the natural rate?

A final point to consider when using HLW's model for policy analysis is the issue of shock recovery in state–space models when there are more shocks than observables in a dynamic system with latent variables. A key finding from the recent theoretical literature on shock recovery is that it is *never* possible to recover all the shocks from such a system (see Chahrour and Jurado (2022), Pagan and Robinson (2022), and Forni et al. (2019)). In HLW's model, there are five shocks and only two observable variables (namely, inflation and GDP growth) that enter the observed measurement variable $\mathbf{y}_t$ in (5).[14] In Buncic and Pagan (2022) and Buncic et al. (2023) it is shown that the shock belonging to '*other factor*' z_t and trend growth g_t cannot be recovered from HLW's model. In fact, the Kalman smoothed estimates of the change in the natural rate, that is, Δr_t^*, can be equally well represented by the following two identities:

$$E_T \Delta r_t^* = 4E_T \varepsilon_t^g + E_T \varepsilon_t^z \tag{19}$$

$$E_T \Delta r_t^* = E_T \Delta r_{t-1}^* - 0.0495\varepsilon_t^{\tilde{y}} + 0.0003 E_T \varepsilon_t^{\pi} - 0.0037 E_T \varepsilon_t^{y^*} + 0.0231 E_T \varepsilon_{t-1}^{\tilde{y}} - 0.0078 E_T \varepsilon_t^{y^*}, \tag{20}$$

where the relation in (19) follows directly from the definition of the natural rate, and the dynamic relation in (20) involves the smoothed demand, inflation, and technology (or trend) shocks. Evidently, having two such identities makes it impossible to disentangle the nature of the shocks driving the (growth rate of the) natural rate.

5. Conclusion

This paper raises a number of econometric issues in the estimation of the natural rate of interest. More specifically, this paper shows that HLW's Stage 2 model is misspecified. Using a simulation experiment, I show that this misspecification leads to an excessively large estimate of the signal-to-noise ratio that drives the downward trend in the natural rate when applied to data generated from a model where the true value is zero. Correcting the misspecification results in a substantially smaller point estimate of the signal-to-noise ratio parameter, and thereby a more subdued trend in '*other factor*' z_t and the natural rate.

The paper also discusses various other problems that make the model unsuitable for policy analysis. For instance, estimates of the natural rate are extremely sensitive to the starting date of the sample used to fit the model. Using a sample that begins in 1972:Q1 (or 1967:Q1) leads to negative estimates of the natural rate for the U.S. These negative estimates are again driven by the exaggerated downward trending behaviour of ' *other factor*' z_t. Moreover, due to the Kalman Filtered (or Smoothed) estimates of the state vector being a function of all observable variables that enter into the model, any central bank induced change in the policy rate i_t is mechanically transferred to the natural rate r_t^*. This makes it impossible to answer causal questions regarding the relationship between r_t^* and i_t, as one responds to changes in the other. Lastly, due to the excess number of shocks in the model relative to the number of observables, shock recovery and hence natural rate recovery is problematic. This implies that it will not be possible to identify which shocks are driving the natural rate, as there exist (at least) two identities that equally well describe the evolution.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

[14] The nominal interest rate enters as an exogenous variable, thus does not add information to the state–space model.

Data availability

Data will be made available on request.

References

Berger, Tino, Kempa, Bernd, 2019. Testing for time variation in the natural rate of interest. J. Appl. Econometrics 34 (5), 836–842.

Buncic, Daniel, 2021. Econometric Issues with Laubach and Williams Estimates of the Natural Rate of Interest. Sveriges Riksbank Working Paper No. 397, Sveriges Riksbank, Available from: https://www.riksbank.se/globalassets/media/rapporter/working-papers/2019/no.-397-econometric-issues-with-laubach-and-williams-estimates-of-the-natural-rate-of-interest2.pdf.

Buncic, Daniel, 2022. On a standard Method for Measuring the Natural Rate of Interest. SSRN Working Paper No. 3725151, Available from: http://ssrn.com/.

Buncic, Daniel, Pagan, Adrian, 2022. Discovering Stars: Problems in Recovering Latent Variables from Models. CAMA Working Paper 52/2022, Centre for Applied Macroeconomic Analysis, Australian National University, Available from: https://cama.crawford.anu.edu.au/sites/default/files/publication/cama_crawford_anu_edu_au/2022-09/52_2022_buncic_pagan.pdf.

Buncic, Daniel, Pagan, Adrian, Robinson, Tim, 2023. Recovering Stars in Macroeconomics. CAMA Working Paper 43/2023, Available from: https://cama.crawford.anu.edu.au/sites/default/files/publication/cama_crawford_anu_edu_au/2023-09/43_2023_buncic_pagan_robinson.pdf.

Chahrour, Ryan, Jurado, Kyle, 2022. Recoverability and expectations-driven fluctuations. Rev. Econom. Stud. 89 (1), 214–239.

Clark, Peter K., 1987. The cyclical component of U.S. economic activity. Q. J. Econ. 102 (4), 797–814.

Elliott, Graham, Rothenberg, Thomas J., Stock, James H., 1996. Efficient tests for an autoregressive unit root. Econometrica 64 (4), 813–836.

Fiorentini, Gabriele, Galesi, Alessandro, Pérez-Quirós, Gabriel, Sentana, Enrique, 2018. The Rise and Fall of the Natural Interest Rat. Banco de Espana Working Paper No. 1822, Banco de Espana.

Forni, Mario, Gambetti, Luca, Sala, Luca, 2019. Structural VARs and noninvertible macroeconomic models. J. Appl. Econometrics 34 (2), 221–246.

Hamilton, James D., 1994. Time Series Analysis. Princeton University Press.

Harvey, Andrew C., 1989. Forecasting, Structural Time Series Models and the Kalman Filter. Cambridge University Press, Cambridge.

Hodrick, Robert, Prescott, Edward C., 1997. Post-war U.S. business cycles: A descriptive empirical investigation. J. Money, Credit, Bank. 29 (1), 1–16.

Holston, Kathryn, Laubach, Thomas, Williams, John C., 2017. Measuring the natural rate of interest: International trends and determinants. J. Int. Econ. 108 (Supplement 1), S59–S75.

Holston, Kathryn, Laubach, Thomas, Williams, John C., 2023. Measuring the Natural Rate of Interest after COVID-19. FRB of New York Staff Report (1063).

Laubach, Thomas, Williams, John C., 2003. Measuring the natural rate of interest. Rev. Econ. Stat. 85 (4), 1063–1070.

Lewis, Kurt F., Vazquez-Grande, Francisco, 2018. Measuring the natural rate of interest: A note on transitory shocks. J. Appl. Econometrics 34 (3), 425–436.

Martínez-García, Enrique, 2021. Get the lowdown: The international side of the fall in the U.S. natural rate of interest. Econ. Model. 100, 105486.

Pagan, Adrian, Robinson, Tim, 2022. Excess shocks can limit the economic interpretation. Eur. Econ. Rev. 145, 104120.

Pagan, Adrian, Wickens, Michael, 2019. Checking If the Straitjacket Fits. CEPR Discussion Paper No.: DP14140.

Quandt, Richard E., 1960. Tests of the hypothesis that a linear regression system obeys two separate regimes. J. Amer. Statist. Assoc. 55 (290), 324–330.

Shephard, Neil G., Harvey, Andrew C., 1990. On the probability of estimating a deterministic component in the local level model. J. Time Series Anal. 11 (4), 339–347.

Stock, James H., Watson, Mark W., 1988. Variable trends in economic time series. J. Econ. Perspect. 2 (3), 147–174.

Stock, James H., Watson, Mark W., 1998. Median unbiased estimation of coefficient variance in a time-varying parameter model. J. Amer. Statist. Assoc. 93 (441), 349–358.