Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

A Simple Correction for Misspecification in Trend-Cycle Decompositions with an Application to Estimating *r**

CAMA Working Paper 2/2022
January 2022 (Original Title: *The Decline in r* according to a Robust Multivariate Trend-Cycle Decomposition*)

Revised in March 2023

James Morley
University of Sydney
Centre for Applied Macroeconomic Analysis, ANU

Trung Duc Tran
University of Sydney
Centre for Applied Macroeconomic Analysis, ANU

Benjamin Wong
Monash University
Centre for Applied Macroeconomic Analysis, ANU

Abstract

We propose a simple correction for misspecification in trend-cycle decompositions when the stochastic trend is assumed to be a random walk process but its estimated path displays serial correlation in its first differences. Possible sources of this misspecification, which might otherwise be hard to detect, include unaccounted for measurement error in the data, omitted variables, or incorrect assumptions about dynamics in the original model used to estimate trend. Our proposed correction is conducted via application of a univariate Beveridge-Nelson decomposition to the preliminary estimated trend and we show in Monte Carlo analysis that this approach can work as well as if the model used to estimate trend were correctly specified. We demonstrate the empirical relevance of our proposed correction in an application to estimating the trend path of the short-term risk-free real interest rate, *r** (a.k.a. "*r*-star*"). Our corrected estimate of *r** is considerably smoother than the preliminary estimate from a multivariate Beveridge-Nelson decomposition based on a vector error correction model, consistent with the presence of measurement error in at least some of the variables in the model.

Keywords

Trend-cycle decomposition, misspecification, natural rate of interest

JEL Classification

C13, C53, E43

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

A simple correction for misspecification in trend-cycle decompositions with an application to estimating r^*

James Morley* Trung Duc Tran† Benjamin Wong‡

March 2023§

Abstract

We propose a simple correction for misspecification in trend-cycle decompositions when the stochastic trend is assumed to be a random walk process but its estimated path displays serial correlation in its first differences. Possible sources of this misspecification, which might otherwise be hard to detect, include unaccounted for measurement error in the data, omitted variables, or incorrect assumptions about dynamics in the original model used to estimate trend. Our proposed correction is conducted via application of a univariate Beveridge-Nelson decomposition to the preliminary estimated trend and we show in Monte Carlo analysis that this approach can work as well as if the model used to estimate trend were correctly specified. We demonstrate the empirical relevance of our proposed correction in an application to estimating the trend path of the short-term risk-free real interest rate, r^* (a.k.a. *"r-star"*). Our corrected estimate of r^* is considerably smoother than the preliminary estimate from a multivariate Beveridge-Nelson decomposition based on a vector error correction model, consistent with the presence of measurement error in at least some of the variables in the model.

JEL codes: C13, C53, E43

Keywords: Trend-cycle decomposition, misspecification, natural rate of interest

*University of Sydney and CAMA. Email: james.morley@sydney.edu.au

†University of Sydney and CAMA. Email: trungduc.tran@sydney.edu.au

‡Monash University and CAMA. Email: benjamin.wong@monash.edu

§This paper previously circulated under the title "The decline in r^* according to a robust multivariate trend-cycle decomposition". We thank Claus Brand, Todd Clark, Jean-Marie Dufour, Sandra Eickmeier, Yunjong Eo, Leland Farmer, Greg Kaplan, Mengheng Li, Elmar Mertens, Tatsuyoshi Okimoto, Tatevik Sekhposyan, Simon van Norden, Nick Sanders, Jake Wong, and conference and seminar participants at the Applied Time Series Econometrics workshop hosted by the Federal Reserve Bank of St. Louis (virtual), the IAAE 2021 Annual Conference (virtual), the Macroeconometrics and Financial Econometrics workshop hosted by Hitotsubashi University (virtual), the Sydney Macro Reading Group workshop (Sydney), the Bank of Canada, the Deutsche Bundesbank, the European Central Bank, the BI Norwegian Business School, the Montreal Econometrics Seminar (virtual), the University of Adelaide (virtual), Monash University, and the University of Tasmania (virtual) for helpful comments. This research was supported by ARC DP190100202 and DE200100693. The usual disclaimers apply.

1 Introduction

Estimates of stochastic trends do not always exhibit behavior consistent with an underlying random walk assumption. Such aberrations could result from otherwise hard-to-detect misspecification when estimating a trend. For example, serial correlation in the first differences of an estimated trend could be due to unaccounted for measurement error in the data, omitted variables, or incorrect assumptions about dynamics in the original model used in estimation. To address this misspecification, we propose a simple correction via the application of a univariate Beveridge and Nelson (1981) (BN) decomposition to the preliminary trend estimates. As an illustrative example, we show how and why the correction works based on the law of iterated expectations in the presence of even a small amount of unaccounted for measurement error. Our proposed correction is easy to apply as a robustness procedure for any trend-cycle decomposition that assumes a random walk trend and Monte Carlo analysis shows it can work as well as if the original model used to estimate trend were correctly specified.

We demonstrate the empirical relevance of our proposed correction in an application to estimating the trend path of the short-term risk-free real interest rate. This application builds on a large literature that links lower interest rates in recent decades to an underlying decline in the low-frequency "equilibrium" or "natural" real rate of interest, often referred to as r^* (or *"r-star"*); see, for example, Cúrdia et al. (2015); Lubik and Matthes (2015); Hamilton et al. (2016); Del Negro et al. (2017); Holston et al. (2017); Fiorentini et al. (2018); Brand and Mazelis (2019); Brand et al. (2019); Berger and Kempa (2019); Lewis and Vazquez-Grande (2019); Bauer and Rudebusch (2020); Kiley (2020a); Johannsen and Mertens (2021); Fu (2023). Notwithstanding some differences regarding the formal definition of r^*, the main debate in the literature has been as to why and how much it has fallen.[1] One prominent view, inspired by the semi-structural unobserved components (UC) model in Laubach and Williams (2003), is that the decline is primarily a supply-side phenomenon associated with lower trend output growth. A competing explanation considers a persistent fall in interest rates due to insufficient aggregate demand, as argued by Summers (2015). Financial market portfolio considerations related to an increase in demand for safe assets have also been put forward by Caballero et al. (2017) and

[1]As discussed in Lunsford and West (2019), there are many definitions of r^*. However, most theoretical settings imply an equivalence between movements in the common low-frequency level of real interest rates (when a set of interest rates includes an essentially risk-free measure) and the low-frequency level of any theoretically-defined r^*. Thus, we focus on trend movements to understand persistent changes in r^* over recent decades.

Del Negro et al. (2017). A basic quantitative question regardless of the hypothesized source is whether r^*, which was previously thought to be approximately 2% over very long periods of time (see, for example, Taylor, 1993), has become negative in the past decade or so as both *ex ante* and *ex post* measures of short-term real interest rates have often been persistently below zero during this time.

For our analysis of r^*, we define it as the common stochastic trend for any set of real interest rates that includes a risk-free short-term rate. Then, to investigate why and how much r^* has fallen, we consider a multivariate version of the BN decomposition based on Morley and Wong (2020). This approach allows us to consider a large set of variables that have been hypothesized to explain changes in r^* and to account for historical movements in r^* based on these variables. Unlike many UC models, the BN decomposition does not assume changes in r^* are orthogonal to structural shocks that also have cyclical implications, which could lead UC models to imply artificially smooth estimates of r^*, especially if the models include many variables linked to cyclical movements in interest rates such as, for example, in Zaman (2023). Our model is a vector error correction model (VECM) that assumes cointegration between short- and long-term real interest rates. Working with *ex ante* real interest rates based on U.S. data makes it viable to model a large system of variables in a linear environment despite the effects of the zero lower bound (ZLB) on nominal interest rates in recent years. However, the constructed *ex ante* real interest rates can be subject to measurement error given the need to proxy for inflation expectations in the bond market, which are generally not directly observed.

Our main empirical findings can be summarized as follows: The preliminary estimated trend for the cointegrated real interest rates is highly informed by the long-term real interest rate, consistent with the expectations hypothesis. However, this estimated trend displays serial correlation in its first differences, suggesting some misspecification in the VECM even though it includes many variables and lags. Then, when we apply our proposed correction, the corrected estimate is considerably smoother than the preliminary estimate, although it still displays a substantial decline over our full sample period from 1973-2019. Notably, the smooth path for r^* is found without imposing smoothness *a priori*, as is often done in the literature (see, for example, Del Negro et al., 2017). Key movements in the estimated r^* can be accounted for by mixture of productivity/demographic and safe asset supply/demand variables.

The rest of this paper is organized as follows. Section 2 proposes our correction for misspecification when there is apparent serial correlation in the first differences of a preliminary estimated trend. Section 3 explains how and why our proposed correction works within the setting of unaccounted for measurement error in the data used to estimate a trend. Section 4 presents the empirical application to the estimation of r^*. Section 5 concludes.

2 A simple correction for misspecification

A standard assumption when conducting trend-cycle decomposition is that stochastic trends follow random walk processes. For example, this assumption is taken in much of the empirical literature on estimating the natural rate of interest, r^*, using semi-structural or reduced-form time series models (see, for example, Laubach and Williams, 2003).[2] However, it sometimes happens that empirical estimates of trend, which we denote as $\{\hat{r}_t^*\}_{t=1}^T$ in the case of estimating r^*, will display some degree of serial correlation in first differences. When this happens, it could be due to misspecification when applying the trend-cycle decomposition. For example, as we show in the next section, failure to account for even a small degree of measurement error in variables used for a multivariate trend-cycle decomposition could result in such serial correlation.

To the extent that a source of misspecification can be detected and addressed by changing the approach to trend-cycle decomposition, including adjusting the model specification, then this would be the obvious way to proceed after finding any preliminary evidence of serial correlation. However, there are settings where the source of misspecification is not clear or cannot be addressed by a straightforward modification to a particular model used when conducting trend-cycle decomposition. In such situations, we propose a simple correction based on applying a univariate BN decomposition directly to the preliminary estimated trend. In particular, as long as the distortion in the preliminary estimated trend due to misspecification is stationary, our proposed correction should provide a more accurate estimated trend that does not display serial correlation in its first differences. As we show in our example in the next section,

[2] In a more structural setting, there may be a model-implied short-run natural rate of interest that is itself subject to transitory dynamics. However, this short-run natural rate should converge to a long-run level that is robust to structural assumptions used to identify transitory movements in the natural rate of interest as distinct from other transitory movements in a short-term risk-free real interest rate. This point about a robust long-run level is related to the motivation for the use of the BN decomposition to evaluate structural models in Rotemberg and Woodford (1996) and Kiley (2013).

this approach works on the basis of the law of iterated expectations, with the univariate BN decomposition removing an estimate of any remaining stationary component in the preliminary estimated trend that was responsible for the apparent serial correlation in its first differences.

To implement our correction, we consider an autoregressive moving-average (ARMA) model for the first differences of the preliminary estimated trend:

$$\phi(L)\Delta\hat{r}_t^* = \theta(L)\epsilon_t, \tag{1}$$

where $\phi(L)$ is an AR lag polynomial with roots outside the unit circle to satisfy stationarity, $\theta(L)$ is an MA lag polynomial with roots outside the unit circle to satisfy invertibility, and ϵ_t is a serially-uncorrelated error term. Then, assuming direct observation of the ARMA errors and based on a univariate BN decomposition, the first difference of the corrected estimated trend, which we denote $\{\tilde{r}_t^*\}_{t=1}^T$, is simply

$$\Delta\tilde{r}_t^* \equiv \frac{\theta(1)}{\phi(1)}\epsilon_t, \tag{2}$$

where assuming the deviations from trend average to zero over the sample, we can estimate the level of r_t^* as

$$\tilde{r}_t^* = \tilde{r}_0^* + \sum_{\tau=1}^{t}\Delta\tilde{r}_\tau^*, \text{where } \tilde{r}_0^* = \frac{1}{T}\sum_{t=1}^{T}\hat{r}_t^* - \frac{1}{T}\sum_{t=1}^{T}\sum_{\tau=1}^{t}\Delta\tilde{r}_\tau^*. \tag{3}$$

Or, as an alternative approach to estimating the level of r_t^*, we could cast the ARMA model into state-space form and follow the calculation in Morley (2002) using the Kalman filter, thus addressing a lack of direct observation of the errors if one were unwilling to make an assumption about their initial values, such as assuming $\epsilon_0 = 0$ for an MA(1) model.

Our proposed correction can be related to the literature on estimating the effects of long-run structural shocks using structural vector autoregressions (SVARs) and non-parametric spectral estimators to address misspecification biases (see, for example, Christiano et al., 2007; Mertens, 2012). However, our focus is specifically on estimating the overall trend rather than identifying structural shocks and our proposed correction is strictly parametric. Related to the approach in Christiano et al. (2007), one could alternatively consider a non-parametric estimator of the spectral density of $\Delta\hat{r}_t$ at frequency zero as a correction, instead of applying a univariate BN decomposition based on a parametric ARMA model. However, as Mertens (2012) points out, such a non-parametric approach provides no panacea against small sample issues when estimating the effects of long-run shocks. Thus, we focus on a parametric approach, in part to maintain consistency with our application of a parametric multivariate BN decomposition to estimate r^* in the first place, although we highlight that the correction can be applied given any preliminary approach that assumes a random-walk stochastic trend.

Another key point with our proposed correction is that as long as we can approximate the serial correlation in the first difference of the preliminary estimated trend with an ARMA model, we can improve inference about a trend in terms of matching its underlying behavior. Meanwhile, if there were actually very little misspecification to start with, then the first difference of the preliminary estimated trend should be close to serially uncorrelated and the corrected estimated trend would be quite similar to the preliminary estimate.[3] Conversely, if the corrected estimates continue to display serial correlation given a poorly-specified ARMA model used for the correction, then additional iterations of our proposed correction could be considered until the estimated errors for the ARMA model appear serially uncorrelated. It is also worth noting that it will generally be more feasible to consider time-varying parameter versions of univariate ARMA models for $\Delta \hat{r}_t^*$ than for high-dimensional multivariate models used in the preliminary estimation of trend if there is any evidence of structural change in the ARMA parameters.[4]

3 An example with measurement error

To illustrate how and why our proposed correction works, we provide a specific example of misspecification due to the presence of unaccounted for measurement error when also applying a BN decomposition to obtain the preliminary estimated trend. However, we emphasize that the correction should also work for other possible sources of misspecification such as omitted variables or incorrect dynamics and is applicable given a preliminary estimated trend from other trend-cycle decompositions such as those based on UC models.

[3]For example, trend growth in output based on the multivariate BN decomposition used to estimate the output gap in Morley and Wong (2020) with log real GDP as the target variable in a large Bayesian VAR displays very little serial correlation. Therefore, applying the correction would have no material impact on the estimates of trend and cycle.

[4]Time variation in ARMA parameters could result from misspecification of a time-varying parameter process with a constant parameter model when conducting the preliminary trend-cycle decomposition. Or it could even occur given a correct specification of a time-varying parameter model, but failure to account for measurement error. In our application, however, we find that the estimated errors from our MA(8) model for $\Delta \hat{r}_t^*$ appear to be serially uncorrelated in both the first and second halves of the sample period from 1973-2019, suggesting that a stable constant parameter MA(8) model is sufficient to capture serial correlation in the preliminary estimated trend based on a constant parameter VECM for the full sample period.

3.1 Beveridge-Nelson decomposition

First, note that the BN decomposition estimates the random walk trend component r_t^* of a time series r_t under the assumption that cyclical deviations from trend have an unconditional mean of zero (i.e., $\mathbb{E}[r_t^c] = 0$, where $r_t^c \equiv r_t - r_t^*$) based on the following equivalence:

$$\mathbb{E}[r_t^* | \Omega_t] = \lim_{h \to \infty} \mathbb{E}[r_{t+h} | \Omega_t], \tag{4}$$

where $\Omega_t = \{\mathbf{x}_t, \ldots, \mathbf{x}_1; \mathbf{f}(\{\mathbf{x}_t\}_{-\infty}^{+\infty})\}$ includes all relevant information at time t for calculating the long-horizon expectation given an assumed data generating process $\mathbf{f}(\{\mathbf{x}_t\}_{-\infty}^{+\infty})$, with $\mathbf{x}_t$ denoting an $n \times 1$ vector that includes the target variable r_t, which is assumed to be in the first row for convenience.[5] The logic of the BN decomposition is that, because the long-horizon expectation of the cyclical deviation from trend is zero, the long-horizon expectation of the overall time series will only reflect an expectation of its trend component. Therefore, in principle, one only needs to specify a forecasting model for a time series to estimate its trend based on the implied long-horizon conditional expectation.

As in Morley and Wong (2020), we assume that conditional expectations for the first difference of the target variable r_t can be fully captured by a stationary linear forecasting model of $\mathbf{\Delta x}_t \sim I(0)$ with the following companion-form representation:

$$\mathbf{\Delta X}_t = \mathbf{F}\mathbf{\Delta X}_{t-1} + \mathbf{H}\mathbf{e}_t, \tag{5}$$

where $\mathbf{\Delta X}_t$ is a $k \times 1$ vector of stationary demeaned variables with $\mathbf{\Delta x}_t - \boldsymbol{\mu}$ in the first n rows with $\boldsymbol{\mu}$ being a vector of unconditional means, $\mathbf{F}$ is a $k \times k$ companion matrix with eigenvalues strictly less than one in modulus, $\mathbf{e}_t$ is a $n \times 1$ vector of serially uncorrelated forecast errors for the variables in $\mathbf{\Delta x}_t$, and $\mathbf{H}$ is a $k \times n$ matrix mapping forecast errors to the companion form with $n \leq k$. We note that this companion form can capture many forecasting models, including VARs and VECMs (see Morley, 2002). We also note that such multivariate forecasting models provide reduced-form representations of many dynamic structural models, with the BN decomposition producing robust estimates of random walk trends for time series across different structural identifications that lead to the same reduced-form representation (see Kiley, 2013).

Following Morley (2002) and Morley and Wong (2020), the BN trend and cycle for r_t given

[5] For this example, we consider trend as a driftless random walk component, as this is typical assumption when estimating r^*. However, it is straightforward to modify the BN decomposition to allow for deterministic drift, as is generally done when applying it to real GDP or other time series with drift.

the forecasting model for $\mathbf{\Delta x}_t$ can be calculated as

$$\mathbb{E}\left[r_t^* \middle| \Omega_t\right] = r_t + \mathbf{s}_{k,1}'\mathbf{F}(\mathbf{I}-\mathbf{F})^{-1}\mathbf{\Delta X}_t, \tag{6}$$

$$\mathbb{E}\left[r_t^c \middle| \Omega_t\right] = -\mathbf{s}_{k,1}'\mathbf{F}(\mathbf{I}-\mathbf{F})^{-1}\mathbf{\Delta X}_t, \tag{7}$$

where we let $\mathbf{s}_{k,j}$ denote a $k \times 1$ selection vector that contains zeros in all rows except for a one in the j^{th} row. Note that, given the forecasting model in (5), the relevant information set can be simplified to $\Omega_t = \{r_t, \mathbf{\Delta X}_t; \mathbf{F}\}$.

A key feature of the BN trend is that its first differences should inherit the same lack of serial correlation as the changes in the true random walk trend. In particular, following Morley and Wong (2020), the first difference of the BN trend is

$$\mathbb{E}\left[r_t^* \middle| \Omega_t\right] - \mathbb{E}\left[r_{t-1}^* \middle| \Omega_{t-1}\right] = \mathbf{s}_{k,1}'(\mathbf{I}-\mathbf{F})^{-1}\mathbf{H}\mathbf{e}_t = \sum_{i=1}^{n} \omega_i e_{it}, \tag{8}$$

where we let ω_i denote a weight that is equal to the i^{th} element of the $1 \times n$ row vector $\boldsymbol{\omega} \equiv \mathbf{s}_{k,1}'(\mathbf{I}-\mathbf{F})^{-1}\mathbf{H}$ and the resulting linear combination of serially-uncorrelated forecast errors in (8) will itself be serially uncorrelated. Thus, any serial correlation in the change in the estimated trend based on the BN decomposition must reflect some misspecification such as unaccounted for measurement error.

3.2 Measurement error and our proposed correction

To consider the effects of measurement error, let $\check{\mathbf{x}}_t \equiv \mathbf{x}_t + \mathbf{u}_t$ denote a vector of observed variables in their accumulated levels with stationary and unconditionally mean-zero measurement error $\mathbf{u}_t \sim I(0)$ and $\mathbb{E}[\mathbf{u}_t] = \mathbf{0}$.[6] For example, it is quite likely that our measure of the real interest rate contains some measurement error as we need to proxy for inflation expectations when constructing an *ex ante* real interest rate. Thus, in practice, we only observe $\check{r}_t$ as a conflation of the underlying true real interest rate r_t and measurement error u_{1t}. As we

[6] The assumption of measurement error in levels, not first differences, is crucial for our correction to work, but we believe it is a reasonable assumption for many macroeconomic variables that the measured and true values are not permanently drifting apart. It is also worth noting that our approach allows the measurement error to be correlated with the underlying structural shocks. The key distinction that defines $\mathbf{u}_t$ as measurement error and distinct from forecast errors reflecting structural shocks is that error-adjusted variables (i.e., $\mathbf{x}_t = \check{\mathbf{x}}_t - \mathbf{u}_t$) are block exogenous with respect to the measurement error. In addition to the standard classical errors-in-variables 'noise' case, this block exogeneity could correspond to the 'news' case in Dungey et al. (2015) where the measurement error is positively correlated with future changes in observed variables instead of negatively correlated. In both cases, the measurement error, if it could be observed, would Granger-cause future changes in the actual observed variables, but not in the error-adjusted variables. Meanwhile, Anderson et al. (2019) show that a high signal-to-noise ratio in the errors-in-variables setting leads to only small distortions in inferences about Granger causality for the observed variables compared to the error-adjusted variables.

will see, this measurement error, if unaccounted for in the forecasting model used for the BN decomposition, would imply serial correlation in the first differences of the BN trend for $\check{r}_t$. However, assuming this implied serial correlation can be well captured by an ARMA model, we also show how to correct for it.

To see how we propose correcting for the effects of unaccounted for measurement error, first consider the implied model for the observed data with measurement error:

$$\mathbf{\Delta}\check{\mathbf{X}}_t = \mathbf{F}\mathbf{\Delta}\check{\mathbf{X}}_{t-1} + \mathbf{H}\check{\mathbf{e}}_t, \tag{9}$$

where $\mathbf{\Delta}\check{\mathbf{X}}_t \equiv \mathbf{\Delta}\mathbf{X}_t + \mathbf{H}\mathbf{\Delta}\mathbf{u}_t$ and $\check{\mathbf{e}}_t \equiv \mathbf{e}_t + \mathbf{\Delta}\mathbf{u}_t - \mathbf{S}'\mathbf{F}\mathbf{H}\mathbf{\Delta}\mathbf{u}_{t-1}$ given the $k \times n$ selection matrix $\mathbf{S} = (\mathbf{s}'_{k,1}, \mathbf{s}'_{k,2}, ..., \mathbf{s}'_{k,n})'$. Note that the $\mathbf{\Delta}\mathbf{u}_t - \mathbf{S}'\mathbf{F}\mathbf{H}\mathbf{\Delta}\mathbf{u}_{t-1}$ term in the vector of residuals $\check{\mathbf{e}}_t$ means these residuals will be serially correlated even if the measurement error in $\mathbf{u}_t$ is itself serially uncorrelated. Thus, estimation of the companion matrix $\mathbf{F}$ without taking into account the serial correlation in the residuals will correspond to a model misspecification and standard estimates given this misspecification will converge to the linear projection matrix $\mathbf{P}$ from the following linear projection:

$$\mathbf{\Delta}\check{\mathbf{X}}_t = \mathbf{P}\mathbf{\Delta}\check{\mathbf{X}}_{t-1} + \mathbf{H}\boldsymbol{\eta}_t, \tag{10}$$

where $\boldsymbol{\eta}_t$ corresponds to serially-correlated projection errors such that $\mathbf{H}\boldsymbol{\eta}_t = (\mathbf{I} - \mathbf{P}L)(\mathbf{I} - \mathbf{F}L)^{-1}\mathbf{H}\check{\mathbf{e}}_t$.[7]

Next, note the following decomposition based on identities for the real interest rate measured with error:

$$\begin{aligned} \check{r}_t &= r_t + u_{1t}, \\ &= \mathbb{E}\left[r_t^* | \, \Omega_t\right] + \mathbb{E}\left[r_t^c | \, \Omega_t\right] + u_{1t}, \\ &= \mathbb{E}\left[r_t^* | \, \Omega_t\right] + \mathbb{E}\left[r_t^c | \, \Omega_t\right] + \hat{r}_t^c - \hat{r}_t^c + u_{1t}, \end{aligned}$$

where $\hat{r}_t^c \equiv -\mathbf{s}'_{k,1}\mathbf{P}(\mathbf{I} - \mathbf{P})^{-1}\mathbf{\Delta}\check{\mathbf{X}}_t$ is the BN cycle for $\check{r}_t$ given a misspecified forecasting model for $\mathbf{\Delta}\check{\mathbf{x}}_t$ that does not take into account serial correlation in the residuals. Then, we get the

[7] The difference between the linear projection matrix $\mathbf{P}$ and the companion matrix $\mathbf{F}$ is analogous to how the population linear projection for a variable following an ARMA(1,1) process on its lag (i.e., the first-order autocorrelation of the process) depends on both the autoregressive and moving-average parameters and will be different from the autoregressive parameter given a non-zero value of the moving-average parameter. Specifically, the first-order autocorrelation for an AR(1) process is simply the autoregressive parameter ϕ, while for an ARMA(1,1) process, it is $\phi + \theta(1 - \phi^2)/(1 + 2\phi\theta + \theta^2)$, where θ is the moving-average parameter on the lagged shock to the ARMA(1,1) process.

following expression for the preliminary BN trend, $\hat{r}_t^* \equiv \check{r}_t - \hat{r}_t^c$:

$$\hat{r}_t^* = \mathbb{E}\left[r_t^* | \, \Omega_t\right] + \mathbb{E}\left[r_t^c | \, \Omega_t\right] - \hat{r}_t^c + u_{1t}. \tag{11}$$

Note that both estimates of the cyclical component on the right hand side of the expression in (11) (i.e., hypothetical $\mathbb{E}\left[r_t^c | \, \Omega_t\right]$ if data were observed without measurement error and actual $\hat{r}_t^c$ given observed data with measurement error) will be stationary and mean zero by construction and the measurement error u_{1t} is $I(0)$ and mean zero by assumption. Thus, the random walk component of the preliminary BN trend $\hat{r}_t^*$ in (11) corresponds to the BN trend $\mathbb{E}\left[r_t^* | \, \Omega_t\right]$ in (4) (i.e., the BN trend if the data were observed without measurement error). As a result, after applying the preliminary BN decomposition to $\check{r}_t$, we can subsequently apply a univariate BN decomposition based on an ARMA model directly to $\hat{r}_t^*$ in order to correct the preliminary estimate for serial correlation in its first differences due to the stationary terms on the right hand side of (11). Specifically, based on the law of iterated expectations, we get the following equivalence:

$$\mathbb{E}\left[\mathbb{E}\left[r_t^* | \Omega_t\right] | \check{\Omega}_t\right] = \mathbb{E}\left[r_t^* | \check{\Omega}_t\right] = \lim_{h \to \infty} \mathbb{E}\left[\hat{r}_{t+h}^* | \check{\Omega}_t\right], \tag{12}$$

where $\check{\Omega}_t = \{\check{\mathbf{x}}_t, \ldots, \check{\mathbf{x}}_1; \mathbf{P}, \boldsymbol{\mu}, \phi(L), \theta(L)\}$, with $\phi(L)$ and $\theta(L)$ corresponding to the lag polynomials from an ARMA model for $\Delta \hat{r}_t^*$.[8] As can be seen, the BN trend for $\hat{r}_t^*$ (i.e., $\lim_{h \to \infty} \mathbb{E}\left[\hat{r}_{t+h}^* | \check{\Omega}_t\right]$) provides an estimate of trend for r_t (i.e., $\mathbb{E}\left[r_t^* | \check{\Omega}_t\right]$) even though we do not directly observe r_t due to measurement error.

We note that the presence of measurement error will generally imply a reasonably complicated ARMA structure for the serial correlation in $\Delta \hat{r}_t^*$. For example, consider a VAR(p) structure for $\mathbf{\Delta x}_t$ and assume that at least some of the variables including the real interest rate are measured with MA(q) error – i.e., $u_{1t} \sim$ MA(q) and $u_{jt} \sim$ MA($\leq q$) for $j > 1$. Then, $\mathbf{\Delta \check{x}}_t$ will have a VARMA($p, q+2$) structure and, following Corollary 11.1.1 in Lütkepohl (2005), we can solve for ARMA orders of the cyclical terms in $\hat{r}_t^*$:

$$\mathbb{E}\left[r_t^c | \, \Omega_t\right] = -\mathbf{s}_{k,1}' \mathbf{F}(\mathbf{I} - \mathbf{F})^{-1} \mathbf{\Delta X}_t \sim \text{ARMA}(\leq pn, \leq pn - 1), \tag{13}$$

$$\hat{r}_t^c = -\mathbf{s}_{k,1}' \mathbf{P}(\mathbf{I} - \mathbf{P})^{-1} \mathbf{\Delta \check{X}}_t \sim \text{ARMA}(\leq pn, \leq pn + q + 1), \tag{14}$$

[8] Instead of an ARMA model, one could alternatively consider a UC model for $\hat{r}_t^*$. However, given an implied correlation between permanent and transitory movements in $\hat{r}_t^*$ based on the general properties of a BN decomposition given by the first two terms on the right hand side of the expression in (11), applying a BN decomposition will be appropriate even when this correlation is not identified for a UC model. See Morley et al. (2003) on the identification of the correlation for UC models.

where the orders follow from an ability to re-write the terms in (13) and (14) in terms of the VAR(p) as $\mathbf{s}'_{n,1}\mathbf{c}_t$ and $\mathbf{s}'_{n,1}\check{\mathbf{c}}_t$, with $\mathbf{c}_t \equiv -\boldsymbol{\Gamma}_{\mathbf{F},1}(\boldsymbol{\Delta}\mathbf{x}_t - \boldsymbol{\mu}) - \ldots - \boldsymbol{\Gamma}_{\mathbf{F},p}(\boldsymbol{\Delta}\mathbf{x}_{t-p+1} - \boldsymbol{\mu})$, $\check{\mathbf{c}}_t \equiv -\check{\boldsymbol{\Gamma}}_{\mathbf{P},1}\boldsymbol{\Delta}(\check{\mathbf{x}}_t - \boldsymbol{\mu}) - \ldots - \check{\boldsymbol{\Gamma}}_{\mathbf{P},p}(\boldsymbol{\Delta}\check{\mathbf{x}}_{t-p+1} - \boldsymbol{\mu})$, and $\boldsymbol{\Gamma}_{\mathbf{F},j}$ and $\boldsymbol{\Gamma}_{\mathbf{P},j}$ denoting the j^{th} set of n columns of $\mathbf{H}'\mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}$ and $\mathbf{H}'\mathbf{P}(\mathbf{I} - \mathbf{P})^{-1}$, respectively. In particular, given $\boldsymbol{\Delta}\mathbf{x}_t \sim$ VAR(p) and $\boldsymbol{\Delta}\check{\mathbf{x}}_t \sim$ VARMA($p, q+2$), $\mathbf{c}_t$ will be VARMA($p, p-1$) and $\check{\mathbf{c}}_t$ will be VARMA($p, p+q+1$), with the ARMA orders for $\mathbf{s}'_{n,1}\mathbf{c}_t$ and $\mathbf{s}'_{n,1}\check{\mathbf{c}}_t$ given by Corollary 11.1.1 in Lütkepohl (2005). Although the ARMA orders provide upper bounds given that roots for the implied autoregressive and moving-average polynomials may cancel for some paramaterizations, the point is that these are highly complicated processes and imply the following even more complicated process for $\Delta\hat{r}^*_t$ when combined with the other terms that show up in $\hat{r}^*_t$ in (11) and first differences are taken:

$$\Delta\hat{r}^*_t \sim \text{ARMA}(\leq 2pn, \leq 2pn + q + 2). \tag{15}$$

In practice, we tend not to know the process for measurement error and it may be hard to detect it using tests for serial correlation in specific elements of the projection error vector $\boldsymbol{\eta}_t$ if the measurement error is relatively small compared to the forecast errors in $\mathbf{e}_t$. To see this issue, consider a case where the measurement error is small enough (i.e., $var(\boldsymbol{\Delta}\mathbf{u}_t) << var(\boldsymbol{\Delta}\mathbf{x}_t)$) such that $\mathbf{P} \approx \mathbf{F}$. In this case, we get the following simplified expression for $\Delta\hat{r}^*_t$:

$$\begin{aligned} \Delta\hat{r}^*_t &\approx \mathbf{s}'_{k,1}(\mathbf{I} - \mathbf{F})^{-1}\mathbf{H}\mathbf{e}_t + \mathbf{s}'_{k,1}\mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}\mathbf{H}\boldsymbol{\Delta}\mathbf{u}_t - \mathbf{s}'_{k,1}\mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}\mathbf{H}\boldsymbol{\Delta}\mathbf{u}_{t-1} + \Delta u_{1t} \\ &= \sum_{i=1}^{n}\omega_i e_{it} + \sum_{i=1}^{n}\gamma_i(\Delta u_{it} - \Delta u_{it-1}) + \Delta u_{1t}, \end{aligned} \tag{16}$$

where we let γ_i denote a weight that is equal to the i^{th} element of the $1 \times n$ row vector $\boldsymbol{\gamma} \equiv \mathbf{s}'_{k,1}\mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}\mathbf{H}$. This expression suggests that $\Delta\hat{r}^*_t$ will have an MA($q + 2$) structure given MA(q) measurement error and the effects of measurement error may be easier to detect than in $\Delta\check{r}_t = \Delta r_t + \Delta u_{1t}$ given a very different signal-to-noise ratio in terms of the effects of shocks (signal) versus measurement error (noise).[9] In particular, if r^*_t is smoother than r_t (i.e., $var(\Delta r^*_t) < var(\Delta r_t)$), the equivalence between the variance of trend shocks and the variance of changes in the BN trend (see Morley, 2011; Kamber et al., 2018) implies that the signal in $\Delta\hat{r}^*_t$ (i.e., $\sum_{i=1}^{n}\omega_i e_{it}$) is weaker than the signal in $\Delta\check{r}_t$ (i.e., Δr_t). Furthermore, the

[9] In our application, we find evidence for MA dynamics in $\Delta\hat{r}^*_t$ instead of a more general ARMA structure such as implied by (15). However, we note that a lack of evidence for AR dynamics does not, on its own, imply the measurement error is particularly small, as there could be near cancellation of some AR and MA roots for some paramaterizations and measurement error processes.

$\sum_{i=1}^{n} \gamma_i(\Delta u_{it} - \Delta u_{it-1})$ term in (16) suggests additional noise related to the measurement error, strictly so given independent measurement error across variables and a positive relationship between the forecast error for r_t and its long-horizon forecast such that $\gamma_1 > 0$. Thus, the lower signal-to-noise ratio for the change in the preliminary estimated trend $\Delta \hat{r}_t^*$ than the data $\Delta \check{r}_t$ means that serial correlation could be considerably easier to detect in $\Delta \hat{r}_t^*$ than in the target-variable projection error η_{1t}, which is what we find in our empirical analysis. We also note that classical measurement error for the target variable with variance σ_{u1}^2 will tend to imply some negative serial correlation in $\Delta \hat{r}_t^*$, especially if the γ_i's are small, with $cov(\Delta \hat{r}_t^*, \Delta \hat{r}_{t-1}^*) = -\sigma_{u1}^2$ when $\gamma_i = 0$, $\forall i$. Indeed, negative serial correlation is what we find in our empirical application and it leads our corrected estimates to be smoother than the preliminary estimates. However, the point is that the smoothing reflects the source of misspecification such as unaccounted for classical measurement error, and is not an inherent feature of the correction.

In terms of applying our proposed correction in practice, it is straightforward to show based on (1), (2), (8), and (10) that the first difference of the corrected BN trend can be written as follows:

$$\Delta \tilde{r}_t^* = \frac{\theta(1)\phi(L)}{\phi(1)\theta(L)} \mathbf{s}_{k,1}' (\mathbf{I} - \mathbf{P})^{-1} (\mathbf{\Delta} \check{\mathbf{X}}_t - \mathbf{P} \mathbf{\Delta} \check{\mathbf{X}}_{t-1}), \tag{17}$$

and the level for $\tilde{r}_t^*$ can be calculated following (3). We also note that there remains an open question about what is the best estimate of the cyclical component of r_t. If we use the corrected BN cycle $\tilde{r}_t^c \equiv \check{r}_t - \tilde{r}_t^*$, it will clearly include some of the measurement error in $\check{r}_t$ in addition to the true cyclical component, while the preliminary BN cycle $\hat{r}_t^c$ may be closer to $\mathbb{E}\left[r_t^c | \Omega_t\right]$. To the extent that the measurement error appears to be small given an apparent lack of serial correlation in the projection error η_{1t}, we would suggest using the corrected BN cycle $\tilde{r}_t^c$.

Given population parameters, applying our correction would not produce as precise of an estimate of r_t^* as if we knew Ω_t (i.e., we observed the data without measurement error) or even as an estimate based on a correctly-specified multivariate model for $\mathbf{\Delta} \check{\mathbf{x}}_t$ that captures the serial correlation in $\check{\mathbf{e}}_t$. For example, if the model for $\mathbf{\Delta x}_t$ is a VAR, the correctly-specified forecasting model for the data measured with error $\mathbf{\Delta} \check{\mathbf{x}}_t$ would be a VARMA for a range of assumptions about stationary processes for the measurement error, including even if it is serially uncorrelated. However, our proposed approach is feasible given that we do not actually observe Ω_t and is much more practicable to implement than consideration of a VARMA given challenges in specifying and estimating parameters for such models when there are many variables and lags. Meanwhile, by incorporating multivariate information in our preliminary estimate of trend, our approach should be more precise than a BN decomposition based on a correctly-specified univariate model $\Delta\check{r}_t$ (or, equivalently, Δr_t under the assumption of no measurement error in the target variable). In particular, our proposed approach can be thought of as partitioning the parameter space into two, using multivariate information $\{\check{\mathbf{X}}_t, \check{\mathbf{X}}_{t-1}\}$ to estimate the projection parameters $\mathbf{P}$ and univariate information $\{\hat{r}_t^*, \hat{r}_{t-1}^* \ldots, \hat{r}_1^*\}$ that more strongly reflects measurement error than $\Delta\check{r}_t$ to estimate univariate ARMA parameters for $\Delta\hat{r}_t^*$ implied by the multivariate VARMA process.[10] The tradeoff is less information than a full multivariate approach (but still more than a full univariate approach) in return for generally far fewer parameters to estimate than for the full multivariate model.

3.3 Monte Carlo analysis

To help illustrate the tradeoffs between different approaches to estimating trend and demonstrate how well our proposed correction works, we conduct Monte Carlo analysis. For simplicity, we consider a data generating process (DGP) in which there are implicitly the same number as structural shocks as observable variables. In this setting, we get the "BN-as-definition" scenario in Morley (2011) where $r_t^* = \mathbb{E}\left[r_t^* \middle| \Omega_t\right]$ and $r_t^c = \mathbb{E}\left[r_t^c \middle| \Omega_t\right]$. We consider a bivariate VAR(1) specification for the underlying data process $\mathbf{\Delta x}_t$ with $\boldsymbol{\mu} = \mathbf{0}$. In this case, $\mathbf{\Delta X}_t = \mathbf{\Delta x}_t$, with the following simplified form:

$$\underbrace{\begin{pmatrix} \Delta r_t \\ \Delta x_{2t} \end{pmatrix}}_{\mathbf{\Delta X}_t} = \underbrace{\begin{pmatrix} 0 & -0.05 \\ 0 & 0.95 \end{pmatrix}}_{\mathbf{F}} \underbrace{\begin{pmatrix} \Delta r_{t-1} \\ \Delta x_{2t-1} \end{pmatrix}}_{\mathbf{\Delta X}_{t-1}} + \underbrace{\begin{pmatrix} e_{1t} \\ e_{2t} \end{pmatrix}}_{\mathbf{e}_t}, \qquad \mathbf{e}_t \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \underbrace{\begin{pmatrix} 0.1125 & 0.1 \\ 0.1 & 0.1 \end{pmatrix}}_{\mathbf{\Sigma}} \right).$$

Then, $r_t = r_t^* + r_t^c$, where $r_t^* = r_0^* + \sum_{\tau=1}^{t} \Delta r_\tau^*$, with $r_0^* = 0$ and $\Delta r_t^* = \mathbf{s}_{2,1}'(\mathbf{I} - \mathbf{F})^{-1}\mathbf{e}_t$, and $r_t^c = -\mathbf{s}_{2,1}'\mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}\mathbf{\Delta X}_t$. For the observed data $\check{\mathbf{x}}_t = \mathbf{x}_t + \mathbf{u}_t$, where $\check{r}_t = r_t + u_{1t}$ and $\check{x}_{2t} = x_{2t} + u_{2t}$, we assume the addition of serially-uncorrelated measurement that has only 5% the variance of the shocks– i.e., $\mathbf{u}_t \sim \mathcal{N}(\mathbf{0}, 0.05 \times \mathbf{\Sigma})$. That is, we focus on the case

[10] In a recent study, Dufour and Pelletier (2022) develop some practical methods for specifying and estimating VARMA models, including considering diagonal MA equations. Estimation of the VAR and MA parameters is split into parts, not unlike our proposed approach, although we focus on univariate ARMA estimation for $\Delta\hat{r}_t^*$ instead of MA estimation for univariate projection errors $\check{e}_{it}$, $i = 1, \ldots, n$, from long autoregressions. See Dufour and Pelletier (2022) for full details of how to estimate VARMA models with diagonal MA equations.

where measurement error is small in the sense that the projection errors for the misspecified model will not display much serial correlation on their own. Again following Corollary 11.1.1 in Lütkepohl (2005), we note that this DGP implies the following processes: $\Delta r_t \sim \text{ARMA}(2,1)$, $\boldsymbol{\Delta}\check{\mathbf{x}}_t \sim \text{VARMA}(1,2)$, $\Delta\check{r}_t \sim \text{ARMA}(2,3)$, and $\Delta\hat{r}_t^* \sim \text{ARMA}(4,6)$. However, given the small degree of measurement error and $\mathbf{P} \approx \mathbf{F}$, we get $\Delta\hat{r}_t^* \approx \text{MA}(2)$.

Given a sample size of $T = 200$, we evaluate the accuracy of different estimates of r_t^* by calculating the root-mean-squared-error (RMSE) for $\hat{r}_t^*$ or $\tilde{r}_t^*$ for the following cases:

1. **VAR(1) for $\boldsymbol{\Delta}\mathbf{x}_t$,** where $\hat{r}_t^* = \mathbb{E}[r_t^*|r_t, \boldsymbol{\Delta}\mathbf{X}_t; \hat{\mathbf{F}}]$. This case corresponds to the underlying data without measurement error being observed, with the RMSE only reflecting the effects of estimation uncertainty about the $n^2 = 4$ parameters in $\mathbf{F}$. Parameter estimation is conducted via OLS and the BN trend is calculated based on (6).

2. **VARMA(1,2) for $\boldsymbol{\Delta}\check{\mathbf{x}}_t$,** where $\hat{r}_t^* = \mathbb{E}[r_t^*|\check{\mathbf{x}}_t, \ldots, \check{\mathbf{x}}_1; \hat{\mathbf{F}}, \hat{\boldsymbol{\Theta}}]$, with $\boldsymbol{\Theta}$ corresponding to the MA parameters in the VARMA model. This case allows us to consider the effects of measurement error and additional parameters to estimate, but assuming the correct model specification for $\boldsymbol{\Delta}\check{\mathbf{x}}_t$. The RMSE reflects noisy information about $\{\mathbf{x}_t, \ldots, \mathbf{x}_1\}$ and estimation uncertainty about the $n^2 = 4$ parameters in $\mathbf{F}$ and the $2n^2 = 8$ parameters in $\boldsymbol{\Theta}$. The VARMA model is cast into state-space form and parameter estimation is conducted via exact MLE based on the Kalman filter and the prediction error decomposition of the likelihood. The BN trend is calculated following Morley (2002).

3. **True VARMA(1,2) for $\boldsymbol{\Delta}\check{\mathbf{x}}_t$,** where $\hat{r}_t^* = \mathbb{E}[r_t^*|\check{\mathbf{x}}_t, \ldots, \check{\mathbf{x}}_1; \mathbf{F}, \boldsymbol{\Theta}]$. This case considers the effects of information loss from only observing the data with measurement error, but assuming the correct model specification for $\boldsymbol{\Delta}\check{\mathbf{x}}_t$ and the true population parameter values for the model. Note that the $\boldsymbol{\Theta}$ parameters depend on $\mathbf{F}$ and $\hat{r}_t^*$ can be calculated based on their implicit values by casting the model in (9) into state-space form and using the Kalman filter following Morley (2002).

4. **VAR(1) for $\boldsymbol{\Delta}\check{\mathbf{x}}_t$,** where $\hat{r}_t^* = \mathbb{E}[r_t^*|\check{r}_t, \boldsymbol{\Delta}\check{\mathbf{X}}_t; \hat{\mathbf{P}}]$. This case considers what happens if preliminary trend estimates are used despite model misspecification. That is, $\boldsymbol{\Delta}\check{\mathbf{x}}_t$ is incorrectly assumed to have no measurement error and $\hat{\mathbf{P}}$ is incorrectly used as an estimate of $\mathbf{F}$. As this is completely analogous to Case 1, just with the incorrect assumptions, parameter estimation is conducted via OLS and the BN trend is calculated based on (6),

but substituting $\mathbf{\Delta}\check{\mathbf{X}}_t$ for $\mathbf{\Delta X}_t$ and $\mathbf{P}$ for $\mathbf{F}$.

5. **VAR(1) for $\mathbf{\Delta}\check{\mathbf{x}}_t$ + MA(2) for $\Delta\hat{r}_t^*$,** where $\tilde{r}_t^* = \mathbb{E}[r_t^*|\check{\mathbf{x}}_t, \ldots, \check{\mathbf{x}}_1; \hat{\mathbf{P}}, \hat{\theta}_1, \hat{\theta}_2]$. This case considers our proposed correction using a parsimonious MA(2) model to capture serial correlation in $\Delta\hat{r}_t^*$. $\hat{\mathbf{P}}$ is estimated via OLS and the MA parameters are estimated via conditional MLE.

6. **VAR(1) for $\mathbf{\Delta}\check{\mathbf{x}}_t$ + ARMA(4,6) for $\Delta\hat{r}_t^*$,** where $\tilde{r}_t^* = \mathbb{E}[r_t^*|\check{\mathbf{x}}_t, \ldots, \check{\mathbf{x}}_1; \hat{\mathbf{P}}, \hat{\phi}_1, \ldots, \hat{\phi}_4, \hat{\theta}_1, \ldots, \hat{\theta}_6]$. This case considers our proposed correction using the true (but not necessarily minimal representation) ARMA(4,6) model specification to capture serial correlation in $\Delta\hat{r}_t^*$. $\hat{\mathbf{P}}$ is estimated via OLS and the ARMA parameters are estimated via exact MLE using the Kalman filter.

7. **ARMA(2,1) for Δr_t,** where $\hat{r}_t^* = \mathbb{E}[r_t^*|r_t, r_{t-1}, \ldots, r_1; \hat{\phi}_1, \hat{\phi}_2, \hat{\theta}]$. This case corresponds to only the target variable being observed, but without measurement error. The RMSE reflects the loss of multivariate information and estimation uncertainty about the ARMA(2,1) parameters. The ARMA model is cast into state-space form and parameter estimation is conducted via exact MLE based on the Kalman filter and the prediction error decomposition of the likelihood. The BN trend is calculated following Morley (2002).

8. **ARMA(2,3) for $\Delta\check{r}_t$,** where $\hat{r}_t^* = \mathbb{E}[r_t^*|\check{r}_t, \check{r}_{t-1}, \ldots, \check{r}_1; \hat{\phi}_1, \hat{\phi}_2, \hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3]$. This case consider the effects of measurement error and additional parameters, but assuming the correct univariate model specification for $\Delta\check{r}_t$. The RMSE reflects the effects of measurement error and estimation uncertainty, in addition to the loss of multivariate information. The ARMA model is cast into state-space form and parameter estimation is conducted via exact MLE based on the Kalman filter and the prediction error decomposition of the likelihood. The BN trend is calculated following Morley (2002).

The results for the Monte Carlo analysis are reported in Table 1 and illuminate the roles of parameter uncertainty, measurement error, model specification, and univariate versus multivariate information in determining the precision of trend estimates. We note that the scale of the reported RMSEs can be related to the fact that, in the absence of measurement error, the error in estimating trend is the same size as the error in estimating the cycle, while the implied standard deviation of the cycle is 1.03 for the DGP. So the RMSEs approximately equal fractions of the standard deviation of the cyclical component of the simulated real interest rate.

Table 1: Accuracy of trend estimates in a Monte Carlo experiment

Case	# of parameters to estimate	RMSE
1. VAR(1) for $\mathbf{\Delta x}_t$	4	0.18
2. VARMA(1,2) for $\mathbf{\Delta \check{x}}_t$	12	0.33
3. True VARMA(1,2) for $\mathbf{\Delta \check{x}}_t$	0	0.05
4. VAR(1) for $\mathbf{\Delta \check{x}}_t$	4	0.37
5. VAR(1) for $\mathbf{\Delta \check{x}}_t$ + MA(2) for $\Delta \tilde{r}_t$	6	0.33
6. VAR(1) for $\mathbf{\Delta \check{x}}_t$ + ARMA(4,6) for $\Delta \tilde{r}_t$	14	0.37
7. ARMA(2,1) for Δr_t	3	0.94
8. ARMA(2,3) for $\Delta \check{r}_t$	5	0.97

Notes: Results are for a Monte Carlo experiment based on a bivariate VAR(1) DGP for $\mathbf{\Delta x}_t$ with serially uncorrelated measurement error in $\mathbf{\Delta \check{x}}_t$ where the variance-covariance for $\mathbf{u}_t$ is equal to 0.05 times the conditional variance-covariance of $\mathbf{\Delta x}_t$. The RMSE for an estimate based on a given model(s) is relative to r_t^* based on the true VAR(1) for $\mathbf{\Delta x}_t$. Given the error in estimating trend is equivalent (but opposite sign) to the error in estimating the cycle (in the absence of measurement error), the RMSE can naturally be compared in scale to the implied standard deviation of the true cycle, which is 1.03. The sample size is $T = 200$ and the number of Monte Carlo simulations is 100.

First, even given Case 1 of the correct model and no measurement error, the RMSE of 0.18 suggests that parameter estimation introduces a nontrivial amount of error in the estimates of r_t^* due to the sample size $T = 200$ and the persistence in the DGP leading to bias in OLS estimates of VAR parameters (see, for example, Kilian, 1998). Introducing measurement error, but assuming the correct model in Case 2, almost doubles the RMSE to 0.33 due the noisier information given the measurement error and the increase in parameter uncertainty given 12 VARMA parameters instead of 4 VAR parameters. However, most of the increase in RMSE appears to be due to parameter uncertainty as the RMSE in Case 3 given the true VARMA parameter values is only 0.05 compared to 0.33 for the estimated VARMA parameters, suggesting that noisier information has a smaller effect than parameter uncertainty, which reflects our assumption for the DGP of relatively small measurement error.

Second, despite an important role of parameter uncertainty, model misspecification appears to have a larger effect on the RMSE given that the RMSE for the misspecified VAR in Case 4 of 0.37 is larger than for the correctly-specified VARMA model with more parameters in Case 3 of 0.33. Crucially, however, our proposed correction with the parsimonious MA(2) model for the first differences of the preliminary estimated trend in Case 5 is able to completely address the effects of misspecification, reducing the RMSE to the same level of 0.33 as for the correctly-specified VARMA model. This reduction from the RMSE for the misspecified model occurs even though $\mathbf{P} \approx \mathbf{F}$ given a small amount of measurement error. At the same time, when the much more parameterized ARMA(4,6) model is considered for our proposed correction in Case 6, the RMSE of 0.37 is no better than for the misspecified VAR model. This reflects estimation of even more parameters than in the true VARMA model. However, a near cancellation of roots for the true ARMA model given relatively small measurement error means that a more parsimonious MA model would likely be chosen in practice when conducting our proposed correction.

Third, the results are considerably worse for the univariate estimates, even when there is no measurement error and the true ARMA is tightly parameterized in Case 7 with an RMSE of 0.94. Adding in measurement error and additional ARMA parameters in Case 8 slightly worsens the RMSE to 0.97, but either way, these RMSE given a univariate model is essentially three times worse than the RMSE of 0.33 for the correctly-specified VARMA model and our proposed correction with a parsimonious model for the first differences of the preliminary estimated trend. This dramatic deterioration in the RMSEs makes it clear just how crucial incorporating multivariate information is for producing a relatively accurate estimate of r_t^*.

The fact that our proposed correction can lead to as accurate estimates as if we had the correctly-specified VARMA model is particularly notable given that our bivariate DGP is such that the VARMA model is reasonably tightly parameterized. If we considered $n > 2$ and $p > 1$ instead, the number of VARMA(p,2) parameters to estimate, which is $n^2(p+2)$, would rapidly proliferate. Indeed, even for $n = 3$, we would have 27 parameters for a VARMA(1,2). When considering simulated data in this case, we found that the RMSE exploded even given a much larger sample size of $T = 2000$, while the RMSE for the parsimonious two-step approach with 11 parameters to estimate remained relatively small given a sample size of $T = 200$ and better than for the misspecified VAR model with 9 parameters. Meanwhile, we note that a VARMA model in our empirical application with $n = 15$ and $p = 4$ would have 1,350 parameters, which we do not consider practicable to estimate, at least without much stronger shrinkage or other restrictions than we consider in our estimation.

3.4 Other possible sources of misspecification

Despite the focus on measurement error in our example of misspecification, we highlight that our proposed approach can also handle misspecification due to omitted variables or incorrect dynamics. In particular, if only a subset of variables in $\mathbf{x}_t$ were included in $\check{\mathbf{x}}_t$, it would imply serial correlation in $\check{\mathbf{e}}_t$ even if there were no measurement error (i.e., $\mathbf{u}_t = \mathbf{0}$). It is also possible that the underlying data generating process for $\mathbf{x}_t$ simply involves moving-average dynamics such that the errors for a finite-order VAR will exhibit serial correlation even if there were no measurement error in the observed data.

Consider, for example, a VAR(p) structure for $\mathbf{\Delta x}_t$, but $\check{\mathbf{x}}_t$ only includes the first $m < n$ variables in $\check{\mathbf{x}}_t$. Again, assume no measurement error. Then following Corollary 11.1.2 in Lütkepohl (2005), $\mathbf{\Delta}\check{\mathbf{x}}_t$ will have a VARMA($\leq p(m-n+1), \leq p(m-n)$) structure and the cyclical terms in (11) will have complicated ARMA dynamics, although there will be a slight simplification compared to the measurement error case given that $u_{1t} = 0$. Adding more lags to the VAR might help address omitted MA dynamics and leave little detectable serial correlation in the estimated errors, similar to the case of a small amount of measurement error. The point is that $\Delta \hat{r}_t^*$ could still exhibit serial correlation given model misspecification and the result in (12) would still hold. What ultimately matters is the serial correlation in the first difference of the preliminary estimated trend and the correction for it with a univariate BN decomposition based on an ARMA model, not what the source of the serial correlation is. Thus, our proposed correction is more broadly applicable given misspecification than just the case of unaccounted for measurement error. In our application, the consideration of a large set of variables should help mitigate concerns about omitted variables, but it is certainly possible that the apparent serial correlation in the first difference of the preliminary estimated trend could be due to an incorrect specification of model dynamics as well as, or even instead of, measurement error.

In terms of the Monte Carlo analysis in the previous subsection, these other sources of misspecification can be thought of as corresponding to the case where the true trend would be defined as the BN trend for the VARMA(1,2) with population parameters (i.e., Case 3). If we use the BN trend for this case as defining the true trend, then the RMSE for the preliminary trend estimate based on a VAR(1) (i.e., Case 4) is 0.38, while the RMSE for the corrected trend estimate using an MA(2) model for the correction (i.e., Case 5) is 0.34, corresponding to the same improvement in terms of reducing the RMSE by 0.04 as in the measurement error case presented in Table 1. Similar to before, the RMSE when estimating the VARMA (i.e., Case 2) is 0.33, so the corrected estimate is almost as accurate as knowing the true model structure in this scenario and, again, can be expected to perform even better comparatively than estimating a VARMA if there were more variables and lags.

4 Application to estimating r^*

In this section, we show the empirical relevance of our proposed correction when applied to estimating r^* using a multivariate BN decomposition based on a medium-scale vector error correction model (VECM).

4.1 Data

We model both short- and long-term interest rates when conducting our trend-cycle decomposition because we are interested in a trend that is common across maturities and because the long-term interest rate can provide information about that trend even when the short-term nominal rate is constrained by the ZLB, although it should be emphasized that we consider real, not nominal, interest rates in our analysis. In particular, modeling nominal interest rates and inflation separately in a linear setting is potentially problematic because nominal interest rates are subject to the ZLB, while U.S. inflation appears to be subject to structural breaks (see, for example, Levin and Piger, 2004; Kang et al., 2009) that, given a Fisher effect on nominal interest rates, should cancel out when considering real interest rates. Meanwhile, to the extent that the ZLB also alters the behavior of the short-term real rate, estimation of a common stochastic trend is helped by the inclusion of a long-term real rate that is less affected by the ZLB, as argued by Del Negro et al. (2017) and Bauer and Rudebusch (2020). Notably, our results are robust to allowing for a possible structural break in the term premium between the real interest rates when considering a sample period that includes the ZLB. The key point is that our linear VECM appears to be a better model to capture a common stochastic trend in real interest rates than it would be of nominal interest rates and inflation separately.[11]

We construct the short (long) *ex ante* real interest rate as the 3-month (10-year) U.S. Treasury nominal yield minus a short (long) measure of inflation expectations. For inflation, we consider the year-on-year growth rate of the core personal consumption expenditure (PCE)

[11] However, see González-Astudillo and Laforte (2020) and Johannsen and Mertens (2021) for approaches that model nominal interest rates and inflations separately. Also see the discussion in Carriero et al. (2021) how results suggesting monetary policy is unconstrained by the ZLB given unconventional policies, such as argued in Swanson and Williams (2014), support the use of constant-parameter linear models even when including data from the ZLB, such as we do with our VECM.

price deflator. We then use a 4-quarter (40-quarter) rolling average of past inflation as a proxy for short (long) inflation expectations. A downside of working with *ex ante* real interest rates is that they are subject to measurement error given the need to proxy for inflation expectations in the bond market. Even given the existence of real bonds, which in any event were not available for U.S. Treasury debt prior to 1997, there may be liquidity and covariance features that mean a 'break-even' rate formed from the difference between nominal and real yields at the same maturity does not just reflect inflation expectations and the real yields may not be cointegrated with the risk-free rate. Meanwhile, we note that the *ex post* long-term real interest rate based on U.S. data appears to be affected by nonstationary inflation expectation errors, likely due to the structural breaks in the inflation process. Unlike with the *ex ante* real interest rates, the *ex post* long-term real interest rate does not test as being cointegrated with the *ex post* short-term real interest rate, which itself exhibits sizable structural breaks in volatility presumably also related to structural breaks in the volatility of inflation expectation errors given their absence in *ex ante* real interest rates. *Ex post* real rates also have the problem of missing observations at the end of the sample, making estimates of r^* based on them less useful for any current analysis in a policy setting.

Given the challenges of measuring *ex ante* real interest rates, we also consider the robustness of our results to alternative proxies for inflation expectations, including a 4-quarter-ahead forecast based on an AR(3) model following Laubach and Williams (2003), a 4-quarter-ahead SPF forecast of GDP deflator inflation, and a 10-year-ahead SPF forecast of PCE deflator inflation, as well as 1-month and 10-year real interest rates constructed by the Cleveland Fed using their model-based measures of expected inflation. Meanwhile, the possibility of some measurement error in the *ex ante* real interest rates directly motivates the application of our proposed correction, as does the inclusion of many other variables in our model, as discussed next, given that some of these variables may also be subject to measurement error.

The choice of other variables to include in our model used for trend-cycle decomposition is motivated by the various potential 'correlates' of real interest rates outlined in Lunsford and West (2019), noting that a multivariate BN decomposition only requires variables that have informational content in forecasting real interest rates, not necessarily causal effects. Although the variables considered in Lunsford and West (2019) are annual and trace back to the 1890s, we focus on those that are available at a quarterly frequency starting at least from the 1970s. The variables can be placed into two broad categories. First, we consider supply-side productivity/demographic variables that could help explain changes in real interest rates, specifically consumption growth per capita, TFP growth, stock returns, real investment growth, employment growth, hours growth, the change in the unemployment rate, the second difference of age dependency, and the change in income inequality. Second, we consider safe asset/global savings glut variables, specifically the change in macroeconomic uncertainty, the change in the excess bond premium, the U.S. current account deficit, the change in U.S. government debt-to-GDP ratio, the return on the trade-weighted U.S. dollar exchange rate, and the change in global central bank foreign reserves-to-GDP. Full details of the specific variables in these categories and their motivating theoretical links to r^* from the literature are provided in the appendix, along with the original data sources and transformations to the raw data.

Our sample period covers 1973Q2 to 2019Q4. We have a balanced panel of the quarterly variables in our baseline model for the full sample period, although the real interest rates from the Cleveland Fed considered in our robustness analysis are only available from 1982Q1.[12] Also, the variables capturing age dependency, income inequality, and global reserves are only available at an annual frequency and so are only included in an annual version of our VECM for the sample period of 1973 to 2019 that is also considered for robustness.

4.2 Estimation

For the purposes of specifying the VECM used to estimate r^* via a multivariate BN decomposition, we denote our measures of the short- and long-term *ex ante* real interest rates as r_t^s and r_t^l, respectively, while $\mathbf{x}_{3:n,t}$ denotes the vector of other variables (in levels) that are hypothesized to drive r^*.

To impose cointegration, we add an error-correction term into the equations for the first differences of the short- and long-term real interest rates, with the error-correction coefficients denoted as β^s and β^l, respectively.[13] Assuming cointegration between the short- and long-term

[12]The 10-year-ahead SPF forecast of PCE deflator inflation that is also used in our robustness analysis is only available from the Philadelphia Fed from 2007Q1. However, we thank Todd Clark for providing earlier data for this measure used in the Fed's FRB/US macroeconometric model.

[13]We impose that the means of the first differences of the real interest rates are zero, corresponding to an assumption of no deterministic drift in levels. Thus, historical downward movements in r^* will be attributed to prediction errors rather than deterministic drift when we conduct the informational accounting in Section 4.4.

real interest rates implies that both are driven by only one stochastic trend that differs only by a constant α. The assumption of a $(1\ -1)'$ cointegrating vector in the model is supported by the stationarity of the spread between long and short real rates and an estimated coefficient when regressing the short real rate on the long real rate (including a constant in the regression) of 1.005 with a standard error of 0.071. The VECM is then specified as follows:

$$\mathbf{\Delta x}_t = \boldsymbol{\mu} + \mathbf{\Phi}_1(\mathbf{\Delta x}_{t-1} - \boldsymbol{\mu}) + \ldots + \mathbf{\Phi}_p(\mathbf{\Delta x}_{t-p} - \boldsymbol{\mu}) + \boldsymbol{\beta}\left(r^l_{t-1} - r^s_{t-1} - \alpha\right) + \mathbf{e}_t, \tag{18}$$

where $\mathbf{\Delta x}_t = (\Delta r^s_t, \Delta r^l_t, \mathbf{\Delta x}_{3:n,t}{}')'$ and $\boldsymbol{\beta} = (\beta^s, \beta^l, \mathbf{0}_{(n-2)\times 1}{}')'$. The model with the error-correction term in (18) can be cast into companion form and the BN trend for the short-term real interest rate calculated following (6), where, for the companion form, $\mathbf{\Delta X}_t = \left\{(\mathbf{\Delta x}_t - \boldsymbol{\mu})', \cdots, (\mathbf{\Delta x}_{t-p} - \boldsymbol{\mu})', \left(r^l_t - r^s_t - \alpha\right)\right\}'$ and $\mathbf{F}$ and $\mathbf{H}$ are as given in the VECM example in Morley (2002).

Given possible measurement error in the observed data or some other source of misspecification of the VECM, estimation of the companion matrix for the VECM could actually correspond to estimation of a linear projection matrix $\mathbf{P}$ instead of $\mathbf{F}$, with the preliminary BN trend estimates $\{\hat{r}^*_t\}_{t=1}^T$ displaying serial correlation in their first differences. Applying the simple correction proposed in Section 2, the corrected estimates $\{\tilde{r}^*_t\}_{t=1}^T$ can then be calculated following (17) and (3) under the assumption that the data are possibly measured with error, with the companion vector for the data denoted as $\mathbf{\Delta \check{X}}_t$ instead of $\mathbf{\Delta X}_t$.

Because we consider a medium-scale VECM that has a large number of parameters, we need to take into account the possibility of overfitting in estimation. To address this, we rely on Bayesian shrinkage following Morley and Wong (2020) using a natural conjugate Normal-Inverse Wishart prior in conjunction with a Minnesota Prior with the shrinkage hyperparameter, λ. Given the VECM setup with the error-correction term only appearing in the interest rate equations, estimation is conducted via MCMC with Gibbs sampling.[14] For computational convenience, then, we set to $\lambda = 0.2$, as in Sims and Zha (1998) and Carriero et al. (2015). We also set an "expectations hypothesis" prior mean for the error-correction coefficient in the short-rate equation of $E[\beta^s] = 0.5$, consistent with the short-rate adjusting to restore the long-

[14] A Bayesian VAR that includes the spread $r^l_t - r^s_t$ instead of the change in the long-rate Δr^l_t and error correction terms for the interest rate equations can be estimated analytically and produces reasonably similar estimates of r^*_t. However, we take the VECM as our baseline specification to allow for general error correction in terms of the short and long real interest rates and to impose no adjustment of the other variables to the error-correction term.

run cointegrating equilibrium, although this prior is not particularly tight and we find the results are completely robust to setting the prior mean to zero instead.[15] The Bayesian VECM is estimated with four lags, as is typical for quarterly data. The full details of the Bayesian estimation for the VECM are reasonably standard and can be found in the appendix.

In terms of inference about r^*, we note the following possible factorization of the joint posterior distribution of $\hat{\mathbf{r}}^* \equiv \{\hat{r}_t^*\}_{t=1}^T$, the VECM coefficients $\mathbf{F}$, and the residual variance-covariance matrix $\mathbf{\Sigma}$ in the case of no measurement error or other source of misspecification:

$$p(\hat{\mathbf{r}}^*, \mathbf{F}, \mathbf{\Sigma}|\mathbf{r}, \mathbf{\Delta X}) = p(\hat{\mathbf{r}}^*|\mathbf{F}; \mathbf{r}, \mathbf{\Delta X})p(\mathbf{F}, \mathbf{\Sigma}|\mathbf{\Delta X}). \tag{19}$$

Given an initial $\mathbf{\Sigma}^{(0)}$ and letting the superscript in parentheses denote the i^{th} draw from an MCMC sampler, we can obtain draws from $p(\mathbf{F}, \mathbf{\Sigma}|\mathbf{\Delta X})$ using a Gibbs sampler and calculate implied draws of $\hat{\mathbf{r}}^*$ using (6) by repeating the three following steps:

1. Draw from $p(\mathbf{F}^{(i)}|\mathbf{\Sigma}^{(i-1)}; \mathbf{\Delta X})$
2. Draw from $p(\mathbf{\Sigma}^{(i)}|\mathbf{F}^{(i)}; \mathbf{\Delta X})$
3. Calculate implied draw from $p(\hat{\mathbf{r}}^{*(i)}|\mathbf{F}^{(i)}; \mathbf{r}, \mathbf{\Delta X})$

However, if the preliminary estimated change in trend displays serial correlation, it suggests some misspecification such as unaccounted for measurement error. Then, the joint posterior factorization we are interested in is given as follows:

$$p(\tilde{\mathbf{r}}^*, \theta(L), \phi(L), \mathbf{P}, \mathbf{\Sigma}|\check{\mathbf{r}}, \mathbf{\Delta}\check{\mathbf{X}}) = p(\tilde{\mathbf{r}}^*|\theta(L), \phi(L), \mathbf{P}; \check{\mathbf{r}}, \mathbf{\Delta}\check{\mathbf{X}})p(\theta(L), \phi(L)|\mathbf{P}, \mathbf{\Sigma}; \mathbf{\Delta}\check{\mathbf{X}})p(\mathbf{P}, \mathbf{\Sigma}|\mathbf{\Delta}\check{\mathbf{X}}). \tag{20}$$

In this case, given an initial $\mathbf{\Sigma}^{(0)}$, we can obtain draws from $p(\mathbf{P}, \mathbf{\Sigma}|\mathbf{\Delta}\check{\mathbf{X}})$ using a Gibbs sampler, draws from $p(\theta(L), \phi(L)|\mathbf{P}, \mathbf{\Sigma}; \mathbf{\Delta}\check{\mathbf{X}})$ using a Metropolis step, and calculate implied draws of $\tilde{\mathbf{r}}^* \equiv \{\tilde{r}_t^*\}_{t=1}^T$ using (17) and (3) by repeating the four following steps:

[15] Although we consider an "expectations hypothesis" prior in our baseline estimation, we generally take an agnostic view in terms of whether or how r^* is related to supply-side variables that capture a link to productivity growth or demographics, financial variables that capture global demand for safe assets, or government debt related to a possible crowding-out effect. Specifically, we do not impose a structural link between r^* and trend output growth, as done in Laubach and Williams (2003), but instead we test such links while including a much larger set of other potential explanatory variables than considered in typical structural or semi-structural models (although see Fu, 2023). In principle, we could incorporate priors based on Granger causation implied by structural models in the Bayesian VECM if we wanted to impose more of that structure when estimating r^*. However, our more agnostic priors allow us to treat posterior inferences about signs of effects as tests of theoretical predictions, while our inferences about r^* are precise without having to rely on potentially restrictive structural assumptions.

1. Draw from $p(\mathbf{P^{(i)}}|\mathbf{\Sigma}^{(i-1)}; \mathbf{\Delta\check{X}})$

2. Draw from $p(\mathbf{\Sigma}^{(i)}|\mathbf{P}^{(i)}; \mathbf{\Delta\check{X}})$

3. Draw from $p(\theta(L)^{(i)}, \phi(L)^{(i)}|\mathbf{P}^{(i)}, \mathbf{\Sigma}^{(i)}; \mathbf{\Delta\check{X}})$

4. Calculate implied draw from $p(\tilde{\mathbf{r}}^{*(i)}|\theta(L)^{(i)}, \phi(L)^{(i)}, \mathbf{P}^{(i)}; \check{\mathbf{r}}, \mathbf{\Delta\check{X}})$

The potentially problematic aspect of the sampler under misspecification is obtaining the draws from $p(\mathbf{P}, \mathbf{\Sigma}|\mathbf{\Delta X})$. Specifically, Miranda-Agrippino and Ricco (2021) consider Bayesian local projections with unmodeled serial correlation in the projection errors and propose using an "artificial" Gaussian posterior centered at MLE for the pseudo-true value $\mathbf{P}$ with a HAC covariance matrix, following a proposal by Müller (2013). For simplicity, we do not make this correction in our application given very little detectable serial correlation in the prediction errors. In particular, the HAC correction is based on sample autocorrelations, which means that corrected inferences for $\mathbf{\Sigma}$ would be very similar to the inferences assuming no serial correlation, while the posterior without the HAC correction is still centered at MLE for the pseudo-true value $\mathbf{P}$. That is, our proposed sampler can be thought of as appropriate when $\mathbf{P} \approx \mathbf{F}$ and there is little detectable serial correlation in the prediction errors, consistent with the idea that the misspecification is due to a small amount of hard-to-detect measurement error in any given variable in our VECM and the effects of the measurement error are only evident in the first differences of the preliminary estimated trend when aggregating prediction errors across many variables. However, we note that, if one were concerned about more substantial serial correlation in the projection errors, the HAC adjustment based on Müller (2013) and Miranda-Agrippino and Ricco (2021) could be applied.

Meanwhile, the relatively non-standard element of the sampler under misspecification is the third step to draw from $p(\theta(L), \phi(L)|\mathbf{P}, \mathbf{\Sigma}; \mathbf{\Delta\check{X}})$. Noting that $\hat{r}_t^*$ in (3) can be calculated using (6) given $\mathbf{P}$, $\check{\mathbf{r}}$, and $\mathbf{\Delta\check{X}}$, this step can be thought of as estimating a univariate ARMA model for $\mathbf{\Delta\hat{r}}^*$ with a known unconditional variance $\sigma^2_{\Delta\hat{r}_t^*} = \mathbf{s}'_{k,1}(\mathbf{I}-\mathbf{P})^{-1}\mathbf{H\Sigma H}'((\mathbf{I}-\mathbf{P})^{-1})'\mathbf{s}_{k,1}$. Indeed, it may be useful to think more generally about a sampler in terms of a correction for $\hat{\mathbf{r}}^*$ instead of being based on the factorization in (20) given that this would also be relevant for some other approaches to the preliminary estimation of trend than just the BN decomposition. In particular, given a joint posterior for the preliminary trend estimates and the unconditional variance of the changes $p(\hat{\mathbf{r}}^*, \sigma^2_{\Delta\hat{r}^*_t}|\Omega_T)$, where Ω_T is the relevant information set to conduct the preliminary trend estimation, we can factorize the joint posterior of the corrected trend estimates $\tilde{\mathbf{r}}^*$, the ARMA parameters, and the preliminary trend estimates $\hat{\mathbf{r}}^*$ as follows:

$$p(\tilde{\mathbf{r}}^*, \theta(L), \phi(L), \hat{\mathbf{r}}^*|\Omega_T) = p(\tilde{\mathbf{r}}^*|\theta(L), \phi(L), \hat{\mathbf{r}}^*)p(\theta(L), \phi(L)|\hat{\mathbf{r}}^*, \sigma^2_{\Delta\hat{r}^*_t})p(\hat{\mathbf{r}}^*, \sigma^2_{\Delta\hat{r}^*_t}|\Omega_T). \quad (21)$$

Then the nonstandard step is to draw from the second conditional distribution in the factorization, $p(\theta(L)^{(i)}, \phi(L)^{(i)}|\hat{\mathbf{r}}^*, \sigma^2_{\Delta\hat{r}^*_t})$, given $\hat{\mathbf{r}}^{*(i)}$ and $\sigma^2_{\Delta\hat{r}^*_t}$. We do so using a Metropolis step given a uniform prior over the stationary and invertible space for the ARMA parameters and a scaled multivariate Normal proposal for the change in the ARMA parameters centered at zero and unscaled variance equal to the inverse Hessian from MLE based on $\Delta\hat{r}^*_t$ calculated at the posterior mean of the projection matrix $\mathbf{P}$.[16] In particular, if a proposed draw of the ARMA parameters satisfies the stationarity and invertibility conditions (i.e., eigenvalues of the companion matrix for the AR component are less than one in modulus and the factorization of the MA polynomial has roots all outside the unit circle), we set the acceptance probability for the draw to the minimum of either the ratio of the proposed and previous likelihood ordinates for the ARMA model under the assumption of normal errors or one. Otherwise, if the draw does not satisfy the stationarity and invertibility conditions, we set the acceptance probability to zero. The proposed ARMA parameters are drawn from a multivariate Normal distribution with a tuned scale parameter to achieve an acceptance rate within the 20-30% range.

The complexity for the Metropolis step is in calculating the likelihood ordinates given that we need the conditional variance σ^2_ϵ, rather than the unconditional variance $\sigma^2_{\Delta\hat{r}^*_t}$, to do so. If just MA dynamics are considered for the correction, this calculation is straightforward given that $\sigma^2_\epsilon = \sigma^2_{\Delta\hat{r}^*_t}/(1 + \theta_1^2 + \cdots + \theta_q^2)$ for an MA(q) process. However, if we consider ARMA(p,q) dynamics, then the calculation is somewhat more complicated and can be made using a state-space form of the ARMA model such that $\Delta\hat{r}^*_t = \boldsymbol{\lambda}'\boldsymbol{\zeta}_t$, where $\boldsymbol{\lambda}$ includes MA parameters, and $\boldsymbol{\zeta}_t = \boldsymbol{\Phi}\boldsymbol{\zeta}_{t-1} + \mathbf{v}_t$, where $\boldsymbol{\Phi}$ includes AR parameters and $\boldsymbol{\zeta}_t$ is $\max(p, q+1) \times 1$. Letting $\boldsymbol{\Gamma}$ denote the unconditional variance of $\boldsymbol{\zeta}_t$, the unconditional variance of $\Delta\hat{r}^*_t$ is $\sigma^2_{\Delta\hat{r}^*_t} = \boldsymbol{\lambda}'\boldsymbol{\Gamma}\boldsymbol{\lambda}$, where $vec(\boldsymbol{\Gamma}) = (\mathbf{I} - \boldsymbol{\Phi} \otimes \boldsymbol{\Phi})^{-1}vec(\mathbf{Q})$ and $\mathbf{v}_t \sim (\mathbf{0}, \mathbf{Q})$, with $\mathbf{Q}[1,1] = \sigma^2_\epsilon$ and zeros elsewhere.

[16] The Gibbs sampler to obtain draws of $\mathbf{P}$ and $\boldsymbol{\Sigma}$ can be run first and then the estimation steps associated with the correction run conditional on the draws from the Gibbs sampler. Thus, the posterior mean of $\mathbf{P}$ is available when constructing the proposal for this Metropolis step. Note, however, that we inherently consider the same number of draws of the corrected trend estimates as for the preliminary trend estimates.

Figure 1: Preliminary BN trends for short- and long-term real interest rates

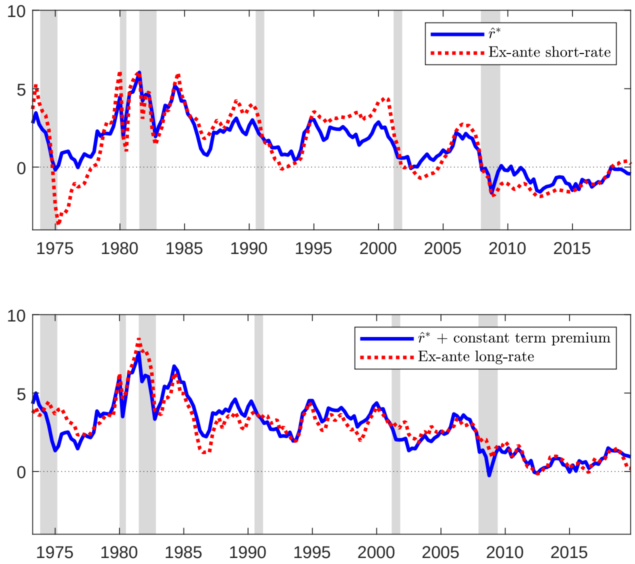

Notes: Posterior means are reported. NBER recession dates are shaded.

Thus, we can solve for σ_ϵ^2 given $\boldsymbol{\lambda}$, $\boldsymbol{\Phi}$, and $\sigma^2_{\Delta \hat{r}^*_t}$ by solving for the unconditional variance of a hypothetical ARMA(p,q) process z_t with error variance equal to 1 as $\sigma_z^2 = \boldsymbol{\lambda}' \boldsymbol{\Gamma}_z \boldsymbol{\lambda}$, where $vec(\boldsymbol{\Gamma}_z) = (\mathbf{I} - \boldsymbol{\Phi} \otimes \boldsymbol{\Phi})^{-1} vec(\mathbf{Q_z})$, with $\mathbf{Q_z}[1,1] = 1$, and then solving $\sigma_\epsilon^2 = \sigma^2_{\Delta \hat{r}^*_t} / \sigma_z^2$.

4.3 Estimates of r^*

Figure 1 plots the posterior means of the preliminary BN trends for the short- and long-term real interest rates based on a Bayesian VECM for the available quarter data. Given cointegration, the trend for the long-term real interest rate only differs from the trend for the short-term real interest rate by a constant α that reflects the long-run level of the term premium. The trend estimates appear to be highly informed by the long-rate, suggesting that the short-rate does most of the adjustment to restore the cointegrating relationship over time. The presence of a cointegrating relationship between the short- and long-term rates provides important information in estimating the real interest rates trends during two episodes in particular: during 1975-1980 and during the ZLB between 2009-2015. First, although the measured short-rate is quite negative during 1975-1980, the positive long-term real interest rate helps identify a higher and generally positive level of trend. Second, during the years that the short-term nominal interest rate was constrained by the ZLB between 2009-2015, the estimated trend is still persistently higher than the short-term real interest rate, although it is sometimes negative. Related, the multivariate BN decomposition implies a fair degree of persistence in

Figure 2: Corrected BN trend for the short-term real interest rate

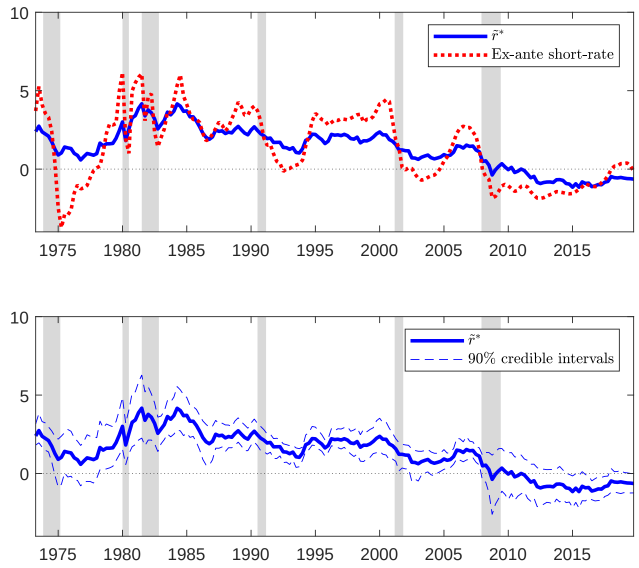

Note: Posterior mean and 90% equal-tailed credible intervals are reported. NBER recession dates are shaded.

the estimated deviations of the real interest rates from trend compared to traditional univariate BN decompositions, a more general finding with multivariate BN decompositions that is similar to Evans and Reichlin (1994) and Morley and Wong (2020).

Figure 2 presents the posterior mean of the corrected BN trend for the short-term real interest rate, along with the 90% equal-tailed credible intervals.[17] Changes in the preliminary trend estimates in Figure 1 exhibited some serial correlation despite the random walk assumption when applying the multivariate BN decomposition. We find that an MA(8) model is sufficient to capture this serial correlation.[18] Then, because we find that the MA parameter estimates imply a substantial amount of negative serial correlation in the change in the preliminary trend estimates, the corrected BN trend in Figure 2 ends up being considerably smoother than the preliminary BN trend for the short-rate in Figure 1. However, despite its relative smoothness, we can observe that r^* still varies over longer periods of time and has fallen considerably since the 1980s. The overall pattern of our r^* estimate is consistent with the previous literature (e.g.,

[17]The credible intervals capture parameter uncertainty. Specifically, given the observation of the data in (17), the only uncertainty about the corrected BN trend is due to uncertainty about the population parameters in $\mathbf{P}$, $\phi(L)$, and $\theta(L)$. Accounting for any filtering uncertainty about r^* would require making an assumption about the number of underlying structural shocks driving r. If it is the same as the number of variables in the model, then we have the "BN-as-definition" scenario as in Morley (2011) and the Monte Carlo analysis in Section 3.3. In this case, there is no filtering uncertainty. However, if there are more structural shocks than variables, we have a "BN-as-estimate" scenario and it would be necessary to specify a UC process and use the Kalman filter to capture any filtering uncertainty.

[18]The sample ACF and PACF for the posterior mean of the change in the preliminary BN trend suggest an MA(8) specification, while allowing for additional lags did not affect inferences about the corrected BN trend.

Figure 3: Comparison with UC estimates

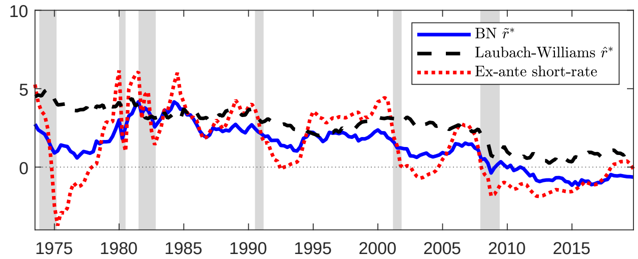

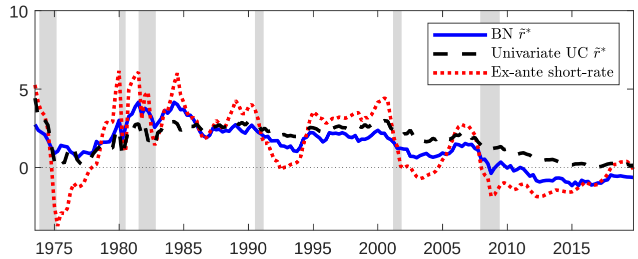

Note: NBER recession dates are shaded.

Cúrdia et al., 2015; Lubik and Matthes, 2015; Hamilton et al., 2016; Del Negro et al., 2017; Holston et al., 2017; Fiorentini et al., 2018; Brand and Mazelis, 2019; Brand et al., 2019; Berger and Kempa, 2019; Lewis and Vazquez-Grande, 2019; Bauer and Rudebusch, 2020; Kiley, 2020a; Johannsen and Mertens, 2021; Fu, 2023), and it should be emphasized that we do not impose smoothness in our estimation *a priori* unlike some other approaches (e.g., Del Negro et al., 2017). In line with our main motivation for the correction in the previous section, the negative serial correlation is consistent with the presence of classical measurement error in some of the variables in the model.

We also highlight that r^* appears to have been persistently low since the Great Recession and the estimated level has even been slightly negative since around 2013, although the estimates are only briefly significantly negative according to the 90% credible intervals. Other estimates of r^*, such as estimates based on semi-structural models in Brand and Mazelis (2019), Kiley (2020a), and Fu (2023) turn negative a few years after the Great Recession. We note that a negative r^* could occur for a variety of possible theoretical reasons, including a very low marginal product of capital and/or high levels of market power (see, for example, the neoclassical growth model with market power in Ball and Mankiw, 2023).

In Figure 3, we directly compare our estimated r^* to two particular estimates based on UC models, namely the multivariate semi-structural model in Laubach and Williams (2003) as reported on the New York Fed website and a univariate UC model with a random walk trend and AR(2) cycle estimated via MLE.[19] The Laubach and Williams (2003) estimate of r^* behaves somewhat similarly to our estimate, also declining by about 3.5ppt over the same sample period, but it has a higher average level in part due to its semi-structural link to trend output growth. The Laubach and Williams (2003) estimate also displays a similar degree of smoothness as our estimate, although Kiley (2020b) shows that the smoothness of r^* is not particularly well-identified for semi-structural models. Encouragingly, the Laubach and Williams (2003) estimate does not appear to have any serial correlation in its first difference, implying no need to apply our correction. Meanwhile, the univariate UC model also produces similar results, although it is likely misspecified. This is because, as with the Laubach and Williams (2003) model, we assume a zero correlation between movements in trend and cycle, while this restriction can be strongly rejected based on a likelihood ratio (LR) test for the univariate model (the LR statistic is 41.95), similar to the findings for real GDP in Morley et al. (2003). Furthermore, the change in the preliminary filtered estimate of r^* for the univariate UC model displays positive serial correlation, thus leading to a somewhat more volatile estimate after applying our correction, which is the estimate reported in Figure 3.

In the appendix, we consider the robustness of our estimated r^*, including in terms of (i) alternative priors, (ii) different measures of inflation expectations when constructing real interest rates, (iii) different information sets, and (iv) time variation in the long-run level of the term premium. First, we find that estimates are highly robust given a more agnostic prior about the adjustment of the short-rate to the spread and to optimizing the shrinkage hyperparameter λ to minimize the one-step-ahead out-of-sample forecast errors for the short-rate, consistent with the approach for output growth in Morley and Wong (2020). This robustness suggests our results are not driven by any particularly informative or arbitrary priors. Second, we find that estimates are reasonably robust to consideration of alternative measures of inflation expectations when constructing real interest rates, suggesting that our proposed correction is successful at addressing misspecification due to measurement error in real interest rate measures given that

[19] We should also note that our estimate bears a striking resemblance to the medium-run forecast of r^* for the dynamic stochastic general equilibrium (DSGE) model in Del Negro et al. (2017) reported and compared to the Laubach and Williams (2003) estimate in their Figure 12. This medium-run forecast is a more comparable estimate to our estimate of r^* than an infinite-horizon forecast for their Bayesian DSGE model given that the infinite-horizon forecast is a constant by construction, while our Bayesian VECM suggests that predictable momentum in the short-term real interest rate dies out over a medium-run horizon. Our estimate is also quite similar to some of the reported estimates in Fiorentini et al. (2018), Kiley (2020a), and Fu (2023) based on semi-structural models.

different implied measurement error for different measures of inflation expectations still results in similar corrected estimates. Third, the estimates are reasonably robust for smaller information sets, being almost identical for a seven-variable VECM that includes the real interest rates and the five most informative variables according to the procedure in Morley and Wong (2020). When we consider a bivariate VECM with just the real interest rates, the corrected estimates are a bit more variable than in our baseline case. But they are still considerably smoother than the preliminary estimates, suggesting that our correction can also handle misspecification due to omitted variables and that a likely source of misspecification comes from measurement error in the *ex ante* real interest rate measures in particular, although the fact the baseline estimates are a bit smoother suggests the possible presence of some measurement error in some of the other variables and a more general benefit of including relevant variables rather than addressing their omission via the correction. Fourth, the estimates are extremely robust to consideration of time variation in the long-run level of the term premium α by estimating it using a backward-looking 40-quarter rolling window, which Kamber et al. (2018) refer to as 'dynamic demeaning'. Thus, the possibility of changes in the long-run relationship between short- and long-term real interest rates over time, including during the ZLB, does not appear to distort our baseline estimates that assume no permanent movements in the spread.

We also conduct three additional exercises in the appendix that are more straightforward to consider with the bivariate VECM that includes just the short- and long-term real interest rates. First, we examine the real-time reliability in an Orphanides and van Norden (2002) of our trend estimates. In this case, we construct the real interest rates using the SPF survey measures of inflation expectations because they are not subject to revision, thus revealing the role of estimation uncertainty in the real-time reliability of the trend estimates. We find that the real-time estimates are highly reliable. Second, we allow for the possibility of stochastic volatility and find that the trend estimates are highly robust to those for the bivariate VECM assuming constant volatility. There is clear evidence of some time variation in volatility. However, the BN trend corresponds to a long-horizon point forecast of the the target variable rather than a density forecast and point forecasts can be quite robust to controlling for heteroskedasticity, while those based on an assumption of constant volatility are often easier to calculate, especially given a medium-scale VECM as considered in our baseline analysis.

Third, we consider frequentist inference with bootstrapping to account for uncertainty in our corrected trend estimate and show that the inferences are very similar to the Bayesian case for the bivariate model. The reason we focus on Bayesian inference in our baseline analysis is the reasonably large number of parameters in the medium-scale VECM that allows us to consider which variables are informative about movements in r^*, which is what we examine next.

4.4 Why did r^* change over time?

To examine which variables in our medium-scale Bayesian VECM are informative about the historical changes in r^*, we conduct an informational decomposition along the lines of Morley and Wong (2020).[20] A complication with doing so is that our proposed correction produces estimated trend changes that are based on the univariate error ϵ_t for $\Delta\hat{r}_t^*$ from an ARMA model rather than underlying forecast errors for the different variables in the VECM, as in (8). However, we note that, given the invertible representation for $\theta(L)$, we can solve for the error in the ARMA model for $\Delta\hat{r}_t^*$ in terms of current and past projection errors $\boldsymbol{\eta}_t$ from the linear projection for $\boldsymbol{\Delta}\check{\mathbf{X}}_t$:

$$\epsilon_t = \theta(L)^{-1}\phi(L)\Delta\tilde{r}_t = \theta(L)^{-1}\phi(L)\mathbf{s}'_{k,1}(\mathbf{I}-\mathbf{P})^{-1}\mathbf{H}\boldsymbol{\eta}_t, \tag{22}$$

Then the informational contributions of the *current* projection errors to the corrected estimated change in trend $\Delta\tilde{r}_t^*$ are given by

$$\Delta\tilde{r}_{it}^* \equiv \frac{\theta(1)}{\phi(1)}\tilde{\omega}_i\eta_{it}, \tag{23}$$

where the $\tilde{\omega}_i$ weights are the elements of the $1\times n$ row vector $\tilde{\boldsymbol{\omega}} \equiv \mathbf{s}'_{k,1}(\mathbf{I}-\mathbf{P})^{-1}\mathbf{H}$ and, assuming each individual AR and MA coefficient is relatively small, we have the following approximation:

$$\sum_{i=1}^{n}\Delta\tilde{r}_{it}^* \approx \Delta\tilde{r}_t^*. \tag{24}$$

[20] By considering a large set of variables that have been hypothesized to explain changes in r^*, our analysis is similar to Lunsford and West (2019), although their analysis focused on the sign of long-run correlations between measured real interest rates and the hypothesized variables using a long sample of annual data from 1870 to 2016, while we focus on estimating r^* using quarterly data over a more recent sample period starting from the 1970s and quantifying the contributions of hypothesized variables to historical movements in r^* in particular. Our analysis also provides a complement to Rachel and Smith (2017), who apply a simple accounting framework based on reduced-form elasticity estimates, and Rachel and Summers (2019), who combine evidence from several structural models, in terms of accommodating a wide range of explanations for why interest rates are low. Also see Marx et al. (2021) for a decomposition of the role of different driving forces based on a structural model that nests various hypothesized reasons for the decline in interest rates.

That is, the informational contributions of the projection errors for each variable to the corrected estimate $\Delta\tilde{r}_t^*$ are approximately proportional to those for the preliminary estimate, with the factor of proportionality equal to the long-run multiplier $\theta(1)/\phi(1)$ for the ARMA model. In practice, we find that the approximation in (24) is quite accurate, consistent with $\mathbf{P} \approx \mathbf{F}$.

Figure 4 plots posterior densities for the contributions of different groups of variables to changes in r^*. Over the whole sample, the productivity/demographic variables appear to contribute the most to the decline in r^*, with the bulk of the posterior density below zero. Safe asset supply (i.e., change in government debt as a % of GDP) also seems to have contributed to at least a small amount of the overall decline, while safe asset demand (all of the other quarterly variables in the safe asset demand/supply category) shows no obvious evidence of contributing much to the decline over the full sample period. However, if we look at the contributions over the three key episodes corresponding to (i) the onset of the Great Moderation (1977Q1-1989Q4), the end of the Great Moderation (1990Q1-2005Q4), and (ii) Secular Stagnation (2006Q1-2019Q4), we can see that each of the different groups of variables made substantial contributions to changes in r^* over time. For example, the posterior densities for productivity/demographics and safe asset demand/supply all suggest positive contributions to an estimated 1.4 percentage point increase in r^* with the onset of the Great Moderation, with a high probability that productivity/demographics and safe asset supply contribute to the bulk of this increase. Likewise, productivity/demographics and safe asset supply appear to contribute to a subsequent estimated 1.4 percentage point decline in r^* up to the end of the Great Moderation, with safe asset supply having the highest probability of contributing to the decline. Finally, productivity/demographics and safe asset demand appear to contribute sizeable and similar amounts to the estimated 2.1 percentage point decline in r^* during the Secular Stagnation era, while safe asset supply appears to have a partially offsetting effect.

To confirm that the variables in the model reflect the underlying forces of productivity/demographics and safe asset demand/supply, we also consider the posterior probability of the direction of correlation between the projection error for a variable and the implied change in the estimate of r^* based on the decomposition in (23). In particular, the sign of the correlation between the overall change in trend and projection error η_{it} corresponds to the sign of $\tilde{\omega}_i$ in (23) and is determined by $(\mathbf{I} - \mathbf{P})^{-1}$. As reported in the appendix, the posterior probabilities

Figure 4: Contributions to r^*: Posterior densities

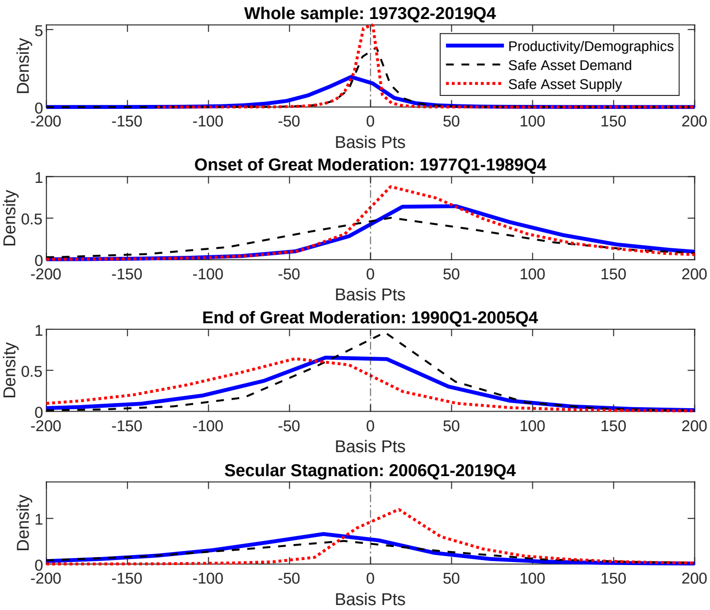

Note: Densities are calculated by applying Matlab *ksdensity* function to MCMC draws of contributions.

support the theoretically predicted signs for all of the quarterly variables that we consider. Individual contributions of each variable to the estimated r^* during the three subsample episodes considered in Figure 4 are also reported in the appendix.[21]

Because certain variables are only available at an annual frequency, we repeat our analysis using annual data to check the robustness of our baseline results when also including these additional variables. For brevity, the results for annual data are also reported in the appendix. To summarize these results, we find that the corrected estimate of r^* is even smoother than in the quarterly case, possibly reflecting overfitting the real interest rates in sample with a VECM and MA model for the correction given the small number of annual observations, even given Bayesian shrinkage for the VECM. The key finding, though, is that the variables which are only available annually and capture income inequality, demographics (i.e., age dependency), and the global savings glut (i.e., global reserves-to-GDP), are found to only contribute negligibly to the estimated movements in r^* and their inclusion in the model does not alter inferences about contributions of the other variables. Thus, we argue that their omission from the baseline quarterly VECM is not an issue. In any event, the application of our correction when conducting

[21]We note that the informational contributions correspond to lower bound effects of driving factors if the forecast errors for interest rates are correlated with the other forecast errors and the correlation is driven by causation from other variables to interest rates. In this sense, our informational decomposition can be thought of as mainly providing information on relative importance of variables rather than their overall importance.

trend-cycle decomposition using the quarterly model helps address the possibility of omitted variables, as previously discussed in Section 3.4.

Finally, we extend our estimation of r^* and informational contributions to cover the onset of the COVID-19 pandemic. The full results are presented in the appendix. The main finding for this extension is that there is a sharp decrease in the estimated level of r^* with the onset of the pandemic, but a quick recovery to a slightly higher level than the pre-pandemic level. These movements reflect the behaviour of productivity variables, but also the offsetting effects of higher asset supply and demand due to debt-financed fiscal stimulus and heightened uncertainty during the pandemic.

5 Conclusion

We have proposed a simple correction to estimating a random-walk trend of a time series when there is apparent evidence of misspecification in preliminary estimates in the form of serial correlation in first differences. The correction is based on applying a univariate Beveridge-Nelson decomposition to the preliminary estimated trend. We show how and why the correction works based on the law of iterated expectations in the setting of a small amount of otherwise hard-to-detect measurement error in some of the variables considered for trend-cycle decomposition. However, the correction is also applicable given other possible sources of misspecification. Our application to estimating r^* with a medium-scale vector error correction model reveals the empirical relevance of the correction, with the corrected estimates for r^* turning out to be considerably smoother than the preliminary estimates. Given a number of variables and lags for the model, the likely source of misspecification in the preliminary estimated trend is measurement error, including in *ex ante* measures of real interest rates. An informational decomposition following Morley and Wong (2020) suggests productivity/demographics and safe asset supply/demand explain major historical movements in r^*. Our estimates are comparable to others in the literature, but importantly we do not impose smoothness in our estimation, rather we find it when including a wide range of multivariate information motivated by different theories of what drives r^* and applying our correction for misspecification. Future research will apply the correction to other estimates of random-walk trends that exhibit serial correlation in their first differences.

References

Anderson, B. D., M. Deistler and J.-M. Dufour, "On the Sensitivity of Granger Causality to Errors-in-Variables, Linear Transformations and Subsampling," *Journal of Time Series Analysis* 40 (2019), 102–123.

Baker, D., J. B. De Long and P. R. Krugman, "Asset Returns and Economic Growth," *Brookings Papers on Economic Activity* 2005 (2005), 289–330.

Ball, L. and N. G. Mankiw, "What Do Budget Deficits Do?," NBER Working Paper No. 5263, National Bureau of Economic Research, 1995.

———, "Market Power in Neoclassical Growth Models," *Review of Economic Studies* 90 (2023), 572–596.

Banbura, M., D. Giannone and L. Reichlin, "Large Bayesian Vector Auto Regressions," *Journal of Applied Econometrics* 25 (2010), 71–92.

Bauer, M. D. and G. D. Rudebusch, "Interest Rates under Falling Stars," *American Economic Review* 110 (2020), 1316–1354.

Berger, T. and B. Kempa, "Testing for Time Variation in the Natural Rate of Interest," *Journal of Applied Econometrics* 34 (2019), 836–842.

Bernanke, B. S., "The Global Saving Glut and the US Current Account Deficit, Homer Jones Lecture," *St. Louis, Missouri* (2005).

Beveridge, S. and C. R. Nelson, "A New Approach to Decomposition of Economic Time Series Into Permanent and Transitory Components With Particular Attention to Measurement of the 'Business Cycle'," *Journal of Monetary Economics* 7 (1981), 151–174.

Brand, C., M. Bielecki and A. Penalver, "The Natural Rate of Interest: Estimates, Drivers, and Challenges to Monetary Policy," NBP Working Papers 308, Narodowy Bank Polski, 2019.

Brand, C. and F. Mazelis, "Taylor-rule consistent estimates of the natural rate of interest," Working Paper Series 2257, European Central Bank, 2019.

Caballero, R. J., E. Farhi and P.-O. Gourinchas, "The Safe Assets Shortage Conundrum," *Journal of Economic Perspectives* 31 (2017), 29–46.

Carriero, A., J. Chan, T. E. Clark and M. Marcellino, "Corrigendum to "Large Bayesian Vector Autoregressions with Stochastic Volatility and Non-Conjugate Priors" [J. Econometrics 212 (1) (2019) 137–154]," *Journal of Econometrics* 277 (2022), 506–512.

Carriero, A., T. E. Clark and M. Marcellino, "Bayesian VARs: Specification Choices and Forecast Accuracy," *Journal of Applied Econometrics* 30 (2015), 46–73.

Carriero, A., T. E. Clark, M. Marcellino and E. Mertens, "Forecasting with Shadow-Rate VARs," Working Paper 21-09, Federal Reserve Bank of Cleveland,, 2021.

Christiano, L. J., M. Eichenbaum and R. Vigfusson, "Assessing Structural VARs," in *NBER Macroeconomics Annual 2006, Volume 21* (2007), 1–106.

Cúrdia, V., A. Ferrero, G. C. Ng and A. Tambalotti, "Has U.S. Monetary Policy Tracked the Efficient Interest Rate?," *Journal of Monetary Economics* 70 (2015), 72–83.

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

Dufour, J.-M. and D. Pelletier, "Practical Methods for Modeling Weak VARMA Processes: Identification, Estimation and Specification with a Macroeconomic Application," *Journal of Business & Economic Statistics* 40 (2022), 1140–1152.

Dungey, M., J. P. Jacobs, J. Tian and S. van Norden, "Trend in Cycle or Cycle in Trend? New Structural Identifications for Unobserved-Components Models of U.S. Real GDP," *Macroeconomic Dynamics* 19 (2015), 776–790.

Dynan, K. E., J. Skinner and S. P. Zeldes, "Do the Rich Save More?," *Journal of Political Economy* 112 (2004), 397–444.

Eichengreen, B., "Secular Stagnation: The Long View," *American Economic Review* 105 (2015), 66–70.

Evans, G. and L. Reichlin, "Information, Forecasts, and Measurement of the Business Cycle," *Journal of Monetary Economics* 33 (1994), 233–254.

Fernald, J. G., "Productivity and Potential Output before, during, and after the Great Recession," *NBER Macroeconomics Annual* 29 (2015), 1–51.

FIORENTINI, G., A. GALESI, G. PEREZ-QUIRÓS AND E. SENTANA, "The Rise and Fall of the Natural Interest Rate," Working Papers 1822, Banco de España, 2018.

FU, B., "Measuring the Trend Real Interest Rate in a Data-Rich Environment," *Journal of Economic Dynamics and Control* 147 (2023), 104606.

GAGNON, E., B. K. JOHANNSEN AND D. LÓPEZ-SALIDO, "Understanding the New Normal: The Role of Demographics," *IMF Economic Review* 69 (2021), 357–390.

GILCHRIST, S. AND E. ZAKRAJŠEK, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review* 102(4) (2012), 1692–1720.

GLICK, R., "r* and the Global Economy," *Journal of International Money and Finance* 102 (2020), 102105.

GONZÁLEZ-ASTUDILLO, M. AND J.-P. LAFORTE, "Estimates of r* Consistent with a Supply-Side Structure and a Monetary Policy Rule for the U.S. Economy," Finance and Economics Discussion Series 2020-085, Board of Governors of the Federal Reserve System (U.S.), September 2020.

HAMILTON, J. D., E. S. HARRIS, J. HATZIUS AND K. D. WEST, "The Equilibrium Real Funds Rate: Past, Present, and Future," *IMF Economic Review* 64 (2016), 660–707.

HOLSTON, K., T. LAUBACH AND J. C. WILLIAMS, "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics* 108 (2017), S59–S75.

JOHANNSEN, B. K. AND E. MERTENS, "A Time-Series Model of Interest Rates with the Effective Lower Bound," *Journal of Money, Credit and Banking* 53 (2021), 1005–1046.

JURADO, K., S. C. LUDVIGSON AND S. NG, "Measuring Uncertainty," *American Economic Review* 105(3) (2015), 1177–1216.

KAMBER, G., J. MORLEY AND B. WONG, "Intuitive and Reliable Estimates of the Output Gap from a Beveridge-Nelson Filter," *Review of Economics and Statistics* 100 (2018), 550–566.

KANG, K. H., C.-J. KIM AND J. MORLEY, "Changes in U.S. Inflation Persistence," *Studies in Nonlinear Dynamics and Econometrics* 13 (2009), 1–21.

KILEY, M. T., "Output Gaps," *Journal of Macroeconomics* 37 (2013), 1–18.

———, "The Global Equilibrium Real Interest Rate: Concepts, Estimates, and Challenges," *Annual Review of Financial Economics* 12 (2020a), 305–326.

———, "What Can the Data Tell Us about the Equilibrium Real Interest Rate?," *International Journal of Central Banking* 16 (2020b), 181–209.

KILIAN, L., "Small-Sample Confidence Intervals for Impulse Response Functions," *Review of Economics and Statistics* 80 (1998), 218–230.

KOOP, G. AND D. KOROBILIS, *Bayesian Multivariate Time Series Methods for Empirical Macroeconomics* (Now Publishers Inc, 2010).

LAUBACH, T. AND J. C. WILLIAMS, "Measuring the Natural Rate of Interest," *Review of Economics and Statistics* 85 (2003), 1063–1070.

LENZA, M. AND G. E. PRIMICERI, "How to Estimate a Vector Autoregression after March 2020," *Journal of Applied Econometrics* 37 (2022), 688–699.

LEVIN, A. T. AND J. M. PIGER, "Is Inflation Persistence Intrinsic in Industrial Economies?," Eurosystem Inflation Persistence Network Working Paper No. 334, European Central Bank, April 2004.

LEWIS, K. F. AND F. VAZQUEZ-GRANDE, "Measuring the Natural Rate of Interest: A Note on Transitory Shocks," *Journal of Applied Econometrics* 34 (2019), 425–436.

LITTERMAN, R. B., "Forecasting with Bayesian Vector Autoregressions — Five Years of Experience," *Journal of Business & Economic Statistics* 4 (1986), 25–38.

LUBIK, T. AND C. MATTHES, "Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches," *Richmond Fed Economic Brief* Oct (2015).

LUNSFORD, K. G. AND K. D. WEST, "Some Evidence on Secular Drivers of US Safe Real Rates," *American Economic Journal: Macroeconomics* 11 (2019), 113–139.

LÜTKEPOHL, H., *New Introduction to Multiple Time Series Analysis* (Springer, 2005).

MARX, M., B. MOJON AND F. R. VELDE, "Why Have Interest Rates Fallen far Below the Return on Capital?," *Journal of Monetary Economics* 124S (2021), S58–S76.

MERTENS, E., "Are Spectral Estimators Useful for Long-Run Restrictions in SVARs?," *Journal of Economic Dynamics and Control* 36 (2012), 1831–1844.

Miranda-Agrippino, S. and G. Ricco, "Bayesian Local Projections," SciencePo Working papers hal-03373574, 2021.

Modigliani, F. and R. J. Shiller, "Inflation, Rational Expectations and the Term Structure of Interest Rates," *Economica* 40 (1973), 12–43.

Morley, J. and B. Wong, "Estimating and Accounting for the Output Gap With Large Bayesian Vector Autoregressions," *Journal of Applied Econometrics* 35 (2020), 1–18.

Morley, J. C., "A State–Space Approach to Calculating the Beveridge–Nelson Decomposition," *Economics Letters* 75 (2002), 123–127.

———, "The Two Interpretations of the Beveridge–Nelson Decomposition," *Macroeconomic Dynamics* 15 (2011), 419–439.

Morley, J. C., C. R. Nelson and E. Zivot, "Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different?," *Review of Economics and Statistics* 85 (2003), 235–243.

Müller, U. K., "Risk of Bayesian Inference in Misspecified Models, and the Sandwich Covariance Matrix," *Econometrica* 81 (2013), 1805–1849.

Orphanides, A. and S. van Norden, "The Unreliability of Output Gap Estimates in Real Time," *Review of Economics and Statistics* 84 (2002), 569–583.

Rachel, Ł. and T. D. Smith, "Are Low Real Interest Rates Here to Stay?," *International Journal of Central Banking* 13 (2017), 1–42.

Rachel, Ł. and L. Summers, "On Secular Stagnation in the Industrialized World," *Brookings Papers on Economic Activity* 2019 (2019), 1–54.

Rotemberg, J. J. and M. Woodford, "Real-Business-Cycle Models and the Forecastable Movements in Output, Hours, and Consumption," *American Economic Review* 86 (1996), 71–89.

Schorfheide, F. and D. Song, "Real-Time Forecasting with a (Standard) Mixed-Frequency VAR During a Pandemic," Working paper, 2020.

Sims, C. and T. Zha, "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review* 39(4) (1998), 949–968.

Summers, L. H., "Demand Side Secular Stagnation," *American Economic Review* 105 (2015), 60–65.

Swanson, E. and J. Williams, "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates," *American Economic Review* 104 (2014), 3154–3185.

Taylor, J. B., "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39 (1993), 195–214.

Zaman, S., "A Unified Framework to Estimate Macroeconomic Stars," Working paper, 2023.

Appendix

A Details of the data

In this appendix, we provide the details of the specific variables hypothesized to drive r^* in the broad categories of productivity/demographics and safe asset supply/demand. We then report the data sources and transformations.

Productivity/demographics

Motivated by an intertemporal IS/Euler-type equation, such as in Lunsford and West (2019), we consider real consumption growth per capita. Related, we also consider TFP growth (Fernald, 2015) and S&P 500 stock returns on the basis that they might be additionally informative about expected trend growth for the economy, which Laubach and Williams (2003) highlight as the key positive determinant of r^*. By contrast, Eichengreen (2015) stresses the importance of investment-specific technological change and the subsequent decline in the price of capital goods in driving down real interest rates. Thus, we also consider real investment growth as a proxy for investment-specific technological change and expect it to have a negative relationship with r^*, at least when controlling for consumption growth and TFP growth.

Various labor-market variables reflect demographic factors and are hypothesized to influence r^* through an effect on the marginal product of capital. For example, Baker et al. (2005) note that in certain overlapping-generations models, labor-force growth is positively related to the real interest rate given that higher labor-force participation would lead, all else equal, to a lower level of capital per worker. Thus, we also consider employment growth, hours growth (to capture the intensive margin), and the change in the unemployment rate as additional possible supply-side variables, although clearly decreases in employment and hours and increases in the unemployment rate could be also be related to a decline in r^* via insufficient demand, as argued by Summers (2015). The unemployment rate also serves as a potential control for economic slack that could distort measures of trend growth and generate short-run deviations in the real interest rate from r^*.

Possible heterogeneity in marginal propensities to consume motivates consideration of income inequality and age dependency. Dynan et al. (2004) find that higher income families have lower marginal propensities to consume, suggesting that an increase in inequality will shift the savings schedule out and lower r^*. Gagnon et al. (2021), on the other hand, suggest that an increase in the dependency (older-to-working) ratio reduces aggregate savings and raises r^*. To capture these demographic factors, we consider the share of wealth held by the top 1% and the age dependency ratio, although these series are only available at an annual frequency and so are only considered in robustness analysis using a VECM with annual data in Appendix F.

Safe asset demand/supply

Caballero et al. (2017) and Del Negro et al. (2017) suggest that demand for safe assets has played a key role in lowering r^* in recent decades. To address this, we consider the change in macroeconomic uncertainty (Jurado et al., 2015), the change in the excess bond premium (Gilchrist and Zakrajšek, 2012), and growth of liquid assets held by financial and non-financial corporate businesses.

Also related to demand for safe assets, Bernanke (2005) suggests a relationship between the U.S. current account deficit and the global savings glut. Capital inflows are typically associated with a trade deficit, but the link to r^* depends on whether those capital flows are induced by a high real interest rate or reflect excess global savings. To address this, we consider the change in the U.S. current account balance (as % of GDP), the change in U.S. government debt (as % of GDP), the trade-weighted U.S. dollar exchange rate growth rate, and global central bank foreign reserves (as % of world GDP), although the global reserves variable is only available at an annual frequency and so is only considered in the robustness analysis in Appendix F. An increase in government expenditure or a decrease in tax revenues that lead to a higher level of government debt is usually thought to raise real interest rates through a crowding-out effect (see, for example, Ball and Mankiw, 1995). So the government debt measure can be thought of as reflecting the supply of safe assets, while the other measures are designed to help capture demand for safe assets that push the real interest rate in the opposite direction.

Prior to inclusion in the VECM, the data are transformed to be stationary. As discussed in Morley and Wong (2020), the BN decomposition calculations in (6) and (7) require specification of the forecasting model in a stationary form. The transformations, along with the original data sources, are given in Table A1.

Table A1: Data sources and transformations

Variable Description	**Source**	**Transformation**
3-Month Treasury Bill Secondary Market Rate	FRED:TB3MS	quarterly avg., Δ
Market Yield on U.S. Treasury Securities at 10-Year Constant Maturity	FRED:GS10	quarterly avg., Δ
Personal Consumption Expenditures Excluding Food and Energy (Chain-Type Price Index)	FRED:PCEPILFE	$\%\Delta_4$
Survey of Professional Forecasters 1-Year-Ahead GDP Deflator Inflation Rate, Median Forecast	Phil.Fed:INFPGDP1YR	
Survey of Professional Forecasters 10-Year PCE Inflation Rate, Mean Response, Annual Average	Phil.Fed:PCE10,T.Clark	
Cleveland Fed 1-Month Real Rate using Model-Based Expected Inflation	clevelandfed.org	quarterly avg., Δ
Cleveland Fed 10-Year Real Rate using Model-Based Expected Inflation	clevelandfed.org	quarterly avg., Δ
Real personal consumption expenditures per capita	FRED:A794RX0Q048SBEA	ln, Δ
Business Sector TFP (annualized quarterly % growth rate)	frbsf.org	$\ln(1 + \text{series}/400)$
S&P 500 Index	FRED:SP500	quarterly avg., ln, Δ
Real Gross Private Domestic Investment	FRED:GPDIC1	ln, Δ
All Employees: Total Nonfarm	FRED:PAYEMS	quarterly avg., ln, Δ
Business Sector: Hours Worked for All Employed Persons	FRED:HOABS	ln, Δ
Unemployment Rate	FRED:UNRATE	quarterly avg., Δ
Age Dependency Ratio: Older Dependents to Working-Age Population for the United States	FRED:SPPOPDPNDOLUSA	annual only, Δ^2
Top 1% Share of Pre-Tax National Income	World Inequality Database	annual only, Δ
1-Month-Ahead Economic Macro Uncertainty Index	sydneyludvigson.com	Δ
Excess Bond Premium	federalreserve.gov	quarterly avg., Δ
Nonfinancial Corporate Business and Other Financial Corporations, Money Market Funds, Insurance Companies, and Pension Funds; Liquid Assets (Broad Measure), Level	FRED:BOGZ1FL104001005Q, BOGZ1FL874001005Q	sum, ln, Δ
Balance on Current Account as a Percent of Gross Domestic Product	FRED:NETFI, GDP	ratio, Δ
Nominal Major Currencies U.S. Dollar Index (Goods Only)	FRED:TWEXMMTH	quarterly avg., ln, Δ
Federal Debt: Total Public Debt as Percent of Gross Domestic Product	FRED:GFDEGDQ188S	Δ
Total reserves comprising holdings of monetary gold, special drawing rights, reserves of members held by the IMF, and holdings of foreign exchange under the control of monetary authorities as a percent of world GDP at purchaser's prices (data are in current U.S. dollars, with gold component of reserves valued at year-end prices and GDP converted from domestic currencies using single-year official exchange rates)	IMF:FI.RES.TOTL.CD, NY.GDP.MKTP.CD	annual only, ratio, Δ

B Bayesian estimation of the VECM

In this appendix, we present the details of the Bayesian estimation of the VECM.

Recall that $\boldsymbol{\Delta}\mathbf{x}_t$ consists of the first differences of the interest rates Δr_t^s and Δr_t^l (which we now denote as Δx_{1t} and Δx_{2t} for convenience) and the 'correlates' $\boldsymbol{\Delta}\mathbf{x}_{3:n,t}$. Because of the error correction term, the regressors differ between the 'interest rate' and 'correlates' blocks of the VECM. Thus, we specify the i^{th} equation of the VECM as

$$\Delta x_{it} = \mu_i + \mathbf{w}_{it}'\mathbf{b}_i + e_{it}, \tag{B1}$$

where $\mathbf{w}_{it} = [(\boldsymbol{\Delta}\mathbf{x}_{t-1} - \boldsymbol{\mu})', \ldots, (\boldsymbol{\Delta}\mathbf{x}_{t-p} - \boldsymbol{\mu})', r_{t-1}^l - r_{t-1}^s - \alpha]'$ if $i = 1, 2$ and $\mathbf{w}_{it} = [(\boldsymbol{\Delta}\mathbf{x}_{t-1} - \boldsymbol{\mu})', \ldots, (\boldsymbol{\Delta}\mathbf{x}_{t-p} - \boldsymbol{\mu})']'$ if $i > 2$, with $\mathbf{b}_i$ corresponding to all of the parameters associated with equation i. The unconditional means μ_i are based on sample averages, consistent with diffuse priors on these parameters, except for $i = 1, 2$ where the means of the changes in real interest rates are set exactly to zero, implying no drift in levels.

Defining

$$\boldsymbol{y}_t \equiv \boldsymbol{\Delta}\mathbf{x}_t - \boldsymbol{\mu}, \;\; \boldsymbol{\beta} \equiv \begin{bmatrix} \mathbf{b}_1 \\ \vdots \\ \mathbf{b}_n \end{bmatrix} \text{ and } \boldsymbol{Z}_t \equiv \begin{bmatrix} \mathbf{w}_{1t}' & \mathbf{0} & \ldots & \mathbf{0} \\ \mathbf{0} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \ldots & \mathbf{0} & \mathbf{w}_{nt}' \end{bmatrix}, \;\; \boldsymbol{e}_t \equiv \begin{bmatrix} e_{1t} \\ \vdots \\ e_{nt} \end{bmatrix},$$

we can stack all the equations and regressors in (B1) and rewrite the system as

$$\boldsymbol{y}_t = \boldsymbol{Z}_t\boldsymbol{\beta} + \boldsymbol{e}_t,$$

or

$$\boldsymbol{y} = \boldsymbol{Z}\boldsymbol{\beta} + \boldsymbol{E},$$

where

$$\boldsymbol{y_i} = \begin{bmatrix} y_{i1} \\ \vdots \\ y_{iT} \end{bmatrix}, \;\; \boldsymbol{e_i} = \begin{bmatrix} e_{i1} \\ \vdots \\ e_{iT} \end{bmatrix},$$

and

$$\boldsymbol{y} = \begin{bmatrix} \boldsymbol{y}_1 \\ \vdots \\ \boldsymbol{y}_n \end{bmatrix}, \quad \boldsymbol{Z} = \begin{bmatrix} \boldsymbol{Z}_1 \\ \vdots \\ \boldsymbol{Z}_T \end{bmatrix}, \quad \boldsymbol{E} = \begin{bmatrix} \boldsymbol{e}_1 \\ \vdots \\ \boldsymbol{e}_n \end{bmatrix}.$$

Let $\boldsymbol{\Sigma}$ be an $n \times n$ covariance matrix for the VECM residuals. If one sets a Normal-Wishart prior on $\boldsymbol{\beta}$ and $\boldsymbol{\Sigma}$ (Koop and Korobilis (2010)), where

$$\begin{aligned} \boldsymbol{\beta} &\sim N(\boldsymbol{\beta}_0, \boldsymbol{V}_{\boldsymbol{\beta}}), && \text{(B2)} \\ \boldsymbol{\Sigma}^{-1} &\sim W(\boldsymbol{S}_0^{-1}, \nu_0), && \text{(B3)} \end{aligned}$$

this implies conditional distributions

$$\begin{aligned} p(\boldsymbol{\beta} \mid \boldsymbol{y}, \boldsymbol{\Sigma}^{-1}) &\sim N(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{V}_{\boldsymbol{\beta}}}), && \text{(B4)} \\ p(\boldsymbol{\Sigma}^{-1} \mid \boldsymbol{y}, \boldsymbol{\beta}) &\sim W(\hat{\boldsymbol{S}}^{-1}, \hat{\boldsymbol{\nu}}), && \text{(B5)} \end{aligned}$$

where

$$\begin{aligned} \hat{\boldsymbol{V}_{\boldsymbol{\beta}}} &= \left({\boldsymbol{V}_{\boldsymbol{\beta}}}^{-1} + \sum_{t=1}^{T} \boldsymbol{Z}_t' \boldsymbol{\Sigma}^{-1} \boldsymbol{Z}_t \right), \\ \hat{\boldsymbol{\beta}} &= \hat{\boldsymbol{V}_{\boldsymbol{\beta}}} \left[{\boldsymbol{V}_{\boldsymbol{\beta}}}^{-1} \boldsymbol{\beta}_0 + \sum_{t=1}^{T} \boldsymbol{Z}_t' \boldsymbol{\Sigma}^{-1} \boldsymbol{Z}_t \right], \end{aligned}$$

and

$$\begin{aligned} \hat{\boldsymbol{S}} &= \boldsymbol{S}_0 + \sum_{t=1}^{T} (\boldsymbol{y}_t - \boldsymbol{Z}_t \boldsymbol{\beta}) (\boldsymbol{y}_t - \boldsymbol{Z}_t \boldsymbol{\beta})', \\ \hat{\boldsymbol{\nu}} &= T + \nu_0. \end{aligned}$$

We elaborate how priors $\boldsymbol{\beta}_0$, $\boldsymbol{V}_{\boldsymbol{\beta}}$, $\boldsymbol{S}_0^{-1}$ and ν_0 are elicited below. Given the priors, (B4) and (B5) define a Gibbs-sampling scheme, where one can sequentially take draws from these conditional distributions, conditioning on the previous draw in the chain. We take 12,000 draws with the sampling scheme, discarding the first 2,000 draws and use the remaining 10,000 draws to make inferences about the posterior distribution.

Priors

Our goal in setting the prior is to apply shrinkage to mitigate possible overfitting. To keep the application of shrinkage as standard as possible, we use a "Minnesota Prior" (e.g., see Litterman, 1986). The idea behind this type of prior is to shrink parameters for persistent variables towards a random walk.

Accordingly, given that the variables in the VECM are included in first differences, we set the prior mean $\boldsymbol{\beta}_0$ in (B2) to a vector of zeros, except for the element associated with the error correction term in the short-rate equation. In that case, we set the prior mean to 0.5, consistent with the expectation hypothesis for the term structure of interest rates (see, for example, Modigliani and Shiller, 1973) that motivates our assumption of cointegration between the interest rates. Setting the prior mean for this parameter to zero is somewhat contrary to the assumption of cointegration. However, we note that our posterior inferences are robust to setting this prior mean to zero.

In specifying the prior variance, which dictates how tightly we shrink the coefficients towards zero, we follow the Minnesota prior approach and treat shorter lags as "more important" than longer lags when applying shrinkage. Let $V_{i,j}^k$ be the prior variance on the parameter in the i^{th} equation for the j^{th} variable on the k^{th} lag. Accordingly, we set

$$V_{i,j}^k = \frac{\lambda^2}{k^2}\frac{\sigma_i^2}{\sigma_j^2}. \tag{B6}$$

where σ_i^2 is the sample variance of the residuals from a univariate AR(4) regression fitted using least squares on the i^{th} variable and σ_i^2/σ_j^2 acts as a scaling factor to account for different units of the variables (note that we set $\sigma_j^2 = \sigma_i^2$ in the case of the error correction coefficients). The overall tightness of the prior is then governed by one hyperparameter, λ. We set $\lambda = 0.2$ in our empirical analysis, which is a fairly common choice within the BVAR literature (e.g., Sims and Zha, 1998) and corroborated as a reasonable choice in forecasting settings by Carriero et al. (2015). We stress, however, that our main results are robust to departures from this particular prior, including optimizing λ to minimize the one-step-ahead out-of-sample RMSFE for the short-run interest rate equation along the lines of Morley and Wong (2020).

Once $V_{i,j}^k$ in (B6) is specified, $\boldsymbol{V}_\beta$ is constructed as

$$
\boldsymbol{V}_{\boldsymbol{\beta}} = \begin{bmatrix} \boldsymbol{V}_1 & \boldsymbol{0} & \dots & \boldsymbol{0} \\ \boldsymbol{0} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \boldsymbol{0} \\ \boldsymbol{0} & \dots & \boldsymbol{0} & \boldsymbol{V}_n \end{bmatrix},
$$

where

$$
\boldsymbol{V}_i = \begin{bmatrix} \boldsymbol{V}_i^1 & \boldsymbol{0} & \dots & \boldsymbol{0} \\ \boldsymbol{0} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \boldsymbol{0} \\ \boldsymbol{0} & \dots & \boldsymbol{0} & \boldsymbol{V}_i^p \end{bmatrix} \text{ and } \boldsymbol{V}_i^k = \begin{bmatrix} V_{i,1}^k & \boldsymbol{0} & \dots & \boldsymbol{0} \\ \boldsymbol{0} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \boldsymbol{0} \\ \boldsymbol{0} & \dots & \boldsymbol{0} & V_{i,n}^k \end{bmatrix},
$$

except for $\boldsymbol{V}_1$ and $\boldsymbol{V}_2$, which each have an additional row and column of zeros and $V_{i,n+1}^1 = \lambda^2$ for $i = 1, 2$ on the respective last diagonal as the prior variance for the corresponding error correction coefficient.

For the remaining quantities in (B3), we set

$$
\boldsymbol{S} = \begin{bmatrix} \nu_0\sigma_1^2 & \boldsymbol{0} & \dots & \boldsymbol{0} \\ \boldsymbol{0} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \boldsymbol{0} \\ \boldsymbol{0} & \dots & \boldsymbol{0} & \nu_0\sigma_n^2 \end{bmatrix},
$$

where ν_0 is set to $n + 1$ (i.e., one greater than the total number of variables), σ_i^2 is obtained from the same AR(4) regression on the i^{th} variable as what used in (B6), and the variance is thus scaled up by a factor of ν_0 so that the prior on the sum of squared residuals is consistent with the prior on the degrees of freedom.

C Robustness checks

In this appendix, we consider some robustness checks to demonstrate that our main empirical findings are not particularly sensitive to choices about priors, proxies for inflation expectations, variables to include in the VECM, and possible permanent movements in the term premium.

First, we investigate the sensitivity of our r^* estimates with respect to the prior on the error-correction coefficient for the short-term interest rate and to the shrinkage hyperparameter used in the Bayesian estimation. The results are highly robust to the choice of prior, as seen in the first panel of Figure C1. Thus, while we see our baseline priors as well justified, we also note that our main findings, including the smoothness of our r^* estimates, do not hinge upon them.

Second, we investigate the sensitivity of our r^* estimates with respect to how we proxy inflation expectations when measuring *ex ante* real interest rates. As shown in the second panel of Figure C1, our estimates are generally robust to three alternative measures of inflation expectations. In the first case, following Laubach and Williams (2003), we proxy short-run inflation expectations with the forecast of the four-quarter-ahead percentage change in core PCE prices generated from a univariate AR(3) of inflation estimated over the prior 40 quarters (10 year rolling window). In the second case, we proxy the short-run (long-run) inflation expectations by the SPF short-term (long-run) inflation forecast. In the third case, we investigate the sensitivity of our r^* estimates with respect to the 1-month and 10-year real interest rates constructed by the Cleveland Fed based on their model-based expected inflation measures. The r^* estimate using the Cleveland Fed data appears to be higher in the early part of the sample for which it is available, but it soon converges to our baseline r^* estimate.

Third, to confirm the relevance of different sources of information, we consider two alternative models in terms of which variables are included. In the first case, we consider a smaller 7-variable model that, in addition to the 'interest rate' block, only includes the five most informationally-relevant variables for deviations of the short-term real interest rate from its trend following the variable selection process proposed in Morley and Wong (2020). The selected variables are the change in government debt, hours growth, employment growth, real consumption per capita growth, and stock returns. In the second case, we consider a bivariate model that only includes the 'interest rate' block. As shown in the third panel of Figure C1, the estimates are generally robust. Notably, by dropping the less informationally-relevant variables from the model, the estimated r^* is barely affected. The estimated r^* changes a bit more, however, when we do not include any possible determinants beyond the 'interest rate' block. But the general similarity of the estimates even when only including interest rates has two important implications: (i) measurement error or other sources of model misspecification that generate serial correlation in the first-stage estimates of trend growth appear to be primarily related to the *ex ante* real interest rates in particular and (ii) one could obtain a reasonably

Figure C1: Robustness of r^* estimates

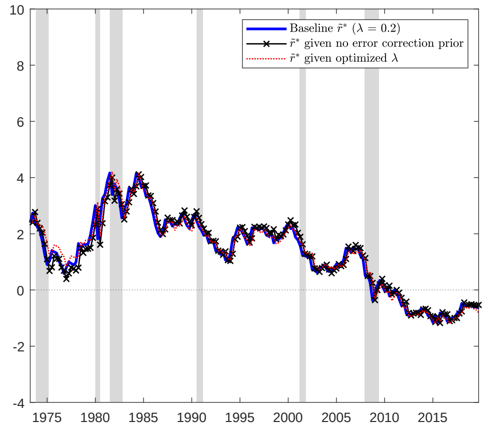

(a) Alternative priors

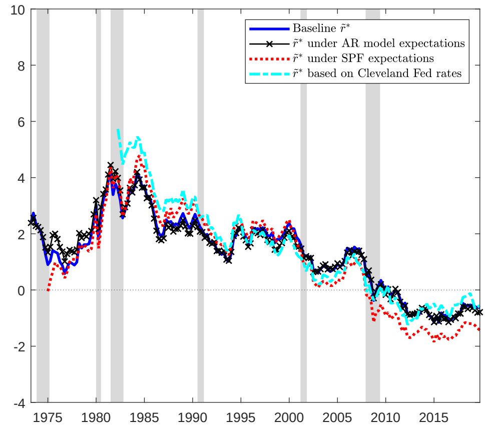

(b) Alternative real rates

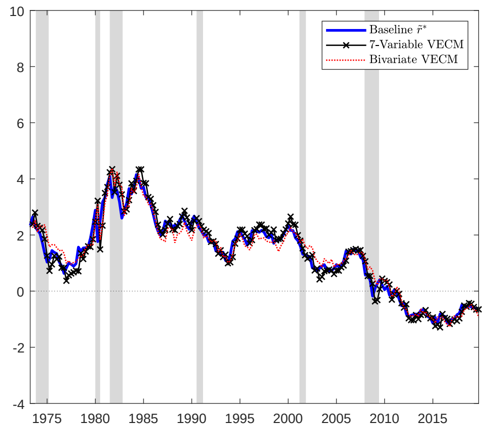

(c) Alternative information sets

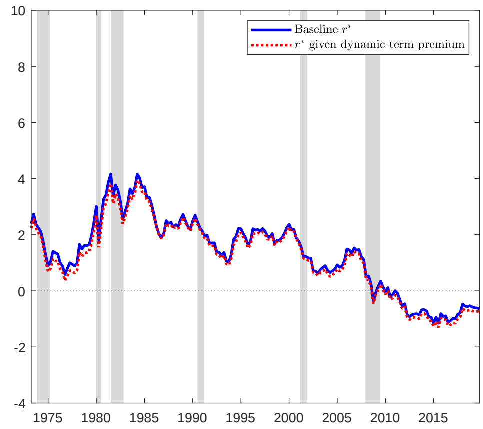

(d) Time-varying long-run term premium

Note: NBER recession dates are shaded.

robust estimate of r^* just by considering a bivariate VAR of the change in the short-term real interest rate and the spread between the long- and short-term real rates given that our correction can handle misspecification due to omitted variables. Of course, our baseline medium-scale Bayesian VECM has the advantage of allowing us to track which economic forces are most important in driving changes in r^*, which we consider in Section 4.4.

Fourth, we investigate the robustness of our result with respect to the assumption of no permanent movements in the term premium. We do so by allowing the long-run level of the term premium to be time varying, effectively by demeaning the error-correction term dynamically with a backward-looking 40-quarter average. As seen in the last panel of Figure C1, the results are robust.

D Additional exercises

In this appendix, we conduct some additional exercises related to the real-time reliability of our estimates, allowing for stochastic volatility, and taking a frequentist approach to inference given our proposed correction.

First, we consider real-time reliability of our estimates of r^*. To abstract from the effect of data revisions, which Orphanides and van Norden (2002) argue are less important than trend-cycle decomposition method for reliability of real-time estimates, we focus on r^* estimates based on a model using only the 'interest rate block' and using SPF survey measures of inflation expectations so that there are no sources of data revision. The first panel of Figure D1 plots the real-time estimate using an expanding window of data for the first ten years of the sample period until the end of the period to estimate r^* and compares it with the *ex post* estimate based on the full sample of data. The real-time estimates are clearly quite reliable, although there is an upward bias earlier in the sample period compared to the revised estimates. This is likely due to some changes in the estimated long-run level of the term premium over the sample period. But the movements and general decline in r^* implied by the real-time estimates are highly robust to consideration of the full sample of data. Notably, there is clearly no end-point problem that plagues other approaches to trend-cycle decomposition such as the Hodrick-Prescott filter.

Second, we consider an additional exercise of allowing for stochastic volatility when estimating the Bayesian VECM. Again, for tractability, we consider a bivariate VECM. We set up the bivariate VECM with stochastic volatility as per Carriero et al. (2022).[1] The MA parameters for the correction are estimated using standardized changes in the preliminary estimated trend, i.e., $\Delta\hat{r}_t^*/\sigma^2_{\Delta\hat{r}_t^*,t}$, where $\sigma^2_{\Delta\hat{r}_t^*,t} = \mathbf{s}'_{k,1}(\mathbf{I}-\mathbf{P})^{-1}\mathbf{H}\mathbf{\Sigma}_t\mathbf{H}'((\mathbf{I}-\mathbf{P})^{-1})'\mathbf{s}_{k,1}$, with $\mathbf{\Sigma}_t$ being the residual variance-covariance matrix for the VECM under stochastic volatility, while the correction is applied to the raw changes, i.e., $\Delta\hat{r}_t^*$. The second panel of Figure D1 plots the results when allowing for stochastic volatility, as well as the original bivariate VECM estimates. It can be seen that the estimates are largely robust.

Third, we sketch out how one might conduct frequentist inference about the corrected trend estimates if one were not inclined to do Bayesian estimation as we do. First, recall that

[1] The Carriero et al. (2022) approach uses the triangular factorization when it sets up the stochastic volatility component, and so we set up the model by ordering the *ex-ante* long-rate before the short-rate in the VECM. All other specifications remain identical to the baseline setting, and we also retain the prior from our baseline analysis for the VECM parameters.

Figure D1: Results for additional exercises

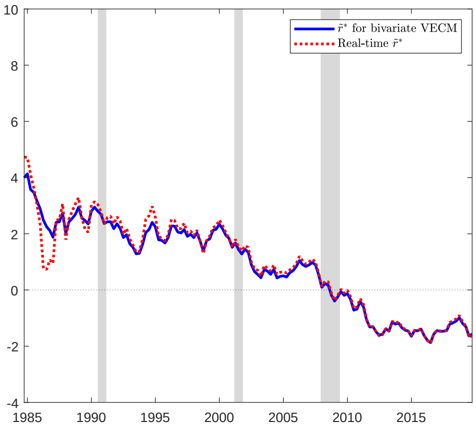

(a) Real-time reliability

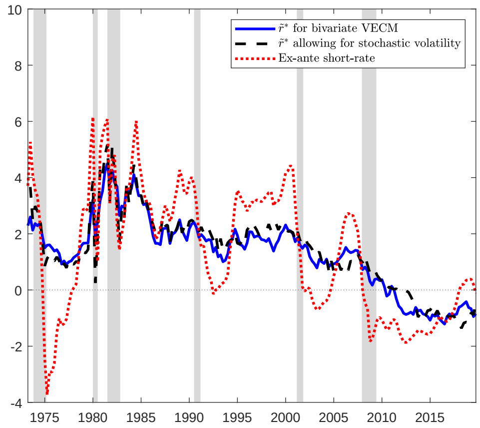

(b) Allowing for stochastic volatility

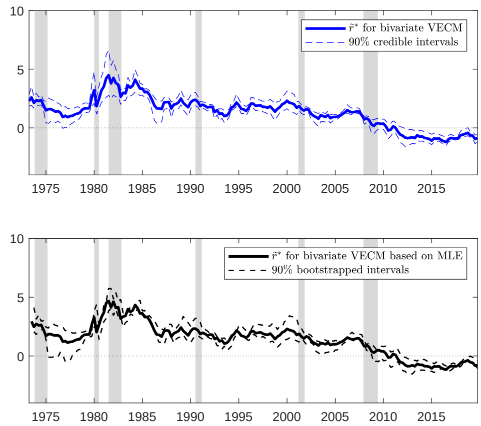

(c) Bayesian versus frequentist inference

Note: NBER recession dates are shaded.

our main motivation for conducting Bayesian inference in our baseline analysis is because we consider a medium-scale model and applying shrinkage by using Bayesian methods is reasonably standard to address in-sample overfitting given parameter proliferation in such settings (e.g., see Banbura et al., 2010). Nonetheless, if one considered the simpler bivariate model, it would be straightforward to consider OLS or MLE, as we also do in our Monte Carlo analysis with a bivariate DGP.

A bootstrap provides a natural approach to conducting frequentist inference about the corrected trend estimates given the complicated mapping from the preliminary estimates. For the bootstrap DGP, we obtain an estimate of the projection matrix $\hat{\mathbf{P}}$ and the projection errors $\hat{\boldsymbol{\eta}}_t$. Then, we propose the following four bootstrap steps:

1. Create an artificial sample of data $\mathbf{\Delta X}_t^{(b)}$ based on $\hat{\mathbf{P}}$ and (10) by drawing with replacement from the projection errors $\hat{\boldsymbol{\eta}}_t$ in a block bootstrap with block size of 5 to capture any small amount of serial correlation due to possible misspecification of the original forecasting model.

2. Obtain a bootstrap estimate of the projection matrix $\hat{\mathbf{P}}^{(b)}$ for the bootstrap sample $\mathbf{\Delta X}_t^{(b)}$.

3. Calculate a preliminary bootstrap estimate of trend $\hat{\mathbf{r}}^{*(b)}$ by using the bootstrap estimate of $\mathbf{P}^{(b)}$ and the original realized data $\mathbf{\Delta X}_t$.

4. Apply our proposed correction to $\hat{\mathbf{r}}^{*(b)}$ based on frequentist estimates of ARMA parameters for the first differences of the preliminary estimated trend. This step provides a bootstrap estimate $\tilde{\mathbf{r}}^{*(b)}$.

We repeat the above steps for 1000 bootstrap replications from $b = 1, \ldots, 1000$. We then take the $\alpha/2$ and $(100 - \alpha/2)$ quantiles across the bootstrapped replications for each $\tilde{r}_t^*$ within $\tilde{\mathbf{r}}^*$ and report these as $(100 - \alpha)\%$ bootstrapped confidence intervals.

It should be noted that our bootstrapped intervals only reflect parameter uncertainty. If we had infinite data and the population parameters, there would be no uncertainty about the corrected estimated changes in trend in (17) as we treat the data as realized values rather than random variables when making our calculation of the corrected estimated trend. In this sense, our bootstrapped intervals are confidence intervals rather than prediction intervals even though the BN decomposition is based on a long-horizon forecast. That is, the forecast is known at time t given population parameters, even though the realized future path is not. This is related also to the idea discussed in the main text that there is no filtering uncertainty about the BN trend, even if there might be about a true underlying r^* given a UC process. Again, it would be necessary to specify the UC process to capture any such filtering uncertainty.

The third panel of Figure D1 compares the Bayesian posterior mean and 90% equal-tailed credible intervals for $\tilde{r}_t^*$ against estimates based on MLE for the bivariate VECM and an MA(8) model for the correction with 90% bootstrapped confidence intervals. Even though there are the obvious conceptual differences, we note that the Bayesian and frequentist estimates and intervals are similar, which also confirms that the degree of uncertainty with our Bayesian inferences is not really due to informative priors.

E Variable-by-variable sign probabilities and informational contributions

In this appendix, we report on variable-by-variable sign probabilities and informational contributions.

Table E1 presents results for the individual variables in terms of (i) the correlation of the projection error for each variable and the implied change in the estimate of r^* based on the decomposition in (23) and (ii) their contributions to estimated r^* during the three subsample episodes considered in Figure 4.

In terms of the signs of correlations, Table E1 reports 62% and 71% posterior probabilities that consumption and TFP growth have a positive relationship with r^*. The broad finding of a positive link between trend growth and r^* corroborates many earlier studies (e.g., Laubach and Williams, 2003; Hamilton et al., 2016; Holston et al., 2017; Berger and Kempa, 2019; Lunsford and West, 2019). Investment growth has the predicted negative relationship consistent with investment-specific technological change with a 60% posterior probability. Consistent with the theoretical prediction on the effect of the labor force on r^* (Baker et al., 2005; Lunsford and West, 2019), there are 69% and 86% posterior probabilities that employment and hours growth have a positive relationship with r^*, while the unemployment rate has a negative relationship with 67% posterior probability, which is also consistent with a labor force effect or possibly insufficient aggregate demand, as suggested in Summers (2015). Meanwhile, consistent with a safe asset demand/flight-to-safety phenomenon, there are 66% and 83% posterior probabilities that macroeconomic uncertainty and the excess bond premium have a negative relationship with r^*. On the contrary, there is only weak evidence that liquid asset growth has a positive relationship with r^*, with only a 53% posterior probability, reflecting a likely mix of supply and demand factors driving this variable. Furthermore, consistent with the global savings glut hypothesis (Bernanke, 2005), there are 88% and 83% posterior probabilities that the current account and a depreciation in the exchange rate have respective positive and negative relationships with r^*. Last, we find that there is a 84% posterior probability of a positive relationship between debt-to-GDP and r^*, consistent with a safe asset supply/crowding-out effect (Ball and Mankiw, 1995).

Table E1: Accounting for changes in r^*

	Sign	Probability	Informational Contributions (bps)		
			Onset of Great Moderation	End of Great Moderation	Secular Stagnation
I. Productivity/Demographics			68 [2, 139]	-21 [-85, 43]	-43 [-116, 24]
Real consumption growth per capita	+	0.62	4 [-14, 23]	3 [-10, 18]	-7 [-38, 20]
TFP growth	+	0.71	2 [-11, 14]	20 [-14, 56]	-18 [-51, 11]
S&P 500 stock returns	+	0.74	6 [-11, 23]	14 [-7, 38]	-8 [-28, 12]
Real investment growth	-	0.60	0 [-10, 9]	-8 [-34, 17]	7 [-17, 32]
Employment growth	+	0.69	8 [-20, 41]	-16 [-49, 14]	-8 [-42, 22]
Hours growth	+	0.86	38 [1, 78]	-27 [-55, 0]	-10 [-30, 8]
Unemployment rate (Δ)	-	0.67	11 [-9, 34]	-8 [-29, 10]	-1 [-17, 18]
II. Safe Asset Demand			35 [-58, 129]	0 [-46, 47]	-39 [-133, 53]
Macroeconomic uncertainty (Δ)	-	0.66	9 [-13, 33]	1 [-7, 10]	-10 [-38, 15]
Excess bond premium (Δ)	-	0.83	0 [-13, 12]	29 [-1, 60]	-31 [-67, 1]
Liquid assets growth	-	0.53	-1 [-89, 86]	1 [-20, 23]	0 [-75, 77]
Current account as % of GDP (Δ)	+	0.88	15 [-3, 34]	-34 [-65, -4]	16 [-4, 38]
Exchange rate return	-	0.80	12 [-6, 33]	3 [-8, 15]	-13 [-35, 7]
III. Safe Asset Supply			52 [-1, 112]	-67 [-140, 1]	29 [-5, 67]
Government debt as % of GDP (Δ)	+	0.84	52 [-1, 112]	-67 [-140, 1]	29 [-5, 67]
Changes in r^*			139 [50, 232]	-141 [-192, -90]	-213 [-244, -180]

Notes: The posterior probability of each reported sign is reported. For the informational contributions, the posterior mean is reported with 67% equal-tailed credible intervals reported in square brackets.

In terms of contributions over the subsamples considered in Figure 4, Table E1 suggests that higher employment and hours growth helped drive the large overall contribution of productivity/demographic factors to the rise in r^* with the onset of the Great Moderation. The individual variables associated with safe asset demand had somewhat offsetting effects with the onset of the Great Moderation, while higher safe asset supply in the form of an increase in government debt-to-GDP during the Reagan years had a clear positive contribution to r^*. The effects of the key individual variables during the onset of the Great Moderation reversed by the end of the Great Moderation, especially with the debt consolidation during the Clinton years, although faster TFP growth and higher stock returns with the so-called 'New Economy' at the time meant the overall drag from productivity/demographic factors was less than otherwise, while the individual safe asset demand variables had largely offsetting effects, with a large positive effect from a lower excess bond premium and large negative effect from a current account deficit due to large capital inflows to the United States related to high savings rates in emerging market economies, especially after the Asian financial crisis and with high revenues earned by oil exporters from booming oil prices (Glick, 2020). Finally, with Secular Stagnation, lower trend growth captured by lower consumption growth, TFP growth, and weaker stock returns, as well as weaker employment and hours growth, all contributed to the fall in r^*, as did the key safe asset demand related variables of macroeconomic uncertainty and the excess bond premium, although the other safe asset demand related variables mostly had offsetting effects, as did the increase in safe asset supply with a higher debt-to-GDP ratio again.

F Estimating the model using annual data

In this appendix, we re-estimate the model using annual data as some of the possible drivers of r^* are only available at an annual frequency, specifically income inequality, age dependency and global reserves-to-GDP. The final dataset starts from 1975 after transformation and also includes all of the variables used in the baseline estimation converted to annual measures.

The r^* estimated using annual data is notably smoother than when using quarterly data as seen in Figure F1. However, this does not reflect the inclusion of the variables only available at an annual frequency, as they appear to contribute only negligibly to movements in the estimated r^*. Instead, the smoothness could reflect possible overfitting in-sample given small sample period of annual data and many parameters, although we only consider one lag for the annual VECM and two lags for the MA model used for the correction.

Figure F1: Corrected BN trend for the short-term real interest rate using annual data

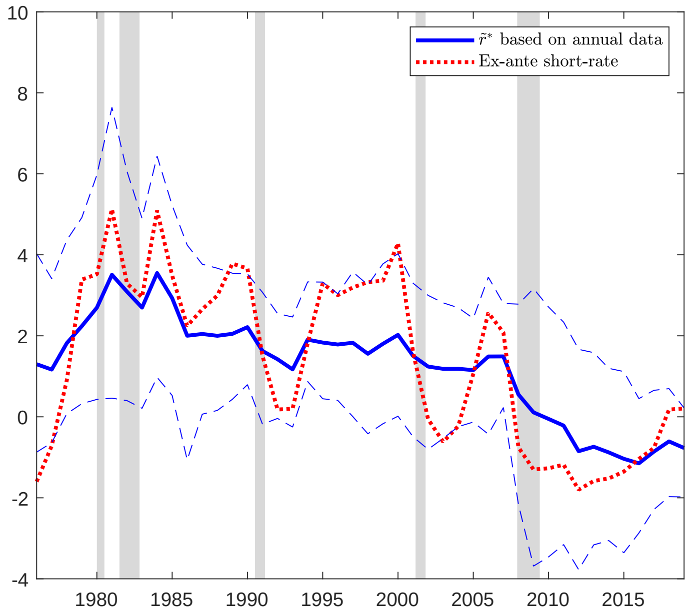

Note: Posterior mean and 90% equal-tailed credible intervals (dashed lines) are reported. NBER recession dates are shaded.

As seen in Table E2, we also find that the informational contributions are roughly similar, but less precise than for the baseline quarterly model. In many cases, variables that were important in the quarterly case have insignificant contributions and sometimes have higher posterior probabilities on the wrong sign in terms of predicted theoretical relationships when considering the annual model. Importantly, because the variables that are only available at an annual frequency do not appear to contribute significantly, we can infer that their omission from the baseline quarterly model does not seem to be a source of distortion of our inferences about informational contributions for our baseline model. In any event, our correction can help address misspecification due to omitted variables in the quarterly model.

Table E2: Accounting for changes in r^* based on annual data

	Sign	Probability	Informational Contributions (bps)		
			Onset of Great Moderation	End of Great Moderation	Secular Stagnation
I. Productivity/Demographics			22 [-60, 122]	-10 [-109, 83]	-40 [-170, 68]
Real consumption growth per capita	+	0.59	3 [-21, 25]	2 [-16, 19]	-10 [-51, 29]
TFP growth	+	0.61	-1 [-20, 19]	8 [-18, 38]	-15 [-65, 31]
S&P 500 stock returns	-	0.52	0 [-17, 18]	-1 [-18, 16]	1 [-24, 27]
Real investment growth	+	0.62	4 [-19, 28]	1 [-16, 20]	-11 [-62, 33]
Employment growth	+	0.63	14 [-25, 54]	-10 [-38, 17]	-12 [-53, 25]
Hours growth	+	0.62	8 [-29, 52]	-7 [-43, 24]	-7 [-43, 23]
Unemployment rate (Δ)	-	0.56	5 [-22, 33]	-3 [-27, 20]	-4 [-32, 23]
Age Dependency (Δ^2)	-	0.53	2 [-23, 26]	8 [-82, 99]	-6 [-76, 65]
Inequality (Δ)	-	0.76	-8 [-30, 13]	-9 [-30, 9]	15 [-11, 45]
II. Safe Asset Demand			-29 [-156, 80]	22 [-45, 103]	14 [-101, 139]
Macroeconomic uncertainty (Δ)	+	0.67	-7 [-27, 12]	-1 [-15, 13]	16 [-14, 46]
Excess bond premium (Δ)	-	0.63	1 [-15, 19]	5 [-13, 25]	-10 [-47, 22]
Liquid assets growth	-	0.57	-17 [-127, 92]	7 [-37, 55]	16 [-77, 111]
Current account as % of GDP (Δ)	-	0.51	-1 [-17, 16]	0 [-17, 19]	1 [-28, 28]
Exchange rate return	-	0.78	12 [-14, 38]	0 [-20, 19]	-7 [-31, 19]
Reserves-to-GDP (Δ)	+	0.80	-24 [-60, 8]	16 [-8, 44]	5 [-19, 32]
III. Safe Asset Supply			0 [-25, 29]	-3 [-64, 53]	1 [-41, 47]
Government debt as % of GDP (Δ)	+	0.53	0 [-25, 29]	-3 [-64, 53]	1 [-41, 47]
Changes in r^*			89 [-20, 210]	-195 [-109, -26]	-289 [-229, -167]

Notes: The posterior probability of each reported sign is reported. For the informational contributions, the posterior mean is reported with 67% equal-tailed credible intervals reported in square brackets.

G How did r^* change during the COVID-19 pandemic?

In this appendix, we extend our analysis to cover the onset of the COVID-19 pandemic. In particular, we update the dataset to 2020Q4, but we use the pre-Covid parameter estimates to avoid possible distortions from large outliers in the data.[2] The first panel of Figure G1 plots the posterior mean of the corrected BN trend for the short-term real interest rate over the latter part of the sample period and up to the end of 2020, noting the pre-2020 estimates of r^* are the same as in Figure 2 given the same parameter estimates. With the onset of the pandemic, the estimated r^* falls sharply to about -2.5% as various indicators related to the marginal product of capital adjusted dramatically and there was a jump up in macroeconomic uncertainty. However, the persistence of these variables was very different than normal given the unusual stop-start nature of economic activity with lockdowns, as well as the unprecedented fiscal stimulus, and the estimated r^* quickly jumped back up slightly above its pre-pandemic level. Looking at the various quarterly variables, we find that demand for safe assets was still a drag on r^* by the end of 2020, contributing an estimated 90 basis point decrease over the year, while supply of safe assets in the form of higher debt-to-GDP mostly offset this effect by contributing an estimated 75 basis point increase over the same period. These contributions are plotted in the second panel of Figure G1.

[2] Lenza and Primiceri (2022) note that excluding data from 2020 when estimating a VAR is a simple approximation to a GLS-type approach, although their main proposal is to model a rescaling of the residual variance-covariance matrix during 2020 when constructing density forecasts during this period. Because we are interested in point forecasts, we take the simpler approach of not including the data from 2020 in parameter estimation, which is also consistent with the findings in Schorfheide and Song (2020) that forecasts based on VAR parameters estimated using only data before the pandemic appear "more stable and reasonable" than those based on updated parameter estimates.

Figure G1: r^* during the pandemic

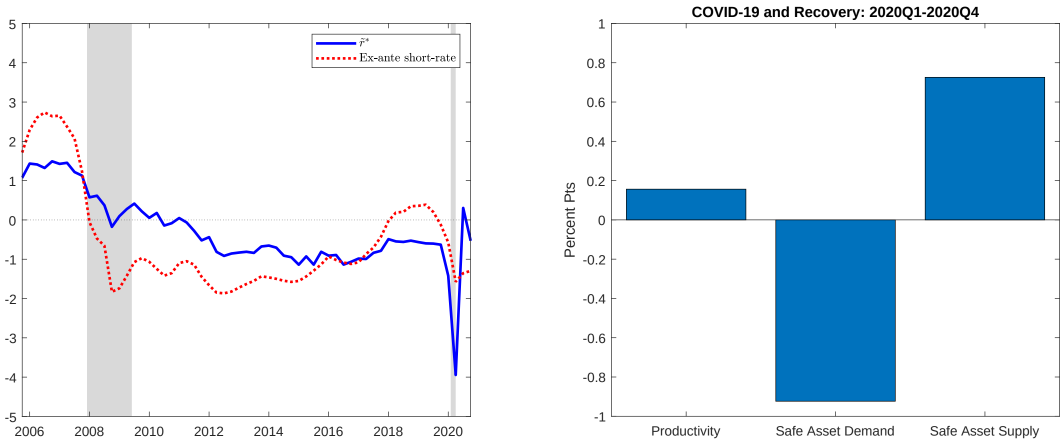

(a) Corrected BN trend including 2020 (b) Informational contributions in 2020

Note: In panel (a), posterior mean is reported and NBER recession dates are shaded. In panel (b), posterior mean informational contributions are reported.