

The decline in r^* according to a robust multivariate trend-cycle decomposition*

James Morley[†] Trung Duc Tran[‡] Benjamin Wong[§]

December 2021, *Preliminary Draft*

Abstract

Interest rates have fallen worldwide in recent decades, a phenomenon that has been linked at least in part to a decline in the natural rate of interest, r^* (a.k.a. “*r-star*”). To investigate this decline, we consider a multivariate trend-cycle decomposition of real interest rates using a large set of variables hypothesized to explain changes in r^* . We develop a robust two-step procedure to address apparent model misspecification that could be due to measurement error in constructed *ex ante* real interest rates or other variables used in the multivariate trend-cycle decomposition. Our estimates suggest a smooth path for r^* without imposing it *a priori*, with the decline since the Great Recession in particular explained more by slower trend growth than increased demand for safe assets, the effect of which appears to be almost completely offset by higher levels of government debt.

JEL codes: C32, E43, E44

Keywords: Natural rate of interest, multivariate trend-cycle decomposition

*This research was supported by ARC DP190100202 and DE200100693. We thank Todd Clark, Yunjong Eo, Leland Farmer, Greg Kaplan, Mengheng Li, Elmar Mertens, Tatsuyoshi Okimoto, Tatevik Sekhposyan, Jake Wong, and conference and seminar participants at the International Association for Applied Econometrics 2021 Annual Conference (virtual), Applied Time Series Econometrics workshop hosted by the Federal Reserve Bank of St. Louis (virtual), the Macroeconometrics and Financial Econometrics workshop hosted by Hitotsubashi University (virtual), the Sydney Macro Reading Group workshop (Sydney), the University of Adelaide (virtual), and the University of Tasmania (virtual) for helpful comments. The usual disclaimers apply. The views expressed in this paper are those of the authors and should not be attributed to the Reserve Bank of Australia.

[†]University of Sydney. Email: james.morley@sydney.edu.au

[‡]Reserve Bank of Australia. Email: tranb@rba.gov.au

[§]Monash University. Email: benjamin.wong@monash.edu

1 Introduction

In recent years, interest rates have been persistently low and many studies have linked this to an underlying decline in the long-run “equilibrium” or “natural” real rate of interest, often referred to as r^* (or “*r-star*”); see, for example, [Del Negro et al. \(2017\)](#); [Hamilton et al. \(2016\)](#); [Holston et al. \(2017\)](#); [Lubik and Matthes \(2015\)](#); [Bauer and Rudebusch \(2020\)](#); [Kiley \(2020\)](#); [Johannsen and Mertens \(2021\)](#); [Cúrdia et al. \(2015\)](#); [Lewis and Vazquez-Grande \(2019\)](#). Notwithstanding differences regarding its formal definition, the main debate has been as to why and how much r^* has fallen. One prominent view, inspired by the semi-structural model in [Laubach and Williams \(2003\)](#), is that the decline is primarily a supply-side phenomenon associated with lower trend 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 [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 long periods of time (see, e.g., [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 been persistently below zero. The answer to this question has important implications for gauging the stance of monetary policy, including during the COVID-19 pandemic.

We define r^* as the common stochastic trend for any set of real interest rates that includes a risk-free short-term rate.¹ Then, to investigate why and how much r^* has fallen, we consider a multivariate version of the [Beveridge and Nelson \(1981\)](#) (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. Our model is a Bayesian VECM that assumes cointegration between short- and long-term real interest rates. In addition to

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

accommodating a large set of variables via a shrinkage prior, the Bayesian approach also allows us to incorporate an “expectations hypothesis” prior consistent with cointegration such that the short rate is assumed to adjust to the long rate, although our main results are robust to relaxing this prior.

Working with *ex ante* real interest rates based on U.S. data makes it viable to model our large system of variables in a linear environment.² However, 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.³ We find evidence of model misspecification with our trend-cycle decomposition that could be due to a small amount of measurement error that is hard to detect directly or to capture with an extended model. To address this, we propose a robust two-step BN decomposition based on the law of iterated expectations to correct for apparent model misspecification. Our two-step procedure is easy to apply and leads to a smooth estimated path for r^* without imposing smoothness *a priori*, as is often done when estimating r^* in the previous literature.⁴

²Modeling nominal interest rates and inflation separately in a linear environment would be problematic because nominal interest rates are subject to the nonlinear constraint of the effective or zero lower bound (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, cancel out when considering real interest rates. Meanwhile, to the extent that the ZLB 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, however, our results are robust to allowing for a possible structural break in the term premium when considering a sample period that includes the ZLB. Put simply, our linear VECM appears to be a much better dynamic model to capture a common stochastic trend in real interest rates than of nominal interest rates and inflation separately. However, see [González-Astudillo and Laforte \(2020\)](#) and [Johannsen and Mertens \(2021\)](#) for approaches that model nominal interest rates and inflations separately and attempt to account for the nonlinearities created by the ZLB.

³Even given real bonds, 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. 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 note in the previous footnote. 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 current analysis of the stance of monetary policy.

⁴See [Del Negro et al. \(2017\)](#) for an example of a smoothness prior being imposed to produce a smooth estimated path of r^* . Also, see [Kiley \(2020\)](#) for a discussion of how the smoothness of r^* is not well-identified for the [Laubach and Williams \(2003\)](#) semi-structural model. The smoothness we find is due to being able to capture the reduced-form serial correlation properties of real interest rates with our VECM that includes a large amount of multivariate information and the two-step correction for potential model

Although we consider an “expectations hypothesis” prior, 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 growth, as done in [Laubach and Williams \(2003\)](#), but instead we test this link while including a much larger set of other potential explanatory variables than considered in typical structural or semi-structural models (although see [Fu, 2020](#)). In principle, we could incorporate DSGE-based Granger-causation priors in the Bayesian VECM if we wanted to impose more structure when estimating r^* . However, our agnostic priors allow us to treat posterior inferences about signs of effects as statistical tests of theoretical predictions, while our inferences about r^* are precise without having to rely on potentially restrictive structural assumptions.

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.⁵ Because some hypothesized variables are only available at an annual frequency, including some that [Lunsford and West \(2019\)](#) identify as being important, we also consider cointegration analysis in terms of annual variables related to income inequality, age dependency, and global reserves-to-GDP and movements in r^* that are not explained by our quarterly variables.

Our main empirical results can be summarized as follows: First, we find that, given equal prior odds, the posterior probabilities suggest that all quarterly variables have their theoretically predicted effects on r^* over the full sample period from 1973-2019. Second,

 misspecification.

⁵Our approach also provides a complement to [Rachel and Smith \(2017\)](#), who applied a simple accounting framework based on reduced-form elasticity estimates, and [Rachel and Summers \(2019\)](#), who combined evidence from several structural models, in terms of accommodating a wide range of explanations for why interest rates are low.

r^* was generally above 2% during the Great Moderation from 1982-2005 due in part to faster trend growth and persistently lower levels of risk and volatility at the time. Third, r^* fell with the Great Recession and during its aftermath due to slower trend growth and increased global demand for safe assets, the effect of which is almost completely offset by higher levels of government debt. Notably, the estimated r^* took on persistently negative values from mid 2011, falling to about -1% from mid 2015, although zero almost always lies within the 95% credible intervals when the point estimate is negative. Supply-side variables related to productivity and demographics accounted for an estimated 41 basis point drop since 2006, while variables related to safe assets demand accounted for an estimated 26 basis point drop, but a rising government debt-to-GDP ratio accounts for an offsetting 25 basis point increase. Fourth, the cointegration analysis with annual data suggests income inequality and age dependency do not play an obvious role in driving r^* , but global reserves-to-GDP ratio is cointegrated with the residual component of r^* that is not explained by quarterly variables and, consistent with the global savings glut hypothesis of [Bernanke \(2005\)](#), contributes about 30 basis points to the decline in r^* since the late 1990s.

We highlight that our estimated r^* can be used to help gauge the stance of monetary policy, including during the COVID-19 pandemic. Extending our estimation of r^* to 2020, we find a sharp decrease in the estimated r^* of a similar scale to the decline in the *ex ante* short-term real interest rate with the onset of the crisis, but an immediate recovery back to a similar level around -1% as before the crisis, which is, notably, back above the *ex ante* short-term real interest rate. The behavior of supply-side variables accounted for the quick reversal, while higher government debt again offset increased demand for safe assets, as it did following the Great Recession. These estimates imply that monetary policy has been at least somewhat accommodative during the pandemic despite the ZLB, but not nearly as much as it would have been if r^* were closer to historical levels of 2% instead of -1%.

The rest of the paper is organized as follows. Section 2 presents our robust two-step BN decomposition and shows how it can correct for model misspecification such as could

be due to even a small amount of measurement error. Section 3 describes the data and the VECM estimation. Section 4 reports our empirical results, including an extension out of sample to consider the behavior of r^* during the COVID-19 pandemic. Section 5 concludes.

2 A robust two-step Beveridge-Nelson trend-cycle decomposition that addresses model misspecification

As described in this section, we develop a robust two-step version of the [Beveridge and Nelson \(1981\)](#) trend-cycle decomposition that can be applied when there appears to be model misspecification due to measurement error or possibly other sources such as omitted variables. The BN decomposition has proven a useful approach to separate trend from cycle in a wide variety of settings (see, for example, [Evans and Reichlin, 1994](#); [Morley and Piger, 2012](#); [Kamber et al., 2018](#); [Kamber and Wong, 2020](#)). Our proposed two-step procedure to address model misspecification is the first to our knowledge for the BN decomposition.

2.1 General framework for the BN decomposition

The key assumption motivating our use of the BN decomposition in the first place is that the natural rate of interest r_t^* corresponds to the driftless random walk trend component of a risk-free short-term real interest rate r_t , with cyclical deviations from trend having an unconditional mean of zero (i.e., $\mathbb{E}[r_t^c] = 0$, where $r_t^c \equiv r_t - r_t^*$).⁶ Under this assumption, the BN decomposition calculates an estimate of trend as follows:

$$\mathbb{E}[r_t^*|\Omega_t] = \lim_{h \rightarrow \infty} \mathbb{E}[r_{t+h}|\Omega_t], \quad (1)$$

⁶In 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 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 r_t . This point is related to a more general one motivating the use of the BN decomposition highlighted by [Rotemberg and Woodford \(1996\)](#) and [Kiley \(2013\)](#), amongst others. We note that the random walk assumption is standard in the empirical literature on r_t^* including [Laubach and Williams \(2003\)](#) and [Del Negro et al. \(2017\)](#).

where $\Omega_t = \{\mathbf{x}_t, \mathbf{x}_{t-1}, \dots, \mathbf{x}_1; \mathbf{f}(\{\mathbf{x}_t\}_{-\infty}^{+\infty})\}$ includes all relevant information at time t for calculating the long-horizon expectation given an assumed data generating process $\mathbf{f}(\{\mathbf{x}_t\}_{-\infty}^{+\infty})$, where \mathbf{x}_t is an $n \times 1$ vector that includes the target variable r_t , which is assumed to be in the first row for convenience. The logic of the BN decomposition is that, because the long-horizon conditional expectation of the cyclical deviation from trend is zero, the long-horizon conditional 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.

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

$$\Delta \mathbf{X}_t = \mathbf{F} \Delta \mathbf{X}_{t-1} + \mathbf{H} \mathbf{e}_t, \quad (2)$$

where $\Delta \mathbf{X}_t$ is a $k \times 1$ vector of stationary demeaned variables with $\Delta \mathbf{x}_t - \boldsymbol{\mu}$ in the first n rows, \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 $\Delta \mathbf{x}_t$, 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 multivariate models such as VARs and VECMs (see, for example, [Morley, 2002](#)). We also note that such forecasting models provide reduced-form representations of many dynamic structural models, with the BN decomposition producing a robust estimate of the random walk trend for a 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 estimates of trend and cycle for r_t given the forecasting model for $\Delta \mathbf{x}_t$ can be calculated as

$$\mathbb{E}[r_t^* | \Omega_t] = r_t + \mathbf{s}'_{k,1} \mathbf{F}(\mathbf{I} - \mathbf{F})^{-1} \Delta \mathbf{X}_t, \quad (3)$$

$$\mathbb{E} [r_t^c | \Omega_t] = -\mathbf{s}'_{k,1} \mathbf{F}(\mathbf{I} - \mathbf{F})^{-1} \Delta \mathbf{X}_t, \quad (4)$$

where we let $\mathbf{s}_{k,1}$ denote a $k \times 1$ selection vector that contains zeros in all rows except for a 1 in the j^{th} row. Note that, given the forecasting model in (2), the relevant information set can be simplified to $\Omega_t = \{\mathbf{X}_t, \mathbf{X}_{t-1}; \mathbf{F}\}$.

A key feature of the BN estimate of trend is that changes in it should inherit the same lack of serial correlation as the changes in true random walk trend. In particular, following [Morley and Wong \(2020\)](#), the change in the BN estimate of trend is

$$\mathbb{E} [r_t^* | \Omega_t] - \mathbb{E} [r_{t-1}^* | \Omega_{t-1}] = \mathbf{s}'_{k,1} (\mathbf{I} - \mathbf{F})^{-1} \mathbf{H} \mathbf{e}_t = \sum_{i=1}^n \omega_i e_{it}, \quad (5)$$

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 (5) will itself be serially uncorrelated. However, if, in practice, changes in a BN estimate of trend display nontrivial serial correlation, we note this could reflect model misspecification due to the presence of measurement error in some variables in the model or other sources such as omitted variables. Next, we consider what happens if there is such model misspecification, focusing on the case of measurement error, and propose a two-step procedure to address it.

2.2 Measurement error and a robust two-step procedure

To consider the effects of measurement error, let $\tilde{\mathbf{x}}_t = \mathbf{x}_t + \mathbf{u}_t$ 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 $\mathbb{E} [\mathbf{u}_t] = \mathbf{0}$.⁷ For example, it is quite likely that our measure of the real interest rate contains measurement error as we have to proxy for inflation expectations when constructing an *ex ante* real interest rate. Thus, in practice, we only observe $\tilde{r}_t = r_t + u_{1t}$. As we will show, such measurement error could imply serial correlation in the changes in the BN estimate of trend for \tilde{r}_t due to model misspecification.

⁷It is worth noting that the BN decomposition approach even 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., $\mathbf{x}_t = \tilde{\mathbf{x}}_t - \mathbf{u}_t$) are block exogenous with respect to the measurement error.

However, assuming this implied serial correlation can be captured by an ARMA model, we also show how to correct for it.

To see how we propose correcting for the effects of measurement error, first consider the misspecified forecasting model applied to the observed data with measurement error:

$$\Delta\tilde{\mathbf{X}}_t = \mathbf{F}\Delta\tilde{\mathbf{X}}_{t-1} + \mathbf{H}\check{\mathbf{e}}_t, \quad (6)$$

where $\Delta\tilde{\mathbf{X}}_t = \Delta\mathbf{X}_t + \mathbf{H}\Delta\mathbf{u}_t$ and $\check{\mathbf{e}}_t = \mathbf{e}_t + \Delta\mathbf{u}_t - \mathbf{S}'\mathbf{F}\mathbf{H}\Delta\mathbf{u}_{t-1}$ given the $k \times n$ selection matrix $\mathbf{S} = (\mathbf{s}'_{k,1}, \mathbf{s}'_{k,2}, \dots, \mathbf{s}'_{k,n})'$. Note that the $\Delta\mathbf{u}_t - \mathbf{S}'\mathbf{F}\mathbf{H}\Delta\mathbf{u}_{t-1}$ term means elements of $\check{\mathbf{e}}_t$ will be serially correlated even if the measurement error \mathbf{u}_t is serially uncorrelated. As a result, estimates of \mathbf{F} based on the false assumption that $\check{\mathbf{e}}_t$ is serially uncorrelated will generally not converge to the true value \mathbf{F} but instead to the population projection of $\Delta\tilde{\mathbf{X}}_t$ on $\Delta\tilde{\mathbf{X}}_{t-1}$, which we denote as \mathbf{P} .⁸

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

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

where $\tilde{r}_t^c \equiv -\mathbf{s}'_{k,1}\mathbf{P}(\mathbf{I}-\mathbf{P})^{-1}\Delta\tilde{\mathbf{X}}_t$ is the first-stage BN estimate of the cycle for \tilde{r}_t given the misspecified forecasting model for $\Delta\tilde{\mathbf{x}}_t$. Then, we can construct a “BN-cycle-adjusted” measure of the real interest rate:

$$\tilde{r}_t \equiv \tilde{r}_t - \tilde{r}_t^c = \mathbb{E}[r_t^* | \Omega_t] + \mathbb{E}[r_t^c | \Omega_t] - \check{r}_t^c + u_{1t}. \quad (7)$$

⁸The difference between the population projection \mathbf{P} of $\Delta\tilde{\mathbf{X}}_t$ on its lag and \mathbf{F} is analogous to how the population 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.

Note that both BN cycle estimates for r_t (hypothetical, if data were observed without measurement error) and \tilde{r}_t (actual, given observed data with measurement error) will be stationary and mean zero by construction and the measurement error is $I(0)$ and mean zero by assumption. Thus, the key point is that the random walk component of BN-cycle-adjusted \tilde{r}_t is the BN estimate of trend $\mathbb{E}[r_t^*|\Omega_t]$. As a result, we can apply a BN decomposition to \tilde{r}_t as a second step following the first-stage estimation of \tilde{r}_t^c from \tilde{r}_t in order to estimate the trend for r_t :

$$\mathbb{E}[\mathbb{E}[r_t^*|\Omega_t]|\tilde{\Omega}_t] = \mathbb{E}[r_t^*|\tilde{\Omega}_t] = \lim_{h \rightarrow \infty} \mathbb{E}[\tilde{r}_{t+h}|\tilde{\Omega}_t], \quad (8)$$

where $\tilde{\Omega}_t = \{\tilde{\mathbf{X}}_t, \tilde{\mathbf{X}}_{t-1}; \mathbf{P}, \phi(L), \theta(L)\}$, with $\phi(L)$ and $\theta(L)$ corresponding to lag polynomials from an ARMA forecasting model for $\Delta\tilde{r}_t$.⁹ As can be seen from the law of iterated expectations, the BN estimate of trend for \tilde{r}_t is equivalent to a BN estimate of trend for r_t even if we do not directly observe r_t due to measurement error.

We note that the presence of measurement error will generally imply a complicated ARMA structure for the serial correlation in $\Delta\tilde{r}_t$. For example, consider a VAR(p) structure for $\Delta\mathbf{x}_t$ and assume that at least some of the variables including the real interest rate are measured with MA(q) error – i.e., $u_{1t} \sim \text{MA}(q)$ and $u_{jt} \sim \text{MA}(\leq q)$ for $j > 1$. Then, $\Delta\tilde{\mathbf{x}}_t$ 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 \tilde{r}_t :

$$\mathbb{E}[r_t^c|\Omega_t] = -\mathbf{s}'_{k,1}\mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}\Delta\mathbf{X}_t \sim \text{ARMA}(\leq pn, \leq pn - 1), \quad (9)$$

$$\tilde{r}_t^c = -\mathbf{s}'_{k,1}\mathbf{P}(\mathbf{I} - \mathbf{P})^{-1}\Delta\tilde{\mathbf{X}}_t \sim \text{ARMA}(\leq pn, \leq pn + q + 1), \quad (10)$$

where the orders follow from an ability to re-write the terms in equations (9) and (10) in terms of the VAR(p) as $\mathbf{s}'_{n,1}\mathbf{c}_t$ and $\mathbf{s}'_{n,1}\check{\mathbf{c}}_t$, with $\mathbf{c}_t \equiv -\mathbf{\Gamma}_{\mathbf{F},1}(\Delta\mathbf{x}_t - \boldsymbol{\mu}) - \dots - \mathbf{\Gamma}_{\mathbf{F},p}(\Delta\mathbf{x}_{t-p+1} - \boldsymbol{\mu})$, $\check{\mathbf{c}}_t \equiv -\check{\mathbf{\Gamma}}_{\mathbf{P},1}\Delta(\tilde{\mathbf{x}}_t - \boldsymbol{\mu}) - \dots - \check{\mathbf{\Gamma}}_{\mathbf{P},p}(\Delta\tilde{\mathbf{x}}_{t-p+1} - \boldsymbol{\mu})$, and $\mathbf{\Gamma}_{\mathbf{F},j}$ and $\mathbf{\Gamma}_{\mathbf{P},j}$ 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

⁹Alternatively, we could consider an unobserved components (UC) model of \tilde{r}_t , but we note that there could be correlation between permanent and transitory movements in \tilde{r}_t , so a BN decomposition will be appropriate even if this correlation is not identified for a UC model.

$\Delta \mathbf{x}_t \sim \text{VAR}(p)$ and $\Delta \check{\mathbf{x}}_t \sim \text{VARMA}(p, q + 2)$, \mathbf{c}_t will be $\text{VARMA}(p, p - 1)$ and $\check{\mathbf{c}}_t$ will be $\text{VARMA}(p, p + q + 1)$, 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 parameterizations, the point is that these are highly complicated processes and imply the following even more complicated process for $\Delta \tilde{r}_t$ when combined with the other terms in \tilde{r}_t and first differences are taken:

$$\Delta \tilde{r}_t \sim \text{ARMA}(\leq 2pn, \leq 2pn + q + 2). \quad (11)$$

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 $\check{\mathbf{e}}_t$ if the measurement error is relatively small compared to the forecast errors in \mathbf{e}_t . To see the issue, consider a case where the measurement error is small enough (i.e., $\text{var}(\Delta \mathbf{u}_t) \ll \text{var}(\Delta \mathbf{x}_t)$) such that $\mathbf{P} \approx \mathbf{F}$. In this case, we get the following expression for $\Delta \tilde{r}_t$:

$$\begin{aligned} \Delta \tilde{r}_t &\approx \mathbf{s}'_{k,1} (\mathbf{I} - \mathbf{F})^{-1} \mathbf{H} \mathbf{e}_t + \mathbf{s}'_{k,1} \mathbf{F} (\mathbf{I} - \mathbf{F})^{-1} \mathbf{H} \Delta \mathbf{u}_t - \mathbf{s}'_{k,1} \mathbf{F} (\mathbf{I} - \mathbf{F})^{-1} \mathbf{H} \Delta \mathbf{u}_{t-1} + \Delta u_{1t} \\ &= \sum_{i=1}^n \omega_i e_{it} + \sum_{i=1}^n \gamma_i (\Delta u_{it} - \Delta u_{it-1}) + \Delta u_{1t}, \end{aligned} \quad (12)$$

where we let γ_i denote a weight that is equal to the i^{th} 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 $\Delta \tilde{r}_t$ will have an $\text{MA}(q + 2)$ structure given $\text{MA}(q)$ measurement error and the effects of measurement error may be easier to detect than in $\Delta \tilde{r}_t = \Delta r_t + \Delta u_{1t}$ given a very different signal-to-noise ratio.¹⁰ In particular, if r_t^* is smoother than r_t (i.e., $\text{var}(\Delta r_t^*) < \text{var}(\Delta r_t)$), 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 $\Delta \tilde{r}_t$ (i.e., $\sum_{i=1}^n \omega_i e_{it}$) is weaker than the signal in $\Delta \tilde{r}_t$

¹⁰In our application, we find evidence for MA dynamics in $\Delta \tilde{r}_t$ instead of a more general ARMA structure such as implied in (11). However, we note that a lack of evidence for AR dynamics does not, on its own, mean the measurement error is particularly small, as there could be near cancellation of some AR and MA roots for some parameterizations and measurement error processes.

(i.e., Δr_t). Furthermore, the $\sum_{i=1}^n \gamma_i (\Delta u_{it} - \Delta u_{it-1})$ term in (12) 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 $\Delta \tilde{r}_t$ than $\Delta \check{r}_t$ means that serial correlation could be considerably easier to detect in $\Delta \tilde{r}_t$ than in the projection error e_{1t} , which is what we find in our empirical analysis.

To implement the second step of our proposed procedure using an ARMA forecasting model for $\Delta \tilde{r}_t$, we determine the AR and MA orders by empirical examination of the “BN-cycle-adjusted” data implied by the first-stage estimation of \hat{r}_t^c . The model is given by

$$\phi(L)\Delta \tilde{r}_t = \theta(L)\epsilon_t, \quad (13)$$

where ϵ_t is a serially-uncorrelated forecast error. Then, assuming direct observation of the ϵ_t forecast errors, the second-stage BN estimate of the change in trend is

$$\Delta \hat{r}_t^* \equiv \frac{\theta(1)}{\phi(1)}\epsilon_t, \quad (14)$$

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

$$\hat{r}_t^* = \hat{r}_0^* + \sum_{\tau=1}^t \Delta \hat{r}_\tau^*, \text{ where } \hat{r}_0^* = \frac{1}{T} \sum_{t=1}^T \tilde{r}_t - \frac{1}{T} \sum_{t=1}^T \sum_{\tau=1}^t \Delta \hat{r}_\tau^*. \quad (15)$$

We note that level of r_t^* can also be estimated by casting the ARMA model into state-space form and following the calculation for the BN estimate of trend in [Morley \(2002\)](#) using the Kalman filter, which would also address the lack of direct observation of ϵ_t if we were unwilling to make an assumption about initial values when conducting maximum likelihood estimation. In practice, the two approaches lead to very similar estimates.

The key point is that as long as we can approximate the serial correlation in $\Delta \tilde{r}_t$ with a simple forecasting model, we can refine the estimate of r_t^* compared the basic BN decomposition for \check{r}_t .¹¹ Meanwhile, if there were actually very little measurement error in

¹¹There remains an open question about what is the best estimate of the cyclical component of r_t . If

the first place, then $\Delta\tilde{r}_t$ should be close to serially uncorrelated and the second-stage BN estimate of trend would be quite similar to the first-stage estimate. Conversely, if $\Delta\hat{r}_t^*$ continues to display serial correlation given a misspecified ARMA model in the second step, then another iteration of our proposed procedure could be conducted or the original forecasting model for $\Delta\tilde{r}_t$ could be improved until the estimated forecast errors appear serially uncorrelated. Given the calculation in (14), it is straightforward to see that the estimated change in r_t^* will be serially uncorrelated if the forecast errors for the ARMA model are serially uncorrelated. It is also worth noting that it would be more feasible to consider time-varying parameter versions of univariate ARMA models for $\Delta\tilde{r}_t$ than for high-dimensional multivariate models used in the first-stage estimation if there is any evidence of structural change in the ARMA parameters for $\Delta\tilde{r}_t$.¹²

2.3 Misspecification due to omitted variables or dynamics

We note that our two-step approach can also handle model misspecification due to omitted variables or dynamics. In particular, if only a subset of variables in \mathbf{x}_t is included in $\tilde{\mathbf{x}}_t$, it will imply serial correlation in $\tilde{\mathbf{e}}_t$ even if there is 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 $\tilde{\mathbf{e}}_t$ inherits these even if there is no measurement error.

Consider, for example, a VAR(p) structure for $\Delta\mathbf{x}_t$, but $\tilde{\mathbf{x}}_t$ only includes the first $m < n$ variables in $\tilde{\mathbf{x}}_t$. Again, assume no measurement error. Then following Corollary 11.1.2 in Lütkepohl (2005), $\Delta\tilde{\mathbf{x}}_t$ will have a VARMA($\leq p(m - n + 1), \leq p(m - n)$) structure and the cyclical terms in (7) 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 projection errors, similar to the

we use the second-stage estimate $\hat{r}_t^c \equiv \tilde{r}_t - \hat{r}_t^*$, it will clearly include some of the measurement error in \tilde{r}_t in addition to the true cyclical component, while the first-stage BN cycle \tilde{r}_t^c may be closer to $\mathbb{E}[r_t^c | \Omega_t]$. To the extent that the measurement error appears to be small given an apparent lack of serial correlation in the projection error \tilde{e}_{1t} , we suggest using the second-stage estimate \hat{r}_t^c .

¹²In our application, we find that \hat{e}_t from our MA(8) model for $\Delta\tilde{r}_t$ appear to be serially uncorrelated in both the first and second halves of the sample period, suggesting that a stable MA(8) model is sufficient to capture serial correlation in $\Delta\tilde{r}_t$.

case of a small amount of measurement error. The point is that $\Delta\tilde{r}_t$ could still exhibit serial correlation given model misspecification and the result in (7) would still hold. Thus, our two-step procedure is more broadly robust to model misspecification than just the case of misspecification due to measurement error. In our application, the consideration of a large set of variables should help assuage concerns about omitted variables, but it is certainly possible that the apparent serial correlation in the changes in the first-stage estimate of the BN trend for r_t is due to misspecification of model dynamics as well as, or even instead of, measurement error.

2.4 Comparison with other approaches

In general, the two-step approach will 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 with population parameters for $\Delta\tilde{\mathbf{x}}_t$ that captures the serial correlation in $\tilde{\mathbf{e}}_t$. For example, if the model for $\Delta\mathbf{x}_t$ is a VAR, the correctly-specified forecasting model for the data measured with error $\Delta\tilde{\mathbf{x}}_t$ 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 two-step approach is feasible given 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. Meanwhile, by incorporating multivariate information in the first step of our proposed approach, it should be more precise than the BN decomposition based on even the correctly-specified univariate model $\Delta\tilde{r}_t$ (or, equivalently, Δr_t under the assumption of no measurement error in the target variable). In particular, our two-step approach can be thought of as partitioning the parameter space into two, using multivariate information $\{\tilde{\mathbf{X}}_t, \tilde{\mathbf{X}}_{t-1}\}$ to estimate the projection parameters \mathbf{P} and univariate information $\{\tilde{r}_t, \tilde{r}_{t-1}, \dots, \tilde{r}_1\}$ for a constructed series that more strongly reflects measurement error than $\Delta\tilde{r}_t$ to estimate univariate ARMA parameters for $\Delta\tilde{r}_t$ implied by the multivariate VARMA process. The tradeoff is less information than the full multivariate approach (but still more than a full univariate approach) in return for

far fewer parameters to estimate than for the full multivariate model.

For purposes of comparison, we conduct a Monte Carlo experiment to consider different approaches to estimating the BN trend in the presence of measurement error. We focus on the case where the measurement error is small in the sense that the projection errors for the misspecified model will not display much serial correlation on their own.

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 $r_t^* = \mathbb{E}[r_t^* | \Omega_t]$ and $r_t^c = \mathbb{E}[r_t^c | \Omega_t]$. Specifically, consider a bivariate VAR(1) for the underlying data process $\Delta \mathbf{x}_t$ with $\boldsymbol{\mu} = \mathbf{0}$. In this case, $\Delta \mathbf{X}_t = \Delta \mathbf{x}_t$, with the following simplified form:

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

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 $r_t^c = -\mathbf{s}'_{2,1} \mathbf{F}(\mathbf{I} - \mathbf{F})^{-1} \Delta \mathbf{X}_t$. For the observed data $\tilde{\mathbf{x}}_t = \mathbf{x}_t + \mathbf{u}_t$, where $\tilde{r}_t = r_t + u_{1t}$ and $\tilde{x}_{2t} = x_{2t} + u_{2t}$, we assume the addition of serially-uncorrelated measurement – i.e., $\mathbf{u}_t \sim \mathcal{N}(\mathbf{0}, 0.05 \times \boldsymbol{\Sigma})$. Again following Corollary 11.1.1 in [Lütkepohl \(2005\)](#), we note that this DGP implies the following processes: $\Delta r_t \sim \text{ARMA}(2, 1)$, $\Delta \tilde{\mathbf{x}}_t \sim \text{VARMA}(1, 2)$, $\Delta \tilde{r}_t \sim \text{ARMA}(2, 3)$, and $\Delta \tilde{x}_{2t} \sim \text{ARMA}(4, 6)$. However, given the small degree of measurement error and $\mathbf{P} \approx \mathbf{F}$, we get $\Delta \tilde{r}_t \approx \text{MA}(2)$.

To evaluate the accuracy of different estimates of r_t^* , we calculate the root-mean-squared-error (RMSE) for \hat{r}_t^* based on the following model or models given a sample size of $T = 200$:

1. **VAR(1) for $\Delta \mathbf{x}_t$** , where $\hat{r}_t^* = \mathbb{E}[r_t^* | \mathbf{X}_t, \mathbf{X}_{t-1}; \hat{\mathbf{F}}]$. This scenario corresponds to the case where the underlying data without measurement error is observed and the RMSE only reflects 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 (3).

2. **VARMA(1,2) for $\Delta\check{\mathbf{x}}_t$** , where $\hat{r}_t^* = \mathbb{E}[r_t^* | \check{\mathbf{X}}_t, \check{\mathbf{X}}_{t-1}, \dots, \check{\mathbf{X}}_1; \hat{\mathbf{F}}, \hat{\Theta}]$, with Θ corresponding to the MA parameters in the VARMA model. This scenario allows us to consider the effects of measurement error and additional parameters to estimate, but assuming the correct model specification for $\Delta\check{\mathbf{x}}_t$. The RMSE reflects noisy information about $\{\mathbf{X}_t, \mathbf{X}_{t-1}\}$ and estimation uncertainty about the $n^2 = 4$ parameters in \mathbf{F} and the $2n^2 = 8$ parameters in Θ . 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 $\Delta\check{\mathbf{x}}_t$** , where $\hat{r}_t^* = \mathbb{E}[r_t^* | \check{\mathbf{X}}_t, \check{\mathbf{X}}_{t-1}, \dots, \check{\mathbf{X}}_1; \mathbf{F}, \Theta]$. This scenario considers the effects of information loss from only observing the data with measurement error, but assumes the correct model specification for $\Delta\check{\mathbf{x}}_t$ and the true parameter values for the model. Note that the Θ parameters depend on \mathbf{F} and \hat{r}_t^* can be calculated based on their implicit values by casting the model in (6) into state-space form and using the Kalman filter following [Morley \(2002\)](#).
4. **VAR(1) for $\Delta\check{\mathbf{x}}_t$** , where $\hat{r}_t^* = \mathbb{E}[r_t^* | \check{\mathbf{X}}_t, \check{\mathbf{X}}_{t-1}; \mathbf{F} = \hat{\mathbf{P}}, \Theta = \mathbf{0}]$. This scenario considers what happens if first-stage estimates are used despite model misspecification. That is, $\hat{\mathbf{P}}$ is incorrectly used to estimate \mathbf{F} and $\Delta\check{\mathbf{x}}_t$ is incorrectly assumed to have no measurement error (i.e., $\Delta\mathbf{x}_t = \Delta\check{\mathbf{x}}_t$ and $\Theta = \mathbf{0}$). As this scenario is completely analogous to the first one, just with the incorrect assumptions, parameter estimation is conducted via OLS and the BN trend is calculated based on (3).
5. **VAR(1) for $\Delta\check{\mathbf{x}}_t + \text{MA}(2)$ for $\Delta\tilde{r}_t$** , where $\hat{r}_t^* = \mathbb{E}[r_t^* | \check{\mathbf{X}}_t, \check{\mathbf{X}}_{t-1}; \hat{\mathbf{P}}, \hat{\theta}_1, \hat{\theta}_2]$. This scenario considers our proposed two-step procedure using the parsimonious MA(2) model to capture serial correlation in $\Delta\tilde{r}_t$. $\hat{\mathbf{P}}$ is estimated via OLS and the MA parameters are estimated via conditional MLE.
6. **VAR(1) for $\Delta\check{\mathbf{x}}_t + \text{ARMA}(4,6)$ for $\Delta\tilde{r}_t$** , where $\hat{r}_t^* = \mathbb{E}[r_t^* | \check{\mathbf{X}}_t, \check{\mathbf{X}}_{t-1}; \hat{\mathbf{P}}, \hat{\phi}_1, \dots, \hat{\phi}_4, \hat{\theta}_1, \dots, \hat{\theta}_6]$. This scenario considers our proposed two-step procedure using the full

true ARMA(4,6) model specification to capture serial correlation in $\Delta\tilde{r}_t$. $\hat{\mathbf{P}}$ 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 $\hat{r}_t^* = \mathbb{E}[r_t^* | r_t, r_{t-1}, \dots, r_1; \hat{\phi}_1, \hat{\phi}_2, \hat{\theta}]$. This scenario corresponds to the case where only the target variable is 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 $\hat{r}_t^* = \mathbb{E}[r_t^* | \check{r}_t, \check{r}_{t-1}, \dots, \check{r}_1; \hat{\phi}_1, \hat{\phi}_2, \hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3]$. This scenario consider the effects of measurement error and additional parameters, but assumes 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 our Monte Carlo experiment 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 are approximately equal to percentages of the standard deviation of the cyclical component of the simulated real interest rate.

First, even given the correct model and no measurement error (i.e., VAR(1) for $\Delta\mathbf{x}_t$), parameter estimation introduces a nontrivial amount of error in the estimates of r_t^* due to

Table 1: Accuracy of \hat{r}_t^* in Monte Carlo Experiment

Model(s)	# of Parameters to Estimate	RMSE
VAR(1) for $\Delta \mathbf{x}_t$	4	0.18
VARMA(1,2) for $\Delta \check{\mathbf{x}}_t$	12	0.33
True VARMA(1,2) for $\Delta \check{\mathbf{x}}_t$	0	0.05
VAR(1) for $\Delta \check{\mathbf{x}}_t$	4	0.37
VAR(1) for $\Delta \check{\mathbf{x}}_t + \text{MA}(2)$ for $\Delta \tilde{r}_t$	6	0.33
VAR(1) for $\Delta \check{\mathbf{x}}_t + \text{ARMA}(4,6)$ for $\Delta \tilde{r}_t$	14	0.37
ARMA(2,1) for Δr_t	3	0.94
ARMA(2,3) for $\Delta \tilde{r}_t$	5	0.97

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

the sample size 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 (i.e., VARMA(1,2) for $\Delta \check{\mathbf{x}}_t$), almost doubles the error in estimating r_t^* due the noisier information given the measurement error and the increase in parameter uncertainty given 12 parameters instead of 4. However, most of the increase in RMSE seems to be due to parameter uncertainty as the RMSE given the true parameter values is only 0.05 compared to 0.33 for the estimated VARMA model, suggesting that noisier information has a smaller effect than parameter uncertainty, reflecting the assumption of relatively small measurement error.

Second, model misspecification has a larger effect on the RMSE than parameter uncertainty given that the RMSE for the misspecified model (i.e., VAR(1) for $\Delta \check{\mathbf{x}}_t$) is larger than for the correctly-specified VARMA model with more parameters. The proposed two-step procedure with the parsimonious second-stage model (i.e., VAR(1) for $\Delta \check{\mathbf{x}}_t + \text{MA}(2)$ for $\Delta \tilde{r}_t$) offsets the effects of misspecification and reduces the RMSE to the same level as for the more highly-parameterized VARMA model.¹³ This reduction compared

¹³We note that the relatively good performance of the VARMA model is quite specific to this bivariate DGP. When considering $n > 2$, the number of VARMA($p,2$) parameters proliferates (i.e., the number of parameters to estimate is $n^2(p+2)$), with the RMSE increasing substantially with 27 parameters for a VARMA(1,2) given $n = 3$ and even given a much larger sample size of $T = 2000$, while the RMSE for

to the misspecified model occurs even though $\mathbf{P} \approx \mathbf{F}$ given a small amount of measurement error. Meanwhile, if the full ARMA model is considered in the second stage (i.e., VAR(1) for $\Delta\tilde{\mathbf{x}}_t + \text{ARMA}(4,6)$ for $\Delta\tilde{r}_t$), the RMSE is higher than with the parsimonious second-stage model and is the same as for the misspecified VAR model as a result of the parameter proliferation to even more parameters than in the true VARMA model in this case. 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.

Third, the results are considerably worse for the univariate estimates, even when there is no measurement error and the true ARMA is tightly parameterized (i.e., ARMA(2,1) for Δr_t). Adding in measurement error and additional ARMA parameters (i.e., ARMA(2,3) for $\Delta\tilde{r}_t$) slightly worsens the RMSE from 0.94 to 0.97, but the fact that this is essentially three times worse than the RMSE of 0.33 for the two-step procedure with the parsimonious second-stage model, which also involves measurement error and has an additional parameter, makes it clear just how crucial incorporating multivariate information is for producing a relatively accurate estimate of r_t^* .

2.5 Informational accounting given the two-step procedure

One issue with our proposed two-step approach that needs to be addressed is how to conduct an informational accounting (or a more structural variance decomposition given identification assumptions) of changes in trend along the lines presented in [Morley and Wong \(2020\)](#) for the multivariate BN decomposition, but allowing for model misspecification. The challenge is that the second step of our proposed procedure produces estimated changes in trend that are based on the univariate forecast error for $\Delta\tilde{r}_t$ rather than underlying forecast errors for different variables, as in (5). However, we note that, given the invertible representation for $\theta(L)$, we can solve for the forecast error from the ARMA model for $\Delta\tilde{r}_t$ in terms of current and past projection errors $\tilde{\mathbf{e}}_t$ from the multivariate

the parsimonious two-step approach with 11 parameters to estimate remains relatively small and better than for the misspecified VAR model with 9 parameters.

forecasting model for $\Delta\check{\mathbf{X}}_t$:

$$\epsilon_t = \theta(L)^{-1}\phi(L)\Delta\tilde{r}_t = \theta(L)^{-1}\phi(L)\mathbf{s}'_{k,1}((\mathbf{I} - \mathbf{P})^{-1}\mathbf{H}\check{\epsilon}_t + (\mathbf{F} - \mathbf{P})(\mathbf{I} - \mathbf{P})^{-1}\Delta\check{\mathbf{X}}_{t-1}), \quad (16)$$

Then the informational contributions of the *current* projection errors from the multivariate model to $\Delta\hat{r}_t^*$ are given by

$$\Delta\hat{r}_{it}^* \equiv \frac{\theta(1)}{\phi(1)}\check{\omega}_i\check{\epsilon}_{it}, \quad (17)$$

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

$$\sum_{i=1}^n \Delta\hat{r}_{it}^* \approx \Delta\hat{r}_t^*. \quad (18)$$

That is, the informational contributions of the projection errors for each variable in the multivariate model to the two-step estimate $\Delta\hat{r}_t^*$ are approximately proportional to those for the first-stage estimate of the BN trend based on the multivariate forecasting model for $\Delta\check{\mathbf{x}}_t$, 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 (18) is quite accurate.

3 Data and VECM estimation

We consider both short- and long-term interest rates in our analysis 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. The short (long) *ex ante* real interest rate is constructed 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) price deflator. We then use a 4-quarter (40-quarter) rolling average of past inflation as a proxy for short (long) inflation expectations. For robustness, we also consider alternative measures of 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.

The choice of other variables to consider in our analysis 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 primarily 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. The first category corresponds to supply-side productivity/demographic type factors. The second category corresponds to safe assets and the global savings glut explanation for changes in real interest rates.

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 pos-

itively 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 our 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. \(2016\)](#), 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 in our analysis, although not in our main model to estimate r^* as these series are only available at an annual frequency. Instead, we include these annual variables in subsequent cointegration analysis.

Safe assets 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^* . To address this, we consider the change in macroeconomic uncertainty ([Jurado et al., 2015](#)), the change in the excess bond premium ([Gilchrist and Zakrajsek, 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). 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. We note that the global central bank foreign reserves measure is only available at an annual frequency, so is not in our main model to estimate r^* , but is only considered with the other annual variables in subsequent cointegration analysis.

Details of the various data series are given in the Appendix. Our sample period covers 1973Q2 to 2019Q4.¹⁴ For the purposes of specifying our VECM, we denote 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 quarterly 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.¹⁵ Assuming cointegration between the short term nominal interest rate and the long term nominal interest rate implies that both are driven by only one stochastic trend that differs only by a constant α . The

¹⁴We have a balanced panel of the quarterly variables for the full sample period, although we note that the real interest rates from the Cleveland Fed considered in our robustness analysis are only available from 1982Q1. 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 macroeconomic model. Variables only available at an annual frequency and therefore not included in our VECM, but only considered in followup cointegration analysis, are the measures of age dependency, income inequality, and central bank reserves of foreign currency as a percentage of world GDP.

¹⁵We impose that the means of the first differences of the interest rates are zero, corresponding to an assumption of no deterministic drift in the levels. Thus, the historical downward movements in r^* will be attributed to forecast errors rather than deterministic drift.

VECM can then be specified as follows:

$$\Delta \mathbf{x}_t = \boldsymbol{\mu} + \Phi_1(\Delta \mathbf{x}_{t-1} - \boldsymbol{\mu}) + \dots + \Phi_p(\Delta \mathbf{x}_{t-p} - \boldsymbol{\mu}) + \boldsymbol{\beta} (r_{t-1}^l - r_{t-1}^s - \alpha) + \mathbf{e}_t, \quad (19)$$

where

$$\Delta \mathbf{x}_t = \begin{bmatrix} \Delta r_t^s \\ \Delta r_t^l \\ \Delta \mathbf{x}_{3:n,t} \end{bmatrix}, \quad \boldsymbol{\beta} = \begin{bmatrix} \beta^s \\ \beta^l \\ \mathbf{0} \end{bmatrix}.$$

The model with the error-correction term in (19) can be cast into companion form and the BN trend for the interest rates calculated and decomposed as described above, where, for the companion form, $\Delta \mathbf{X}_t = \{(\Delta \mathbf{x}_t - \boldsymbol{\mu})', \dots, (\Delta \mathbf{x}_{t-p} - \boldsymbol{\mu})', (r_t^l - r_t^s - \alpha)\}'$ and \mathbf{F} and \mathbf{H} are as implied by the VECM example in [Morley \(2002\)](#).

Because we are dealing with a medium-scale VECM with a large number of parameters, we need to take into account the possibility of overfitting. To address this, we rely on Bayesian estimation using a natural conjugate Normal-Inverse Wishart prior in conjunction with a Minnesota Prior with the shrinkage hyperparameter, λ , set to 0.2, as in [Sims and Zha \(1998\)](#) and [Carriero et al. \(2015\)](#). We also set an “expectations hypothesis” prior that the short-rate adjusts to its spread from the long rate, consistent with cointegration between the long and the short rates. Because of the VECM setup, estimation is conducted via MCMC with Gibbs sampling.¹⁶ The Bayesian VECM is estimated with four lags, as is typical for quarterly data. Details of the Bayesian estimation can be found in the Appendix.

¹⁶A Bayesian VAR that includes the spread $r_t^l - r_t^s$ instead of the change in the long rate Δr_t^l and error correction terms for the interest rate equations can be estimated analytically and produces very similar estimates of r_t^* . However, we take the VECM as our baseline specification even though estimation is more complicated.

4 Empirical results

4.1 Estimates of r^*

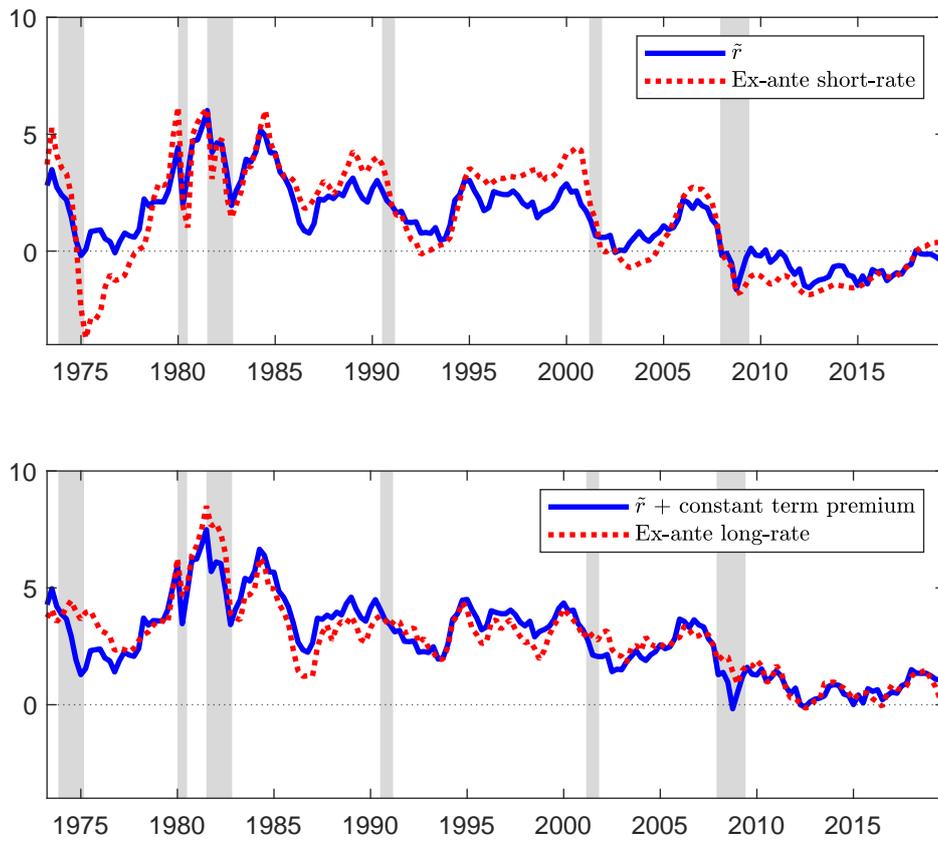
Figure 1 plots the first-stage estimates of BN trends for the short- and long-term real interest rates.¹⁷ The first-stage trend estimate (i.e., what we refer to as \tilde{r} in the assumed presence of measurement error or another source of model misspecification) appears to be highly informed by the long-rate, suggesting that the short-rate adjusts to the term spread. The presence of a cointegrating relationship between the short- and long-term rates provides important information in estimating the first-stage 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 allows for persistence in the estimated deviations of the real interest rates from trend compared to the traditional univariate BN decomposition, a result that follows from [Evans and Reichlin \(1994\)](#) and [Morley and Wong \(2020\)](#).

Figure 2 presents the second-stage BN trend for the short-term real interest rate, along with the 95% credible intervals.¹⁸ The second-stage BN trend estimate for the short-rate (i.e., our estimate of r^*) is noticeably smoother than the first-stage BN trend estimate for the short-rate (\tilde{r}) in Figure 1. 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 estimates is consistent with the previous literature (see, for example,

¹⁷For the first-stage trend estimates, we take the Bayesian VECM parameter estimates at the posterior mean and cast the model into companion form, as in (2). We then use (3) to get an estimate of the trend in the short-term real interest rate. Given cointegration, the level of the BN trend for the long-term real interest rate only differs from the BN trend for the short-term real interest rate by the constant α that presumably reflects a long-run term premium.

¹⁸For the second-stage trend estimate, we again take the Bayesian VECM parameter estimates at the posterior mean and cast the model into companion form, but now conceptually given potential misspecification, as in (6). We then calculate an estimate of r^* following (15) given (14) based on an ARMA model of $\Delta\tilde{r}_t$, as in (13). As noted previously, we find that an MA(8) model estimated via conditional MLE is sufficient to capture the serial correlation in $\Delta\tilde{r}_t$.

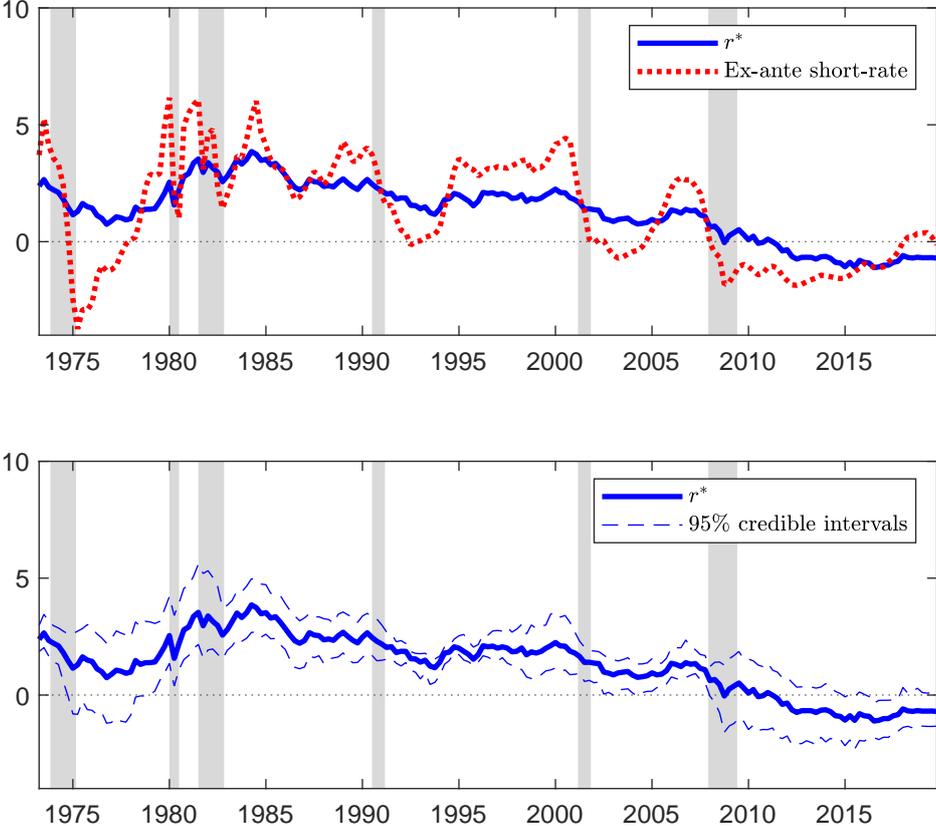
Figure 1: First-stage BN trends for short- and long-term real interest rates



Note: NBER recession dates are shaded.

Del Negro et al. (2017); Hamilton et al. (2016); Holston et al. (2017); Lubik and Matthes (2015); Bauer and Rudebusch (2020); Kiley (2020); Johannsen and Mertens (2021); Cúrdia et al. (2015); Lewis and Vazquez-Grande (2019), but it should be emphasized that we do not impose smoothness in our estimation *a priori*. We also highlight that r^* has been persistently low since the Great Recession and has even been slightly negative since around 2013, although it is only briefly significantly negative according to the 95% credible intervals.

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



Note: NBER recession dates are shaded.

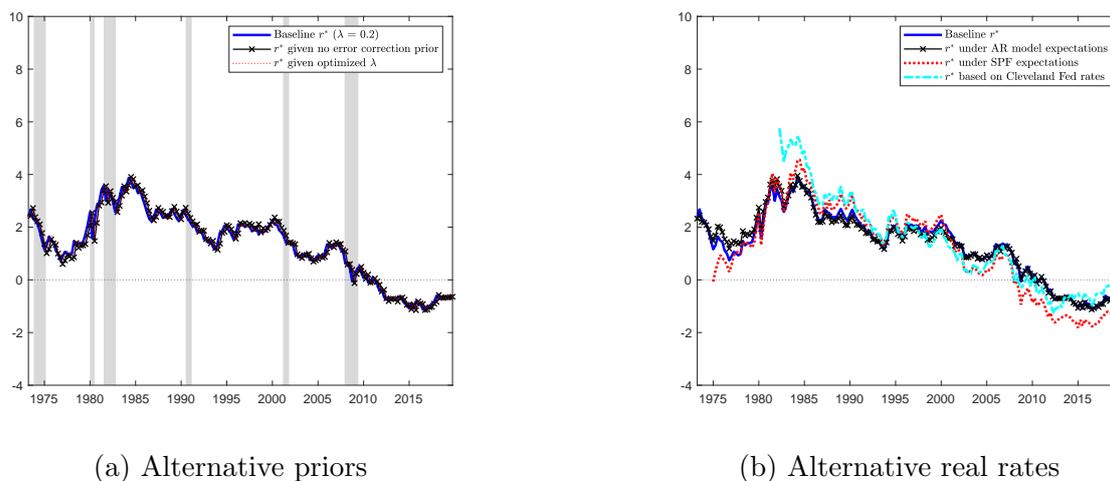
Robustness Checks

For robustness, 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 Figure 3a. 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.

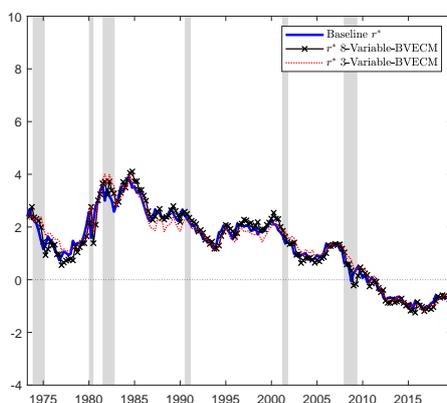
We also investigate the sensitivity of our r^* estimates with respect to how we proxy inflation expectations when measuring *ex ante* real interest rates. As shown in Figure 3b, our estimates are generally robust to three alternative measures of inflation expectations. In the first case, following Laubach 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.

Figure 3: Robustness of r^* estimates



(a) Alternative priors

(b) Alternative real rates



(c) Alternative information sets

Note: NBER recession dates are shaded.

Finally, to confirm the relevance of different sources of information in our model, we consider two alternative models in terms of which variables are included. In the first case, we consider a smaller 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 model that only includes the ‘interest rate’ block. As shown in [Figure 3c](#), 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^*

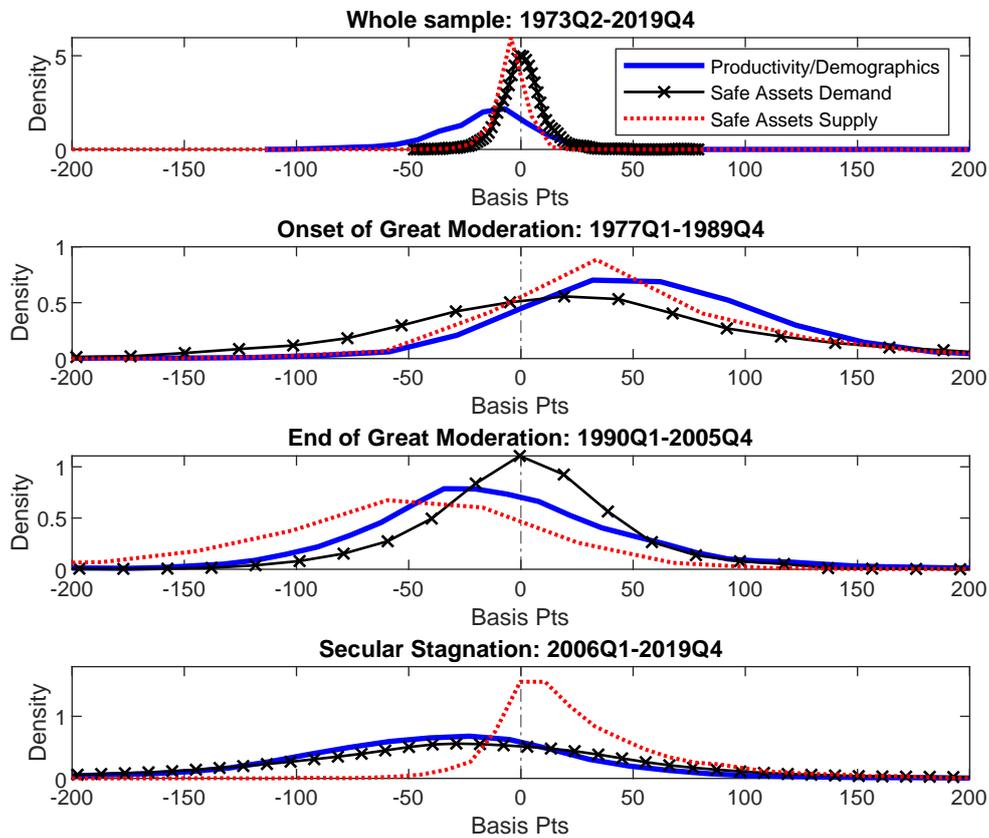
changes a bit more, however, when we do not include any possible determinants beyond the ‘interest rate’ block. But the general similarity of the estimates even when only including interest rates has two important implications: (i) any measurement error or other source of model misspecification that generates serial correlation in the first-stage estimates of trend growth must be primarily related to the *ex ante* real interest rates in particular and (ii) one could obtain a reasonably 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. Of course, the baseline medium-scale VECM has the advantage of allowing us to track which economic forces are most important in driving changes in r^* , which is what we turn to next.

4.2 Why did r^* change?

Having estimated r^* using a multivariate trend-cycle decomposition, we now look at the evidence for possible sources of changes in it over time. 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 to the fall in r^* , with most of the posterior density below zero, although the probability that these variables contributed more than 50 basis points to the overall estimated decline of about 3 percentage points appears to be reasonably low. Safe assets 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 assets demand (all of the other quarterly variables in the safe assets 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 (from 2006Q1), we can see that the different groups of variables made more substantial contributions to large changes in r^* within the full sample period. For example, the posterior densities for productivity/demographics and safe assets demand/supply all suggest positive contributions to the 1.3 percentage point increase in r^* with the onset of the

Great Moderation, with a high probability that productivity/demographics variables contributed more than 50 basis points. Likewise, productivity/demographics and safe assets supply appear to contribute to the 1.4 percentage point decline in r^* up to the end of the Great Moderation, with safe assets supply having the highest probability of contributing more than 50 basis points to the decline. Finally, productivity/demographics and safe assets demand appeared to contribute sizeable and similar amounts to the 2.1 percentage point decline in r^* during the Secular Stagnation era, although safe assets supply clearly has at least a partially offsetting effect.

Figure 4: Contributions to r^* : Posterior densities



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

Table 2 presents results for the individual variables in terms of their effects on estimated r^* and their contributions during the three subsample episodes considered in Figure 4. Before looking at these individual estimated contributions, we first note that, according

to the table, the quarterly variables all have their predicted effects, with ‘Sign’ referring to the implied correlation between the forecast error for a variable and the implied change in the estimate of r^* based on the decomposition in (5) (or (17) if we assume model misspecification).¹⁹ To get a sense of significance, we also report the posterior probability of the reported direction of correlation. We find that there are 63% and 73% 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](#); [Lunsford and West, 2019](#)). Investment growth has the predicted negative relationship consistent with investment-specific technological change, but with only a 59% 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 68% and 85% posterior probabilities that employment and hours growth have a positive relationship with r^* , while the unemployment rate has a negative relationship with 69% 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 assets demand/flight-to-safety phenomenon, there are 67% and 85% 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 55% posterior probability, possibly reflecting more of a mix of supply and demand factors ultimately driving this variable. Furthermore, consistent with the global savings glut hypothesis ([Bernanke, 2005](#)), there are 89% and 81% 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 assets supply/crowding-out effect ([Ball and Mankiw, 1995](#)).

¹⁹We emphasize that this is a correlation only and not necessarily a causal link. However, the signs are consistent with the causal effects discussed in the previous section. We also note that a more structural analysis could be considered given identification restrictions for a structural VAR. However, we leave such analysis for future research.

Table 2: Accounting for changes in r^*

	Sign	Probability	Contribution (bps)		
			Onset of Great Moderation	End of Great Moderation	Secular Stagnation
I. Productivity/Demographics			63	-15	-41
Real consumption growth per capita	+	0.63	5	4	-9
TFP growth	+	0.73	2	18	-16
S&P 500 stock returns	+	0.75	4	12	-5
Real investment growth	-	0.59	0	-5	5
Employment growth	+	0.68	11	-12	-12
Hours growth	+	0.85	33	-23	-8
Unemployment rate (Δ)	-	0.69	9	-9	4
II. Safe Assets Demand			20	2	-26
Macroeconomic uncertainty (Δ)	-	0.67	9	1	-10
Excess bond premium (Δ)	-	0.85	0	24	-28
Liquid assets growth	-	0.56	-10	2	9
Current account as % of GDP (Δ)	+	0.89	12	-28	13
Exchange rate return	-	0.81	10	3	-10
III. Safe Assets Supply			46	-59	25
Government debt as % of GDP (Δ)	+	0.84	46	-59	25

In terms of individual contributions, Table 2 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 assets demand had somewhat offsetting effects with the onset of the Great Moderation, while higher safe assets 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 assets 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, 2019). 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 assets demand related variables of macroeconomic uncertainty and the excess bond premium, although the other safe assets demand related variables mostly had offsetting effects, as did the increase in safe assets supply with a higher debt-to-GDP ratio again.

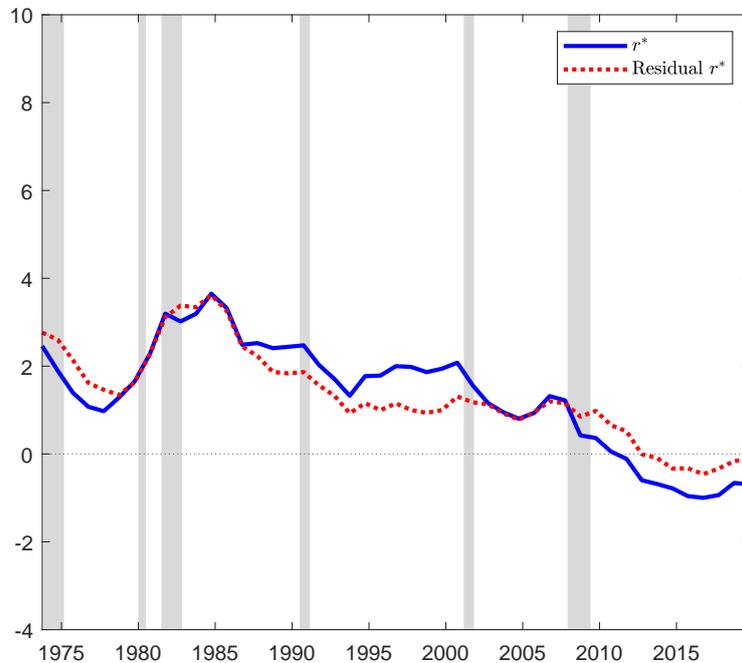
Although the various informational variables explain some of the major movements in r^* over time, we note that there remain large movements in r^* related to forecast errors for the interest rates and, therefore, are unexplained by the other informational variables. To see this, we construct a residual unexplained component of r^* by removing the contributions of the other informational variables as follows:

$$\hat{r}_t^{*resid} = \hat{r}_t^* - \sum_{i>2} \hat{r}_{it}^* \quad (20)$$

Figure 5 then compares our r^* estimate with the residual r^* . The difference between

the residual r^* and r^* confirms the role the various quarterly variables in driving r^* also evident in Figure 4 and Table 2. Before the onset of the Great Moderation, the other informational variables act to lower r^* compared to otherwise, but this is reversed for most of the Great Moderation, at least until near its end. Meanwhile, the quarterly variables clearly reduce r^* compared to otherwise during the Secular Stagnation era. However, the general decline in r^* over the whole sample period is evident in the residual measure and, therefore, is not explained by the quarterly variables. This motivates us to consider additional variables for which only annual data is available in the next subsection.

Figure 5: Residual r^* not explained by quarterly variables



Note: NBER recession dates are shaded.

4.3 Cointegration Analysis with Annual Variables

To investigate possible connections between r^* and the annual variables, we consider cointegration analysis of long-run relationships given that there are relatively few observations with which to determine any higher-frequency links. It should be noted that our previous analysis with quarterly variables implicitly assumed no cointegration other than between the real interest rates. This assumption is confirmed by apparent omitted I(1) variables

given that the residual r^{*resid} measure tests as being I(1). Also, we find evidence of possible cointegration between annual variables and the residual r^{*resid} rather than with r^* itself.

The annual variables that we consider are income inequality, age dependency, and global reserves-to-GDP, all of which we assume to be exogenous with respect to r^* . We note that using ADF tests and Engle-Granger cointegration analysis, age dependency actually tests as trend stationary and has the theoretically ‘wrong’ (i.e., negative) sign when included in cointegration regressions despite not testing as having a unit root.²⁰ Inequality appears to have a unit root and has the predicted positive sign, but cointegration tests with the residual r^{*resid} measure are insignificant. Only global reserves-to-GDP appears to have a unit root and tests as having a cointegrating relationship with the residual r^{*resid} measure, with the predicted negative sign corresponding to it being a proxy for the global savings glut. This significance holds even when including a deterministic time trend in the cointegration test regression. The Engle-Granger test for cointegration has a p -value of 0.03 and the cointegration regression results are given as follows:

$$\hat{r}_t^{*resid} = 2.86 - 0.06 \cdot t - 2.83 \cdot globalreserves_t + \hat{z}_t,$$

where t is a time trend, $globalreserves_t$ denotes global reserves-to-GDP, and \hat{z}_t is the cointegrating error. As seen in Figure 6, these estimates imply global reserves explain an additional 30 basis point decline in r^* beyond the quarterly variables since the Great Moderation.

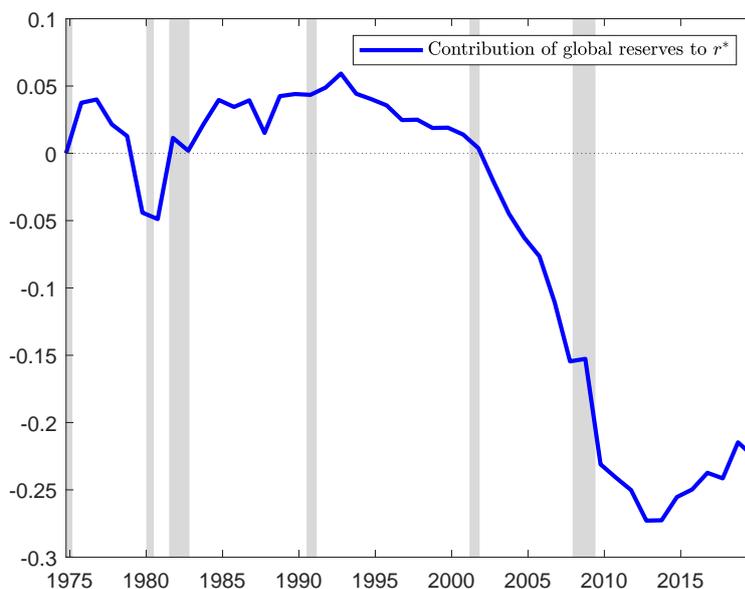
4.4 How has r^* changed since COVID-19?

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. Figure 7 plots the estimate of r^* over the latter part of the sample period and up to the end of 2020.²¹ At the start of

²⁰Gagnon et al. (2016) suggests that an increase in the dependency ratio will reduce aggregate savings and raise r^* .

²¹Given the same parameter estimates, the pre-Covid estimates of r^* are the same as in Figure 2.

Figure 6: Estimated contribution of global reserves to r^*



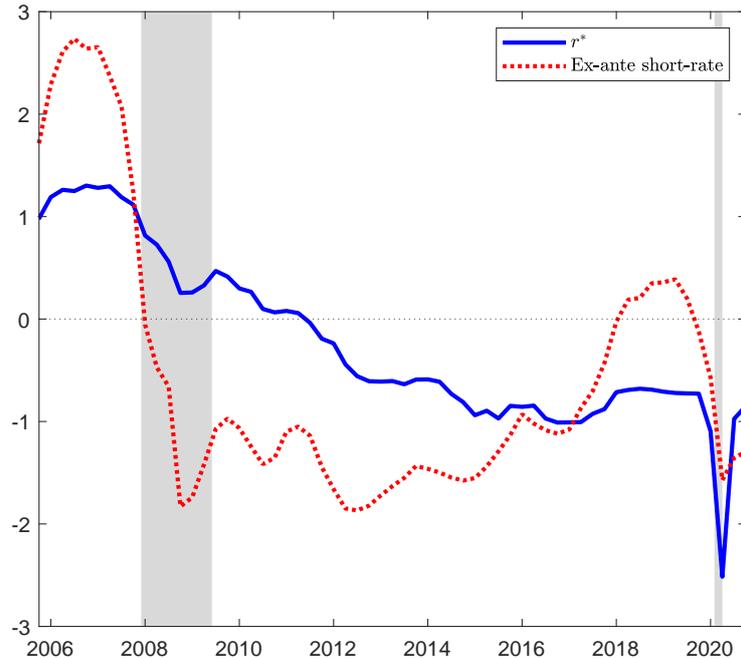
Notes: NBER recession dates are shaded. The initial contribution is normalized to zero.

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 effects was different than in the previous sample given the unusual stop-start nature of economic activity with lockdowns, as well as the massive fiscal stimulus, with the estimated r^* quickly returning to its pre-pandemic level of around -1%. 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 a 43 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 a 33 basis point increase over the same period.

As noted in the introduction, the estimated r^* can be used to help gauge the stance of monetary policy during the pandemic. The initial sharp decrease in r^* is estimated to be of a similar scale to the decline in the *ex ante* short-term real interest rate, but the immediate recovery back to a similar level around -1% as before the pandemic is

back above the *ex ante* short-term real interest rate. These estimates imply that, after little initial change in the positive real interest rate gap, monetary policy quickly became somewhat accommodative with a negative gap despite the ZLB, although not nearly as much as if r^* were closer to its historical levels of 2% instead of -1%.

Figure 7: Second-stage BN trend for the short-term real interest including 2020



Note: NBER recession dates are shaded.

5 Conclusion

In this paper, we have measured r^* as the common stochastic trend in real interest rates using a multivariate BN decomposition based on a Bayesian VECM. We have developed a robust two-step approach to multivariate trend-cycle decomposition that addresses the apparent predictability of changes in the first-stage estimate of the BN trend and leads to a smooth second-stage estimate of r^* . We provide a comprehensive explanation for why r^* has declined over the past few decades by measuring informational contributions of different variables to its historical movements. We find that the decline in r^* is related to supply-side productivity/demographic factors and to safe assets demand/supply.

Slower trend growth and higher safe assets demand as part of the global savings glut both contributed to lower r^* during the recent Secular Stagnation era, with some offsetting positive effects from crowding out due to increase safe assets supply in the form of increased government debt-to-GDP. These offsetting effects continued during the COVID-19 pandemic.

References

Baker, Dean, J Bradford De Long, and Paul R Krugman, “Asset returns and economic growth,” *Brookings Papers on Economic Activity*, 2005, *2005* (1), 289–330.

Ball, Laurence and N Gregory Mankiw, “What do budget deficits do?,” Technical Report, National Bureau of Economic Research 1995.

Bauer, Michael D. and Glenn D. Rudebusch, “Interest Rates under Falling Stars,” *American Economic Review*, May 2020, *110* (5), 1316–54.

Bernanke, Ben, “The Global Saving Glut and the US Current Account Deficit, Homer Jones Lecture, St,” *Louis, Missouri, April*, 2005, *14*, 2005.

Beveridge, Stephen and Charles 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*, 1981, *7* (2), 151–174.

Caballero, Ricardo J, Emmanuel Farhi, and Pierre-Olivier Gourinchas, “The safe assets shortage conundrum,” *Journal of Economic Perspectives*, 2017, *31* (3), 29–46.

Carriero, Andrea, Todd E Clark, and Massimiliano Marcellino, “Bayesian VARs: specification choices and forecast accuracy,” *Journal of Applied Econometrics*, 2015, *30* (1), 46–73.

- Clark, Todd E and Taeyoung Doh**, “Evaluating alternative models of trend inflation,” *International Journal of Forecasting*, 2014, 30 (3), 426–448.
- Cúrdia, Vasco, Andrea Ferrero, Ging Cee Ng, and Andrea Tambalotti**, “Has US monetary policy tracked the efficient interest rate?,” *Journal of Monetary Economics*, 2015, 70, 72–83.
- DiCecio, Riccardo**, “Sticky wages and sectoral labor comovement,” *Journal of Economic Dynamics and Control*, 2009, 33(3), 538–553.
- Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes**, “Do the Rich Save More?,” *Journal of Political Economy*, 2004, 112 (2), 397–444.
- Eichengreen, Barry**, “Secular Stagnation: The Long View,” *American Economic Review*, May 2015, 105 (5), 66–70.
- Evans, George and Lucrezia Reichlin**, “Information, forecasts, and measurement of the business cycle,” *Journal of Monetary Economics*, 1994, 33 (2), 233–254.
- Fernald, John G**, “Productivity and Potential Output before, during, and after the Great Recession,” *NBER Macroeconomics Annual*, 2015, 29 (1), 1–51.
- Gagnon, Etienne, Benjamin K. Johansson, and J. David Lopez-Salido**, “Understanding the New Normal : The Role of Demographics,” Finance and Economics Discussion Series 2016-080, Board of Governors of the Federal Reserve System (U.S.) September 2016.
- Gilchrist, Simon and Egon Zakrajšek**, “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, 2012, 102(4), 1692–1720.
- Glick, Reuven**, “R* and the Global Economy,” *Journal of International Money and Finance*, 2019, p. 102105.
- González-Astudillo, Manuel and Jean-Philippe Laforte**, “Understanding the New Normal : The Role of Demographics,” Finance and Economics Discussion Series 2020-085, Board of Governors of the Federal Reserve System (U.S.) September 2020.

- Hamilton, James D, Ethan S Harris, Jan Hatzius, and Kenneth D West**, “The equilibrium real funds rate: Past, present, and future,” *IMF Economic Review*, 2016, 64 (4), 660–707.
- Holston, Kathryn, Thomas Laubach, and John C Williams**, “Measuring the natural rate of interest: International trends and determinants,” *Journal of International Economics*, 2017, 108, S59–S75.
- Johannsen, Benjamin K. and Elmar Mertens**, “A Time-Series Model of Interest Rates with the Effective Lower Bound,” *Journal of Money, Credit and Banking*, 2021, 53 (5), 1005–1046.
- Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng**, “Measuring Uncertainty,” *American Economic Review*, 2015, 105(3), 1177–1216.
- Kamber, Güneş and Benjamin Wong**, “Global factors and trend inflation,” *Journal of International Economics*, 2020, 122, 103265.
- , **James Morley, and Benjamin Wong**, “Intuitive and reliable estimates of the output gap from a Beveridge-Nelson filter,” *Review of Economics and Statistics*, 2018, 100 (3), 550–566.
- Kang, Kyu Ho, Chang-Jin Kim, and James Morley**, “Changes in U.S. inflation persistence,” *Studies in Nonlinear Dynamics and Econometrics*, 2009, 13 (4), 1–21.
- Kiley, Michael T.**, “Output gaps,” *Journal of Macroeconomics*, 2013, 37, 1–18.
- , “What Can the Data Tell Us about the Equilibrium Real Interest Rate?,” *International Journal of Central Banking*, 2020, 16 (3), 181–209.
- Kilian, Lutz**, “Small-sample confidence intervals for impulse response functions,” *Review of economics and statistics*, 1998, 80 (2), 218–230.
- Koop, Gary and Dimitris Korobilis**, *Bayesian multivariate time series methods for empirical macroeconomics*, Now Publishers Inc, 2010.

- Laubach, Thomas and John C Williams**, “Measuring the natural rate of interest,” *Review of Economics and Statistics*, 2003, 85 (4), 1063–1070.
- Levin, Andrew T. and Jeremy M. Piger**, “Is inflation persistence intrinsic in industrial economies?,” Eurosystem Inflation Persistence Network Working Paper No. 334, European Central Bank April 2004.
- Lewis, Kurt F and Francisco Vazquez-Grande**, “Measuring the natural rate of interest: A note on transitory shocks,” *Journal of Applied Econometrics*, 2019, 34 (3), 425–436.
- Litterman, Robert B**, “Forecasting with Bayesian vector autoregressions—five years of experience,” *Journal of Business & Economic Statistics*, 1986, 4 (1), 25–38.
- Lubik, Thomas and Christian Matthes**, “Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches,” *Richmond Fed Economic Brief*, 2015, Oct.
- Lunsford, Kurt G and Kenneth D West**, “Some evidence on secular drivers of US safe real rates,” *American Economic Journal: Macroeconomics*, 2019, 11 (4), 113–39.
- Lütkepohl, Helmut**, *New Introduction to Multiple Time Series Analysis*, Springer, 2005.
- Modigliani, Franco and Robert J Shiller**, “Inflation, rational expectations and the term structure of interest rates,” *Economica*, 1973, 40 (157), 12–43.
- Morley, James and Benjamin Wong**, “Estimating and accounting for the output gap with large Bayesian vector autoregressions,” *Journal of Applied Econometrics*, 2020.
- **and Jeremy Piger**, “The asymmetric business cycle,” *Review of Economics and Statistics*, 2012, 94 (1), 208–221.
- Morley, James C**, “A state–space approach to calculating the Beveridge–Nelson decomposition,” *Economics Letters*, 2002, 75 (1), 123–127.

– , “The two interpretations of the Beveridge–Nelson decomposition,” *Macroeconomic Dynamics*, 2011, 15, 419–439.

Negro, Marco Del, Domenico Giannone, Marc P Giannoni, and Andrea Tambalotti, “Safety, liquidity, and the natural rate of interest,” *Brookings Papers on Economic Activity*, 2017, 2017 (1), 235–316.

Rachel, Łukasz and Lawrence H Summers, “On Secular Stagnation in the Industrialized World,” Technical Report, National Bureau of Economic Research 2019.

Rachel, Łukasz and Thomas D Smith, “Are low real interest rates here to stay?,” *International Journal of Central Banking*, 2017, 13 (3), 1–42.

Rotemberg, Julio J. and Michael Woodford, “Real-Business-Cycle Models and the Forecastable Movements in Output, Hours, and Consumption,” *American Economic Review*, 1996, 86 (1), 71–89.

Sims, Chris and Tao Zha, “Bayesian Methods for Dynamic Multivariate Models,” *International Economic Review*, 1998, 39(4), 949–968.

Summers, Lawrence H, “Demand side secular stagnation,” *American Economic Review*, 2015, 105 (5), 60–65.

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

Appendix

A Data Sources and Transformations

[Morley and Wong \(2020\)](#) find that the multivariate BN decomposition can be sensitive to including nonstationarity or highly persistent time series. Thus, we transform the data to ensure stationarity. Details of data sources and transformations 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	
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:SPPODPNDOLUSA	annual only
Top 1% Share of Pre-Tax National Income	World Inequality Database	annual only
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	ratio, annual only

B Bayesian Estimation

Recall that $\Delta \mathbf{x}_t$ 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’ $\Delta \mathbf{x}_{3:n,t}$.

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

$$\Delta x_{it} = \mu_i + \mathbf{w}'_{it} \mathbf{b}_i + e_{it}, \quad (\text{B1})$$

where $\mathbf{w}_{it} = [(\Delta \mathbf{x}_{t-1} - \boldsymbol{\mu})', \dots, (\Delta \mathbf{x}_{t-p} - \boldsymbol{\mu})', r_{t-1}^l - r_{t-1}^s - \alpha]'$ if $i = 1, 2$ and $\mathbf{w}_{it} = [(\Delta \mathbf{x}_{t-1} - \boldsymbol{\mu})', \dots, (\Delta \mathbf{x}_{t-p} - \boldsymbol{\mu})']'$ if $i > 2$, with \mathbf{b}_i corresponding to all of the parameters associated with equation i .

Defining

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

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

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

or

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

where

$$\mathbf{y}_i = \begin{bmatrix} y_{i1} \\ \vdots \\ y_{iT} \end{bmatrix}, \quad \mathbf{e}_i = \begin{bmatrix} e_{i1} \\ \vdots \\ e_{iT} \end{bmatrix},$$

and

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

Let Σ be an $n \times n$ covariance matrix for the VECM residuals. If one sets a Normal-Wishart prior on β and Σ respectively (Koop and Korobilis (2010)), where

$$\beta \sim N(\beta_0, \mathbf{V}_\beta), \quad (\text{B2})$$

$$\Sigma^{-1} \sim W(\mathbf{S}_0^{-1}, \nu_0), \quad (\text{B3})$$

this implies conditional distributions

$$p(\beta \mid \mathbf{y}, \Sigma^{-1}) \sim N(\hat{\beta}, \hat{\mathbf{V}}_\beta), \quad (\text{B4})$$

$$p(\Sigma^{-1} \mid \mathbf{y}, \beta) \sim W(\hat{\mathbf{S}}^{-1}, \hat{\nu}), \quad (\text{B5})$$

where

$$\begin{aligned} \hat{\mathbf{V}}_\beta &= \left(\mathbf{V}_\beta^{-1} + \sum_{t=1}^T \mathbf{Z}'_t \Sigma^{-1} \mathbf{Z}_t \right), \\ \hat{\beta} &= \hat{\mathbf{V}}_\beta \left[\mathbf{V}_\beta^{-1} \beta_0 + \sum_{t=1}^T \mathbf{Z}'_t \Sigma^{-1} \mathbf{Z}_t \right], \end{aligned}$$

and

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

We elaborate how priors β_0 , \mathbf{V}_β , \mathbf{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 2,000 draws with the sampling scheme, discarding the first 1,000 draws and use the remaining 1,000 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 β_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 would contradict the assumption of cointegration. However, we note that our results are robust to setting this prior mean to zero (see ‘ECM prior’ case in [Figure 3a](#)).

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

$$V_{i,j}^k = \frac{\lambda^2 \sigma_i^2}{k^2 \sigma_j^2}. \quad (\text{B6})$$

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 \(2020\)](#) (see ‘optimized λ ’ case in [Figure 3a](#)).

Once $V_{i,j}^k$ in [\(B6\)](#) is specified, \mathbf{V}_β is constructed as

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

where

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

except for \mathbf{V}_1 and \mathbf{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

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

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.