

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

**Measuring the Natural Rate of Interest: A Note on Transitory
Shocks**

Kurt F. Lewis and Francisco Vazquez-Grande

2017-059

Please cite this paper as:

Lewis, Kurt F., and Francisco Vazquez-Grande (2017). "Measuring the Natural Rate of Interest: A Note on Transitory Shocks," Finance and Economics Discussion Series 2017-059. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2017.059r1>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Measuring the Natural Rate of Interest: a Note on Transitory Shocks*

Kurt F. Lewis[†] Francisco Vazquez-Grande[‡]

Current version: July 5, 2018

First version: February 15, 2017

Abstract

We present evidence that the natural rate of interest is buffeted by both permanent and transitory shocks. We establish this result by estimating a benchmark model with Bayesian methods and loose priors on the unobserved drivers of the natural rate. When subject to transitory shocks, the median estimate for the U.S. economy is more procyclical, displays a less marked secular decline, and is therefore higher following the Great Recession than most estimates in the literature.

JEL: C32, E43, E52, O40

Keywords: natural rate of interest, monetary policy, Kalman filter, pileup, trend growth

*The authors would like to thank Gianni Amisano, Giovanni Favara, Kathryn Holston, Ben Johansson, Thomas Laubach, John Stevens, John Williams and Egon Zakrajšek for many helpful comments and conversations. The authors also thank the organizers and participants of seminars at the Federal Reserve Board, Banco de España, and those of the following conferences: the Federal Reserve System Macro Committee Conference, the Midwest Macro Spring 2017 Meetings, the 2017 Asian chapter of the Econometric Society, the 2017 International Association for Applied Econometrics Annual Meeting, the 2017 Georgetown Center for Economic Research Biennial Conference, the 2017 Simposio de la Asociación de Economía Española. All remaining errors are our own. This paper was previously circulated under the title “Measuring the Natural Rate of Interest: Alternative Specifications.” The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff, the Board of Governors, or anyone else in the Federal Reserve System. Neither author has any conflicts of interest to declare.

[†]Corresponding Author: Federal Reserve Board. kurt.f.lewis@frb.gov, office: 202-452-3424, fax: 202-452-2301

[‡]Federal Reserve Board. francisco.vazquez-grande@frb.gov

1 Introduction

The natural rate of interest, or r^* , is a central concept in macroeconomics. It measures the opportunity cost of investment in an economy producing at capacity, and it is typically defined as the real interest rate consistent with stable inflation and output equating its long-term potential (Wicksell, 1936). In real business cycle models, with and without nominal or financial frictions, the natural rate of interest is time-varying and is driven by shocks to either aggregate supply or aggregate demand.

As the natural rate of interest is unobservable, empirical researchers make assumptions about the composition of r^* in order to estimate its level. For example, in their seminal contribution, Laubach and Williams (2003) model r^* as driven by two processes: one that affects aggregate supply through the growth rate of potential output (g) and another factor (z) that captures disturbances to aggregate demand, such as shocks to household preferences. They find evidence that both of these components are random walks.

In principle there is no clear theoretical justification why both drivers of r^* , g and z , need to be non-stationary processes. In fact, theory suggests that shocks to aggregate demand, such as fiscal or financial shocks, may weigh on aggregate demand only temporarily. In this paper we re-estimate a benchmark model of r^* under a looser set of prior parameter restrictions in order to let the data determine the statistical properties of its components.

Using standard Bayesian methods and loose priors on the volatility parameters, our estimates confirm earlier work suggesting that g (the growth component of r^*) is appropriately modeled as having a unit root. However, our results also suggest that the non-growth component of r^* (z) should be modeled as having transitory shocks, which stands in contrast to earlier findings. With transitory shocks to z , estimates of r^* implied by our model are markedly more volatile than those of previous studies; moreover we find the level of r^* after the Great Recession to be higher than commonly estimated in the literature.

There are methodological challenges when estimating models with latent factors. The

standard practice under maximum likelihood estimation (MLE) (e.g. Laubach and Williams, 2003; Trehan and Wu, 2007; Clark and Kozicki, 2005) is to use the median unbiased estimator of Stock and Watson (1998), which is designed to avoid the *pileup problem* (i.e. a tendency for the maximum-likelihood estimates of certain volatility parameters to be biased toward zero, see Stock, 1994). In this paper we use a Bayesian approach with uninformative priors on reasonable regions of the parameter space to mitigate the pileup problem. We note that the median unbiased estimator procedure can be viewed, from a Bayesian perspective, as very precise (and asymptotically motivated) *implicit prior beliefs* about the volatilities of the latent factors in order to mitigate the pileup problem. Bayesian methods allow us to mitigate the problem under a less restrictive structure, making visible the effects of these implicit priors on the final estimation of r^* .

The existence of transitory shocks to r^* is economically important. For central banks that use a short-term interest rate as their main policy tool, the difference between r^* and the real short-term rate provides a measure of the stance of monetary policy. Our model estimates for the U.S. economy deliver a more procyclical median estimated path of r^* over the past several recessions. In particular, our median estimate of r^* implies that the natural rate of interest plummeted during the financial crisis of 2008 but has moved back up over the past ten years to a level seen in the periods following the past several recessions. This is in contrast to the most recent point estimates in Holston, Laubach, and Williams (2017) in which r^* fell during the financial crisis and remains well below the levels estimated for earlier time periods.

Our results contrast with those of Laubach and Williams (2003) and Trehan and Wu (2007), who do not find evidence of transitory shocks to r^* . Also, under looser priors, we find that the data do provide some information on the volatility of z , in contrast to Kiley (2015), though uncertainty about this component remains significant. Our finding that transitory shocks affect r^* may also have consequences for Pescatori and Turunen (2016), who estimate r^* using Bayesian methods while attempting to decompose the drivers of z .

Our findings suggest that their use of comparatively restrictive priors, particularly on the volatility parameters and the autoregressive parameters of z , may significantly affect their results.

The paper proceeds as follows: Sections 2 and 3 briefly detail the model and estimation strategy, Section 4 discusses the results and Section 5 concludes.

2 The r^* Model

We estimate an augmented version of the Holston, Laubach, and Williams (2017) model, the most recent update to the original Laubach and Williams (2003) study. This model features an IS equation that drives the dynamics of the output gap, a “Phillips curve,” and evolutions for the unobserved components such as the level and trend-growth rate of potential GDP and the natural rate of interest. The six equations are:

$$\tilde{y}_t = a_1 \tilde{y}_{t-1} + a_2 \tilde{y}_{t-2} + \frac{a_r}{2} (\tilde{r}_{t-1} + \tilde{r}_{t-2}) + \sigma_1 \varepsilon_{1,t} \quad (2.1)$$

$$\pi_t = b_1 \pi_{t-1} + (1 - b_1) \sum_{i=2}^4 \frac{\pi_{t-i}}{3} + b_y \tilde{y}_{t-1} + \sigma_2 \varepsilon_{2,t} \quad (2.2)$$

$$r_t^* = g_t + z_t \quad (2.3)$$

$$z_t = \rho_z z_{t-1} + \sigma_3 \varepsilon_{3,t} \quad (2.4)$$

$$y_t^* = y_{t-1}^* + g_{t-1} + \sigma_4 \varepsilon_{4,t} \quad (2.5)$$

$$g_t = \mu_g (1 - \rho_g) + \rho_g g_{t-1} + \sigma_5 \varepsilon_{5,t} \quad (2.6)$$

where y is log-real GDP, y^* is log-potential GDP and $\tilde{y} \equiv y - y^*$. Similarly, $\tilde{r} \equiv r - r^*$ where r is the real short-term interest rate and r^* is the natural rate of interest.

The specification listed in equations (2.1) to (2.6) allows both g and z to be either random walks or stationary AR(1) processes. We focus on two of the nested specifications of the model, the *baseline specification* with $\rho_g = \rho_z = 1$ by assumption and an *alternative*

specification where we estimate ρ_z and assume $\rho_g = 1$.¹

3 Estimation Procedure

We estimate the model in two ways, using Bayesian methods under loose priors and by maximum likelihood as in Holston, Laubach, and Williams (2017). We see the three-stage MLE process as a way of choosing a specific (and asymptotically motivated) degenerate implicit prior over ratios of the volatilities: $\lambda_g \equiv \frac{\sigma_5}{\sigma_4}$ and $\lambda_z \equiv a_r \frac{\sigma_3}{\sigma_1}$. In order to mitigate the pileup problem, point estimates of these ratios are constructed in each of the first two stages by estimating simplifications of the model.² Those estimates λ_g and λ_z are then imposed during the maximization of the likelihood function in the final step, which reduces the implied level of parameter uncertainty.

The fully Bayesian estimation uses much less restrictive prior distributions on reasonable regions of the parameter space, as discussed in DeJong and Whiteman (1993), Primiceri (2005), and Kim and Kim (2018), to avoid the pileup problem. Formally, after specifying the priors, we construct the likelihood from the linear-Gaussian filter output and use the random-walk Metropolis-Hastings algorithm to generate draws from the posterior distributions of the model parameters. Each draw of the parameters from the posterior distribution implies a sampled path for the unobserved variables, including r_t^* , as in Carter and Kohn (1994) and Frühwirth-Schnatter (1994).³

¹We have also examined the other permutations of these settings. For example, we find that the data supports the assumption in Laubach and Williams (2003) that g is a random walk. Results from a model in which both z and g are estimated AR(1) processes are included in the online appendix.

²For details on the specifics of the three-stage procedure, see Laubach and Williams (2003).

³Textbook treatments of this approach can be found in Geweke (2005) and Herbst and Schorfheide (2015). The online appendix contains the state space representation of the model used in the estimation as well as additional technical details and sources of information about the data.

3.1 Prior Distributions

The marginal prior distributions of the parameters are given in Table 1. These priors were chosen with a mind toward being minimally informative.⁴ The priors on the standard deviations of the unobserved shock processes play a critical role and we choose marginal priors to be uniform between 0 and 5, in contrast to the common usage of inverse gamma priors in the literature. While inverse gamma distributions have a domain that runs along the positive real line, their mass is concentrated in a fairly small area, and are therefore relatively informative in the context of this model, as demonstrated by Kiley (2015). To avoid the pileup problem we restrict λ_g and λ_z to take values in $[0.01, 5]$, which represents much less a priori certainty than previous studies.⁵ Regarding the prior of ρ_z , the choice of $N(1, \frac{1}{2})$ is meant to reflect the a priori belief that the z processes is highly persistent and could have a unit root.⁶

4 Results

Figure 1 displays the higher parameter uncertainty revealed by Bayesian estimation under priors which are looser than the implicit priors used in the MLE procedure. The MLE procedure fixes the ratios of the relevant volatility parameters (λ_g and λ_z) to specific values displayed by the red lines. Under our Bayesian estimation these ratios are free to take values between 0.01 and 5. The posterior distributions have modes which are not far from the values used in MLE, but their standard deviations are considerable. Accounting for this parameter uncertainty has implications for r^* , even in the baseline model.⁷

⁴See the online appendix for additional detail.

⁵The implied prior distributions for λ_g and λ_z (the results of marginal priors of their components and the restriction discussed above) along with the priors from Pescatori and Turunen (2016) and the MLE values are shown in the online appendix.

⁶As noted by Sims (1988), the shape of the likelihood function is not changed by the inclusion of unit or explosive roots, so there is no need to truncate the distribution centered at one.

⁷As a check, we estimated a version of the model that imposes, via degenerate priors, the MLE point-estimates for λ_g and λ_z within the Bayesian estimation. When we did this, we recovered the same median path of r^* as in the MLE estimation.

Figure 2 shows the effects of more completely incorporating parameter uncertainty on the median estimate of r^* . The relaxation of the λ_g and λ_z restrictions imposed by the MLE methodology generates a median path of the natural rate of interest which is more volatile and procyclical than its MLE counterpart. We note that the level of uncertainty about r^* is considerable.

As seen in Figure 3, the majority of the uncertainty about r^* comes from the non-growth component, z , the uncertainty of which we now more fully appreciate. While the priors on the parameters of both g and z are identical, the relative magnitudes of the credible sets shown in the figure indicate that the likelihood function generates considerably more concentration of posterior mass for the parameters of g relative to those of z .

As shown by panel (a) of Table 2, the increased uncertainty about z comes predominantly from the wider range of plausible values for the volatility parameter governing its shocks, σ_3 .⁸ While the peak of the posterior distribution of σ_3 in the baseline specification is near the point estimate of the parameter under MLE, the distribution is skewed and the standard deviation is fairly large. Under MLE, the process required to avoid the pileup problem via a point estimate of λ_z necessarily results in tighter restrictions on the potential values of σ_3 . Bayesian estimation allows the pileup problem to be mitigated with less restrictions, revealing that our uncertainty about z is large. Taking all this into account, we next reexamine a key finding in the earlier literature: that z is a random walk.

Using the Savage-Dickey density ratio (SDDR) we find substantial statistical evidence that the data prefers not to assume that z is a random walk. Dickey (1971) constructs an exact Bayes factor comparing two nested models that differ only insofar as one model (here, the baseline specification) fixes a model parameter at a specific value ($\rho_z = 1$), while the other model (the alternative specification) estimates it. In such a case, the Bayes factor can

⁸Importantly, our posterior distributions for the volatility parameters σ_3 and σ_5 , shown in the online appendix, still show no signs of pileup.

be written in terms of the output of only the *unrestricted* model:

$$B_{alt,baseline} = \frac{p_{alt}(\rho_z = 1)}{p_{alt}(\rho_z = 1|Y)}$$

where $p_{alt}(\rho_z = 1|Y)$ is the value of the pdf of the marginal posterior distribution for ρ_z under the alternative specification at the point $\rho_z = 1$, and $p_{alt}(\rho_z = 1)$ is the value of the pdf of the prior on ρ_z evaluated at 1, also under the alternative specification.

The SDDR provides a very intuitive signal: when the weight of the marginal posterior goes up relative to the prior, the data supports the assumption in the restricted model, and vice-versa. As can be seen in Figure 4, the weight of the marginal posterior on $\rho_z = 1$ is considerably lower than it is in the prior. The ratio, and thus the Bayes Factor in favor of the alternative specification is 10.2, which according to Jefferys (1961), is “substantial” evidence in favor non-permanent shocks to z . Kass and Raftery (1995), who develop their own scale for Bayes factors label this as “positive” evidence in favor of the alternative specification.⁹ This result is robust to alternative prior specifications for ρ_z .

Panel (b) of Table 2 shows the Bayes factor in favor of either model (constructed using the SDDR), along with other model comparison information from the two specifications under both Bayesian and maximum likelihood estimation. We see that, in this model, the choice of procedure imposed to deal with the pileup problem can flip model selection. The log marginal likelihood value, constructed using the Newton and Raftery (1994) methodology, finds values of -533 and -526 for the baseline and alternative specifications, respectively, also evidence generally supportive of choosing the alternative model. While our findings about z contradict those of Laubach and Williams (2003) and Trehan and Wu (2007), we confirm that the divergence between our results and theirs is based on the more restrictive solution to the pileup problem used in MLE of this model. Replicating that three-step process, we found that the log-likelihood value of the baseline model at the maximum likelihood estimates is

⁹In both ranking systems, this grade of evidence is considered the second level, with the next level labeled “strong” and further levels labeled “very strong” or “decisive.”

-518, while it is -517 under the alternative specification. The Bayesian information criterium favors the baseline model over the alternative.

We find economic appeal in a z process subject to persistent, but transitory, shocks because of its behavior in the period around, and following, the crisis. Under the baseline specification (and in the MLE results) there was a fairly sudden decline in z , and thus r^* , in 2008. While many slow-moving phenomena could be invoked to bolster a strong prior belief that z should be a random walk, these proposals need to align with the relatively sharp movement in that time period. Figure 5 shows the median path of z under the alternative specification and Figure 6 shows the corresponding estimate for r^* when z is subject to transitory shocks. In addition to the higher volatility and the much larger impact of the Great Recession on the level of the median path of r^* , the post-crisis profile of r^* is very different than that of the baseline model, largely driven by the different dynamics in z . Most notably, following the sharp dip in the Great Recession, the median path of r^* has generally trended in a positive direction, though it remained broadly below zero for several years following the crisis. This is in contrast to the estimates from Holston, Laubach, and Williams (2017) and others, where the natural rate descends in the 21st century and remains at historically low levels through the end of the sample.¹⁰ The change to the specification for z appears to have had very little effect on the estimate of g , a component of our final discussion below.

4.1 Output Gap Implications

Our statistically preferred specification for z may have implications for economically important objects such as the output gap. Figure 7 shows that while our baseline and alternative estimates of r^* are different from those under MLE, our estimates of the output gap are much more in concert. Figure 7 also includes the estimate of the output gap available from the

¹⁰A related concept, as discussed in Del Negro et al. (2017), is an explicitly long-run, rather than medium term, r^* . A short discussion of the long-run r^* from the alternative specification is included in the online appendix.

Congressional Budget Office (CBO) and from the model of Pescatori and Turunen (2016), who take signal from the CBO.

The CBO estimate may be considered an external check on model output as it can represent a benchmark for judging our estimate of economic slack. While our Bayesian and MLE estimates are similar to the CBO's measure for much of the sample prior to 2000, the estimates diverge at that point, with our measures indicating a higher level of resource utilization in recent years. Figure 8 shows that these post-2000 differences are not the result of dramatically different views of potential output growth by the different models over that period. Rather, the figure shows that the recent divergence in output gap estimates is driven by the CBO's high estimate of potential output growth during a brief period in the late 1990s. This led to a shift in the estimated level of potential GDP, which results in a CBO output gap estimate which ends our sample (mid-2017) at a negative level.

Figure 8 shows a remarkable similarity across the estimates of potential output growth from the five sources. All of the model-based estimates lie well within the 90% credible set from the baseline model and the CBO estimate lies within the set for the majority of the time as well. Additionally, all the models appear to provide similar support for the idea that there has been a secular decline in potential output growth in 21st century, as discussed by Summers (2014), Eggertsson, Mehrotra, and Summers (2016) and others.

5 Conclusion

This paper re-estimates a benchmark model under looser prior assumptions and finds different median estimates of the natural rate of interest. We find that a more complete picture of the parameter uncertainty in the model results in a higher median estimate of the conditional volatility of r^* . We also find that the data prefers r^* to be affected by transitory shocks, in contrast to previous studies. Acknowledging the potential for persistent, but transitory, shocks to r^* will likely help to shape the search for its economic drivers.

6 Tables and Figures

Name	Domain	Density	Parameter 1	Parameter 2
a_1	\mathbb{R}	Normal	0	2
a_2	\mathbb{R}	Normal	0	2
a_r	\mathbb{R}^-	Normal	0	2
b_1	$[0, 1]$	Uniform	0	1
b_Y	\mathbb{R}^+	Normal	0	2
ρ_z	\mathbb{R}^+	Normal	1	$\frac{1}{2}$
σ_1	$[0, 5]$	Uniform	0	5
σ_2	$[0, 5]$	Uniform	0	5
σ_3	$[0, 5]$	Uniform	0	5
σ_4	$[0, 5]$	Uniform	0	5
σ_5	$[0, 5]$	Uniform	0	5

Table 1: The table presents the marginal prior distributions under the individual model parameters for the alternative specification. The prior distribution parameters are the mean (1) and standard deviation (2) for those with Normal distributions and the end-points of the domain interval for uniform distribution. The domains of a_r , b_Y , ρ_g and ρ_z are truncations of the standard form of the prior density. In the baseline specification the prior distribution of ρ_z was set to be degenerate at $\rho_z \equiv 1$. We restrict λ_g and λ_z to take values in $[0.01, 5]$, additional discussion of the prior distributions is included in the online appendix.

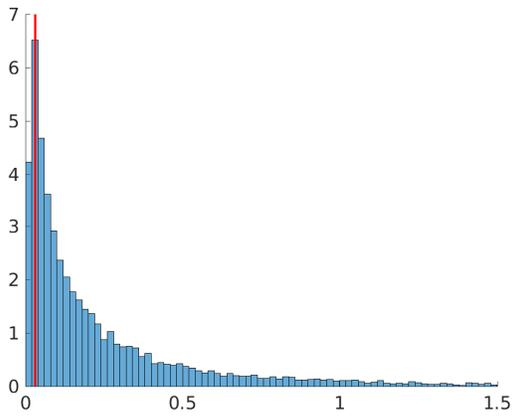
	Bayesian		MLE	
	Baseline	Alternative	Baseline	Alternative
a_1	1.251 [0.97,1.51]	1.270 [0.84,1.52]	1.531 [1.36,1.70]	1.530 [1.36,1.70]
a_2	-0.364 [-0.58,-0.07]	-0.348 [-0.59,0.05]	-0.589 [-0.76,-0.42]	-0.587 [-0.76,-0.41]
a_r	-0.113 [-0.19,-0.06]	-0.093 [-0.18,-0.06]	-0.070 [-0.10,-0.04]	-0.067 [-0.10,-0.04]
b_1	0.682 [0.57,0.79]	0.665 [0.56,0.78]	0.671 [0.60,0.74]	0.670 [0.60,0.74]
b_Y	0.051 [0.03,0.13]	0.071 [0.04,0.15]	0.077 [0.04,0.12]	0.079 [0.04,0.12]
σ_1	0.412 [0.11,0.66]	0.279 [0.08,0.57]	0.355 [0.21,0.50]	0.365 [0.21,0.52]
σ_2	0.794 [0.74,0.87]	0.795 [0.74,0.86]	0.791 [0.75,0.83]	0.791 [0.75,0.83]
σ_3	0.149 [0.07,1.69]	1.755 [0.67,3.95]	0.160 [0.10,0.23]	0.172 [0.10,0.25]
σ_4	0.564 [0.1,0.64]	0.580 [0.25,0.65]	0.571 [0.48,0.66]	0.567 [0.47,0.66]
σ_5	0.036 [0.02,0.13]	0.035 [0.02,0.11]	0.030 [0.02,0.03]	0.030 [0.02,0.03]
ρ_z	1*	0.789 [0.31,0.89]	1*	0.916 [0.77,1.06]

(a) Estimation of the Parameters

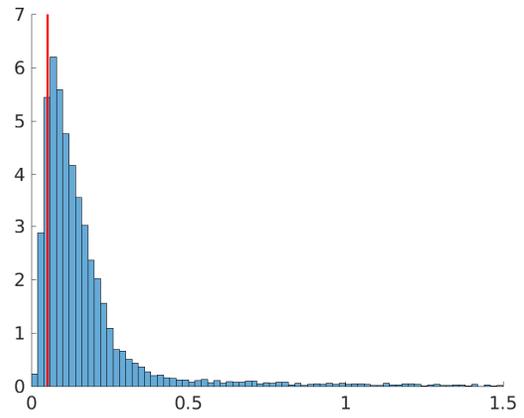
	Bayesian		MLE	
	Log Marg. Like.	BF	Log. Like.	BIC
Baseline	-533	0.1	-518	1088
Alternative	-526	10.2	-517	1093

(b) Model Comparison Under Bayesian and MLE Methods

Table 2: Panel (a) shows the modes of the posterior distributions of the model parameters from the baseline and alternative specifications, along with the MLE. The numbers in brackets represent the 90% credible set from the posterior distributions of the parameters for the Bayesian estimation, and the 90% asymptotic confidence interval for the MLE, the standard errors come from the third estimation stage. Panel (b) shows the model comparison statistics under Bayesian and MLE methods. The Log Marginal Likelihood values are built using the Newton and Raftery (1994) methodology, and the Bayes Factor (BF) in favor of a given model is built using the Savage-Dickey density ratio of Dickey (1971). The Bayesian Information Criteria (BIC) is reported for the two MLE estimates. *In the baseline specification under both estimation methods ρ_z is set to one.



Posterior λ_z (Baseline)



Posterior λ_g (Baseline)

Figure 1: Posterior Distributions of λ_z and λ_g Under Baseline Specification

NOTES: The blue bars conform the histogram of the posterior distribution of λ_z and λ_g . The red lines stand at the median unbiased estimates used in MLE

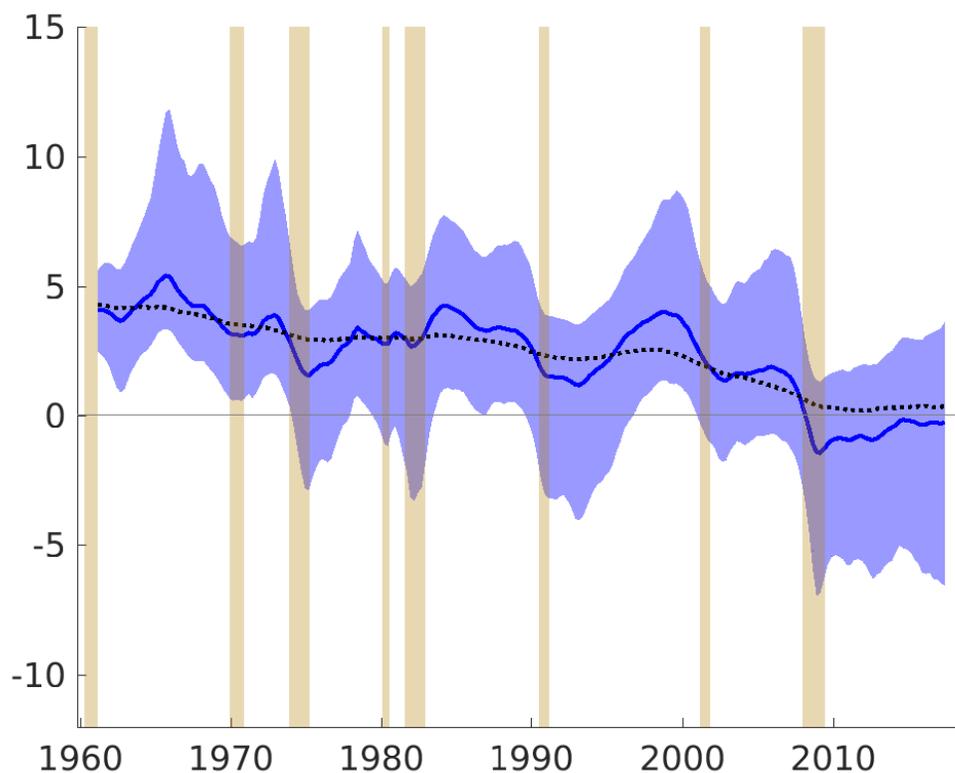


Figure 2: r^* Path (Baseline Model)

NOTES: The path of r^* under the baseline model when $\rho_g = \rho_z = 1$. The solid blue line shows the median path of the smoothed estimate and the blue-shaded area shows the 90% credible set of the estimated path. The black dashed line plots the equivalent series under MLE. The vertical shaded bars represent NBER-dated recessions. For reference, the standard error from the MLE estimate of r^* averages 1.1 percentage points.

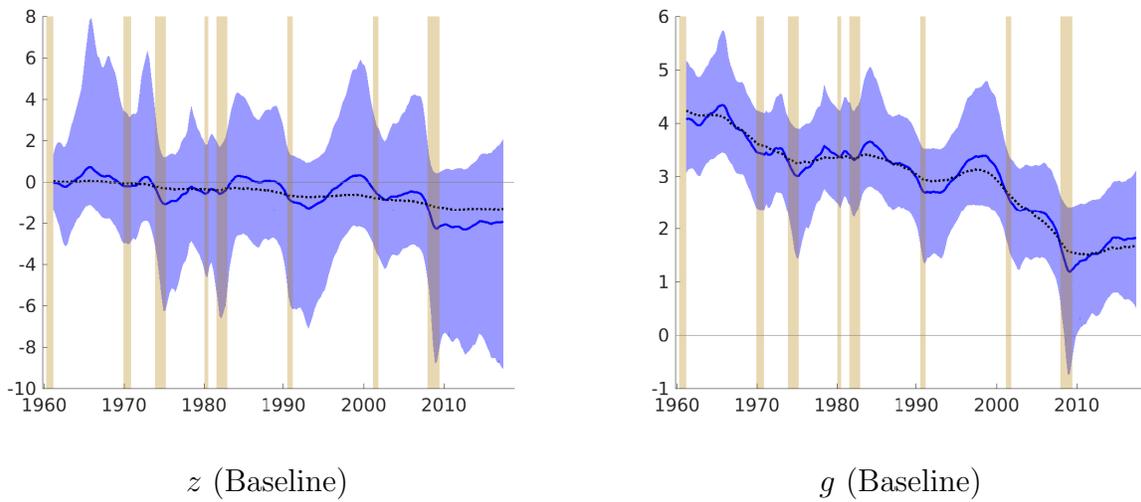


Figure 3: The Paths of the Components of r^* Under Baseline Specification

NOTES: The paths of the components of r^* under the baseline specification. The blue line is the median estimate, the black dotted line is the equivalent series under MLE, the shaded area represents the 90% credible set.

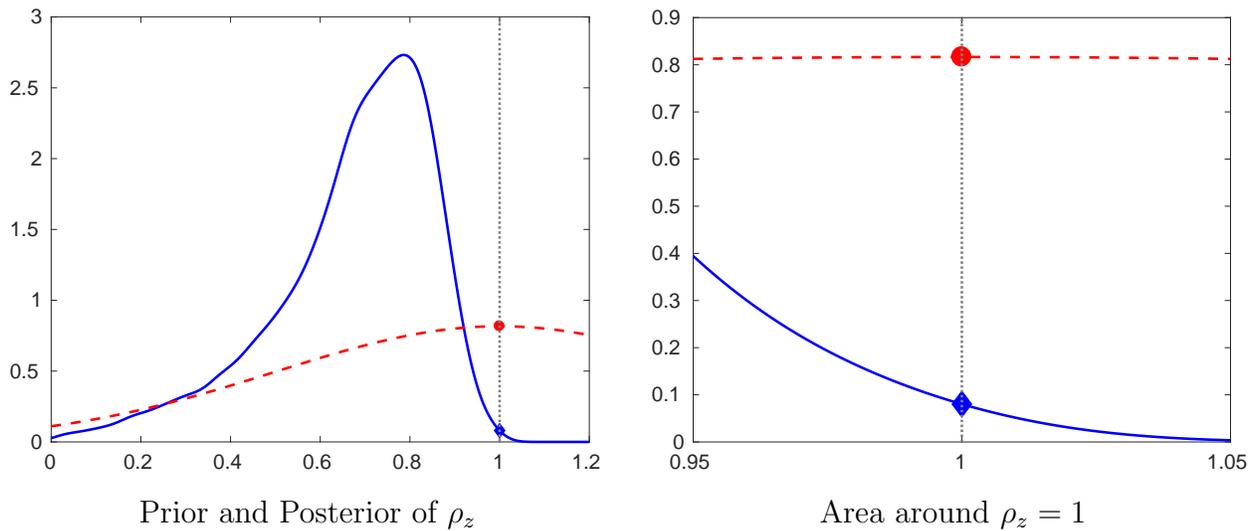


Figure 4: An Illustration of the Savage-Dickey Density Ratio

NOTES: The panel on the left shows the marginal posterior distribution of ρ_z under the alternative specification (the solid blue line) and the prior distribution over the same interval (the dashed red line). The vertical gray dashed line indicated where $\rho_z = 1$. The panel on the right shows the same distributions expanded around the region where $\rho_z = 1$. The red circle indicates the pdf value for the prior at $\rho_z = 1$, and the blue diamond indicates the pdf value for the marginal posterior at 1.

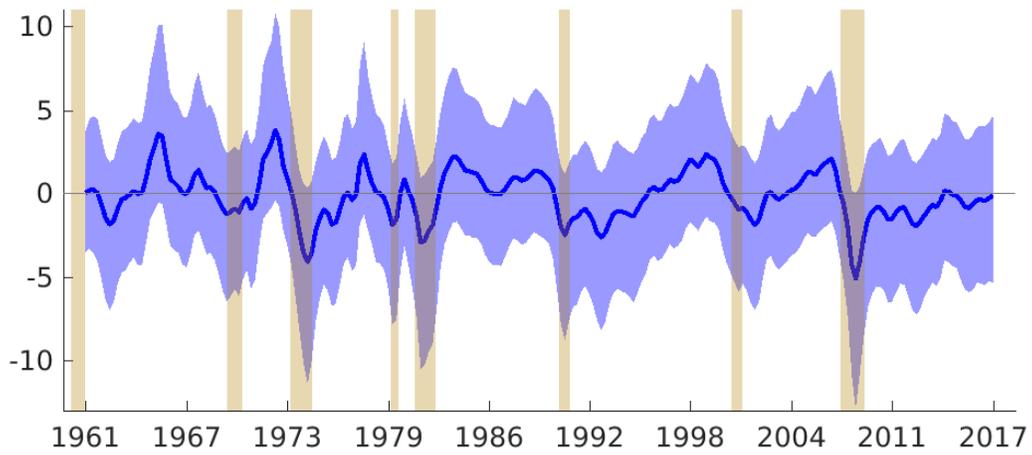


Figure 5: z (Alternative)

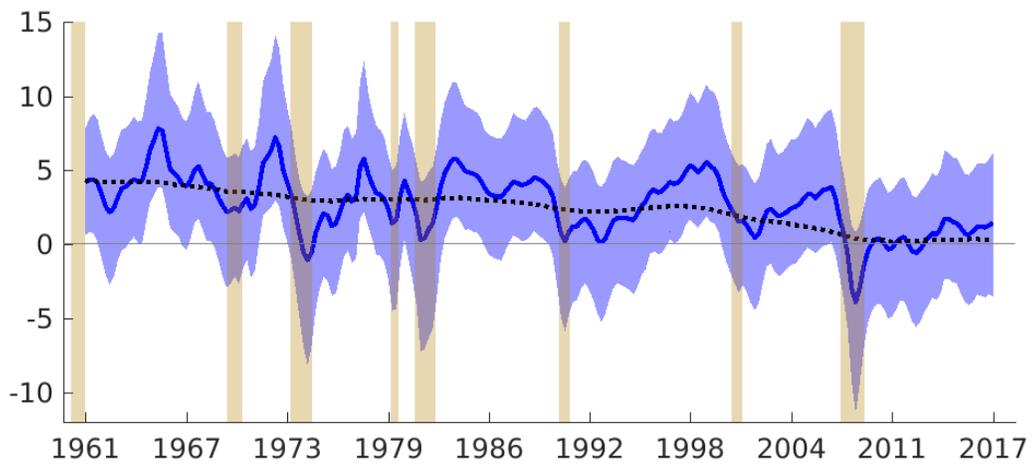


Figure 6: r^* Path (Alternative)

NOTES: The paths of z and r^* under the alternative model when ρ_z is estimated. The solid blue line shows the median path of the smoothed (two-sided) estimate and the blue-shaded areas represent the 90% credible sets. The vertical shaded bars represent NBER-dated recessions.

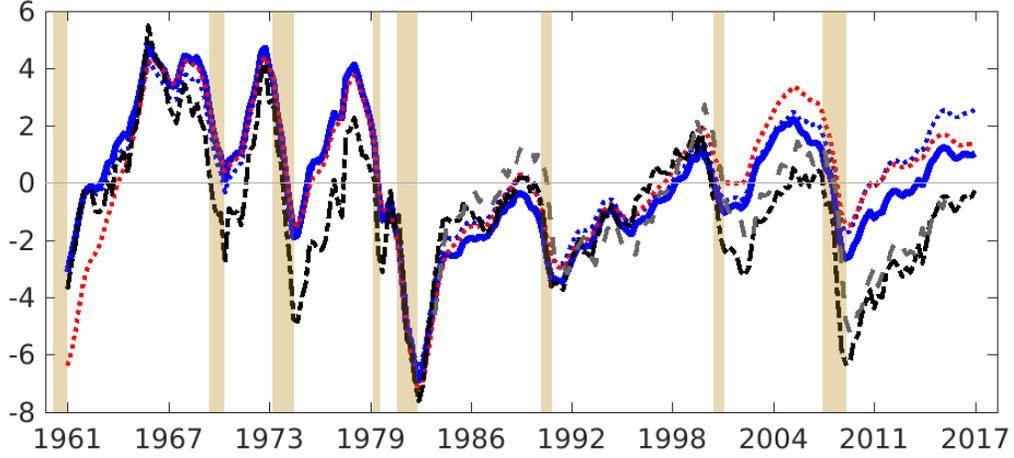


Figure 7: Estimates of the Output Gap

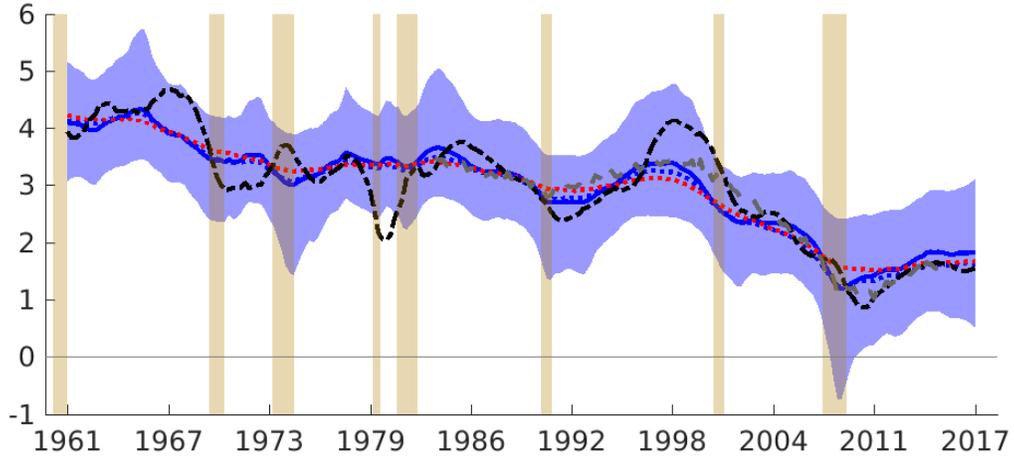


Figure 8: Estimates of Potential Output Growth

NOTES: Estimates of the output gap, Figure 7, and the corresponding level of potential output growth, Figure 8 are shown along with comparable measures from other sources. The solid blue lines show the estimates from the baseline specification, the dotted blue lines show the alternative specification. The red dotted lines show the equivalent estimates under MLE, as in Holston, Laubach, and Williams (2017). The black dashed lines are from the estimates provided by the Congressional Budget Office (CBO), and the gray dash-dotted lines are the estimates from Pescatori and Turunen (2016) (available 1983-2015). The blue shaded area in Figure 8 represents the 90% credible set of the growth rate of potential output in the baseline specification.

References

- Carter, C. K. and R. Kohn (1994). On gibbs sampling for state space models. *Biometrika* 81(3), 541–553.
- Carter, C. K. and R. Kohn (1996). Markov chain monte carlo in conditionally gaussian state space models. *Biometrika* 83(3), 589–601.
- Clark, T. E. and S. Kozicki (2005). Estimating equilibrium real interest rates in real time. *The North American Journal of Economics and Finance* 16(3), 395 – 413. Real-time data and monetary policy.
- DeJong, D. N. and C. H. Whiteman (1993). Estimating moving average parameters: Classical pileups and bayesian posteriors. *Journal of Business & Economic Statistics* 11(3), 311–317.
- Del Negro, M., D. Giannone, M. Giannoni, and A. Tambalotti (2017). Safety, liquidity, and the natural rate of interest. Brookings papers on economic activity, Brookings.
- Dickey, J. M. (1971). The weighted likelihood ratio, linear hypotheses on normal location parameters. *The Annals of Mathematical Statistics* 42, 204–223.
- Eggertsson, G. B., N. R. Mehrotra, and L. H. Summers (2016, May). Secular Stagnation in the Open Economy. *American Economic Review* 106(5), 503–507.
- Frühwirth-Schnatter, S. (1994, December). Applied state space modelling of non-gaussian time series using integration-based kalman filtering. *Statistics and Computing* 4(4), 259–269.
- Geweke, J. (2005). *Contemporary Bayesian Econometrics and Statistics*. Wiley Series in Probability and Statistics. Wiley.
- Herbst, E. and F. Schorfheide (2015). *Bayesian Estimation of DSGE Models*. Princeton: Princeton University Press.
- Holston, K., T. Laubach, and J. C. Williams (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics* 108, S59 – S75. 39th Annual NBER International Seminar on Macroeconomics.
- Jefferys, H. (1961). *Theory of Probability* (3 ed.). Oxford University Press.
- Kass, R. E. and A. E. Raftery (1995). Bayes factors. *Journal of the American Statistical Association* 90(430), 773–795.
- Kiley, M. T. (2015, August). What Can the Data Tell Us About the Equilibrium Real Interest Rate? Finance and Economics Discussion Series 2015-77, Board of Governors of the Federal Reserve System (U.S.).

- Kim, C.-J. and J. Kim (2018). Trend-cycle decompositions of real gdp revisited: Classical and bayesian perspectives on an unsolved puzzle. <https://ssrn.com/abstract=2883438>.
- Laubach, T. and J. C. Williams (2003). Measuring the natural rate of interest. *Review of Economics and Statistics* 85(4), 1063–1070.
- Newton, M. A. and A. E. Raftery (1994). Approximate bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society. Series B (Methodological)* 56(1), 3–48.
- Pescatori, A. and J. Turunen (2016). Lower for longer: Neutral rate in the u.s. *IMF Economic Review* 64(4), 708–731.
- Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72(3), 821–852.
- Sims, C. A. (1988). Bayesian skepticism on unit root econometrics. *Journal of Economic Dynamics and Control* 12(2), 463 – 474.
- Stock, J. H. (1994). Unit roots, structural breaks and trends. In *Chapter 46*, Volume 4 of *Handbook of Econometrics*, pp. 2739 – 2841. Elsevier.
- Stock, J. H. and M. W. Watson (1998, March). Median unbiased estimation of coefficient variance in a time-varying parameter model. *Journal of the American Statistical Association* 93(441), 349–358.
- Summers, L. H. (2014, April). U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound. *Business Economics* 49(2), 65–73.
- Trehan, B. and T. Wu (2007). Time-varying equilibrium real rates and monetary policy analysis. *Journal of Economic Dynamics and Control* 31(5), 1584 – 1609.
- Wicksell, K. (1936). *Interest and Prices*. Macmillan, London. Trans. of 1898 edition by R.F. Kahn.

A Online Appendix Material

A.1 State Space Formulation

Based on the system of equations in section 2 of the paper, substituting the formula for r_t^* into the output gap equation and expanding, we can come to a version of these equations that can be expressed in the traditional observation/transition equation style of the standard state space model. Following some algebraic manipulation, these equations are given as follows. First, the observation equations on real GDP and inflation.

$$\begin{aligned}
 y_t = & y_t^* - a_1 y_{t-1}^* - a_2 y_{t-2}^* - 2a_r \rho_g g_{t-1} - 2a_r \rho_g g_{t-2} \\
 & - \frac{a_r}{2} \rho_z z_{t-1} - \frac{a_r}{2} \rho_z z_{t-2} - 4a_r \mu_g (1 - \rho_g) + a_1 y_{t-1} \\
 & + a_2 y_{t-2} + \frac{a_r}{2} r_{t-1} + \frac{a_r}{2} r_{t-2} + \sigma_1 \varepsilon_{1,t}
 \end{aligned} \tag{A.1}$$

$$\pi_t = -b_Y y_{t-1}^* + b_Y y_{t-1} + b_1 \pi_{t-1} + (1 - b_1) \sum_{i=2}^4 \frac{\pi_{t-i}}{3} + \sigma_2 \varepsilon_{2,t} \tag{A.2}$$

Then, the transition equations for unobserved potential real GDP, its growth rate, and the z process.

$$y_t^* = y_{t-1}^* + \mu_g (1 - \rho_g) + \rho_g g_{t-2} + \sigma_5 \varepsilon_{5,t-1} + \sigma_4 \varepsilon_{4,t} \tag{A.3}$$

$$z_{t-1} = \rho_z z_{t-2} + \sigma_3 \varepsilon_{3,t-1} \tag{A.4}$$

$$g_{t-1} = \rho_g g_{t-2} + \mu_g (1 - \rho_g) + \sigma_5 \varepsilon_{5,t-1} \tag{A.5}$$

These equations can be represented in state space form using the standard structure:

$$s_t = A s_{t-1} + B u_t + C w_t \tag{A.6}$$

$$x_t = D s_t + F u_t + G w_t \tag{A.7}$$

where:

$$s_t = \begin{bmatrix} y_t^* \\ y_{t-1}^* \\ y_{t-2}^* \\ g_{t-1} \\ g_{t-2} \\ z_{t-1} \\ z_{t-2} \end{bmatrix}, \quad x_t = \begin{bmatrix} y_t \\ \pi_t \end{bmatrix}, \quad u_t = \begin{bmatrix} 1 \\ y_{t-1} \\ y_{t-2} \\ r_{t-1} \\ r_{t-2} \\ \pi_{t-1} \\ \sum_{i=2}^4 \frac{\pi_{t-i}}{3} \end{bmatrix}, \quad w_t = \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \varepsilon_{3,t-1} \\ \varepsilon_{4,t} \\ \varepsilon_{5,t-1} \end{bmatrix},$$

and

$$A = \begin{bmatrix} 1 & 0 & 0 & \rho_g & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \rho_g & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho_z & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} \mu_g(1 - \rho_g) & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \mu_g(1 - \rho_g) & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$C = \begin{bmatrix} 0 & 0 & 0 & \sigma_4 & \sigma_5 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_5 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad D = \begin{bmatrix} 1 & -a_1 & -a_2 & -2a_r\rho_g & -2a_r\rho_g & -\frac{a_r\rho_z}{2} & -\frac{a_r\rho_z}{2} \\ 0 & -b_Y & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$F = \begin{bmatrix} -4a_r(1 - \rho_g)\mu_g & a_1 & a_2 & \frac{a_r}{2} & \frac{a_r}{2} & 0 & 0 \\ 0 & b_Y & 0 & 0 & 0 & b_1 & (1 - b_1) \end{bmatrix}, \quad G = \begin{bmatrix} \sigma_1 & 0 & 0 & 0 & 0 \\ 0 & \sigma_2 & 0 & 0 & 0 \end{bmatrix}$$

The ε 's are all assumed to be *i.i.d.* $N(0, 1)$ variables, with the standard deviation of the processes controlled by the σ_i 's.

A.2 Data

The data used in this analysis is the same as the US data used in Holston, Laubach, and Williams (2017), henceforth HLW, and it is transformed in the same way. See the data appendix in HLW for additional specifics on obtaining the data. Real GDP data are obtained from the BEA, inflation is calculated as the annualized quarterly growth rate of the price index for personal consumption expenditures excluding food and energy (commonly referred to as “core PCE inflation”). We follow HLW in using a 4-quarter moving average of inflation in period t as a proxy for inflation expectations in that period. The short-term interest rate is the annualized nominal effective federal funds rate, where the quarterly value is constructed as the average of the monthly values. Prior to 1965, we use the Federal Reserve Bank of New York’s discount rate.

A.3 Bayesian estimation details

Some additional details:

- Restrictions on parameters (primarily inherited from HLW):
 - We enforce that a_r be negative and b_Y positive, following HLW in using the actual restrictions $a_r < -0.0025$ and $b_Y > 0.025$.
 - As the sum of the coefficients on lags of inflation must sum to 1, we restrict b_1 to be between 0 and 1.
 - Because of our expectation of a positive autocovariance for both g_t and z_t in the event of stationarity, we restrict ρ_g and ρ_z to be positive.

The initialization for the states was duplicated from the process used in HLW: the initial values for potential output y^* were constructed by HP filtering the GDP series beginning in 1960Q1, then using the trend component of the filtered output for the observations just before the beginning of the data used in the estimation (1960 Q2, Q3 and Q4); the initial values for g were the changes of that trend component in the second half of 1960. The initial values for z were set to zero, as in HLW.

The estimation is performed in MATLAB using our own code to implement the random-walk Metropolis-Hastings algorithm (see, e.g., Herbst and Schorfheide, 2015). The filter code was written to execute the forward-filter, backward sample methodology of Frühwirth-Schnatter (1994) and Carter and Kohn (1994, 1996) to obtain samples of the unobserved states. We used a burn-in period of 250,000 draws before accepting every tenth draw for a total of 500 kept draws from each of 20 chains, for a total sample of 10,000 draws from the posterior distribution.

A.4 A Flexible Specification Where g and z Are Both AR(1)

Another specification which we tested was to allow both z and g to be estimated as AR(1) processes without forcing either to be a random walk. Allowing ρ_g and μ_g to be estimated along with ρ_z did not dramatically alter the median path of r^* that was estimated as the alternative specification in the paper, as can be seen below in Figure A.1. This is because the posterior distributions provide considerable evidence that the persistence parameter, ρ_g , can be reasonably assumed to be one for the purposes of parsimony, see Figures A.2 and A.3. In fact, when we conduct the same Savage-Dickey density ratio test on ρ_g in this specification that we conduct on ρ_z in the alternative specification of the main text, we find that the data adds weight to the posterior at the value $\rho_g = 1$, see Figure A.4. The posterior distributions are described in Table A.1 and are generally similar to the alternative specification except for the new parameters of g . The Newton and Raftery (1994) log marginal likelihood value is -526, the same as that of the alternative specification.

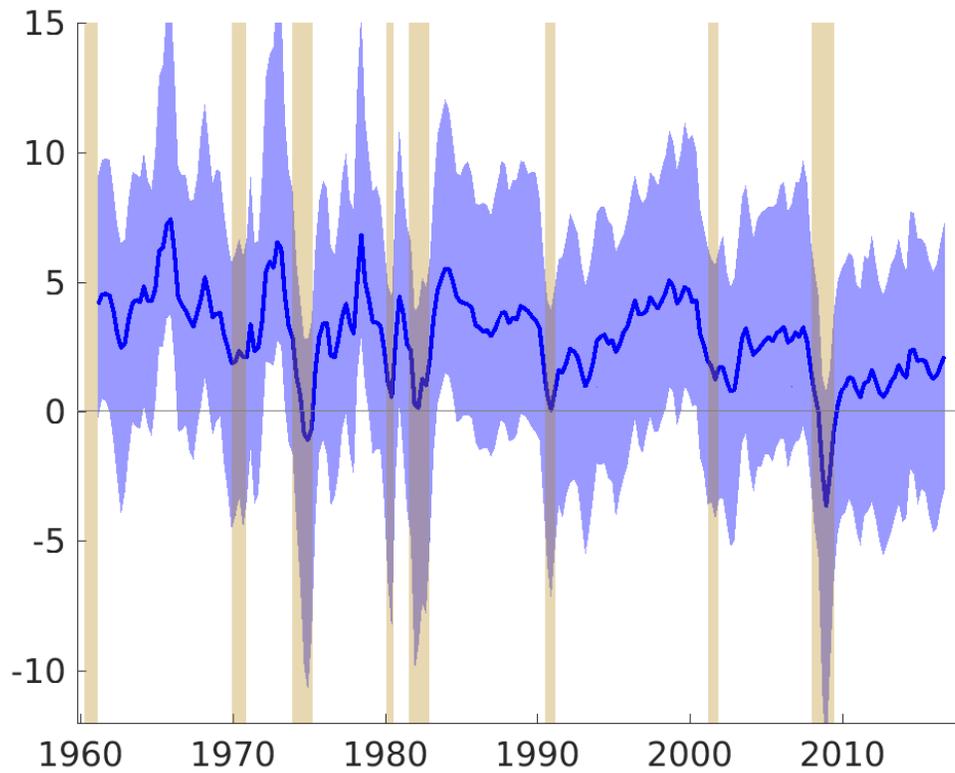


Figure A.1: r^* Path

NOTES: The path of r^* under the a specification in which both g and z are estimated as AR(1) processes. The solid blue line shows the median path of the smoothed (two-sided) estimate and the blue-shaded area is bounded by the 5th and 95th percentiles of the estimated path. The vertical shaded bars represent NBER-dated recessions.

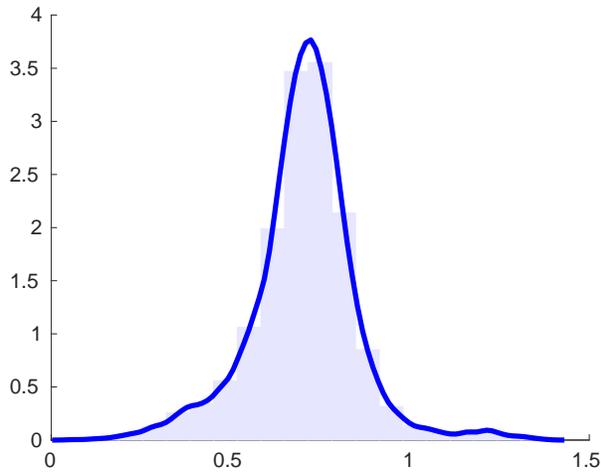


Figure A.2: μ_g from a flexible AR(1) specification for g

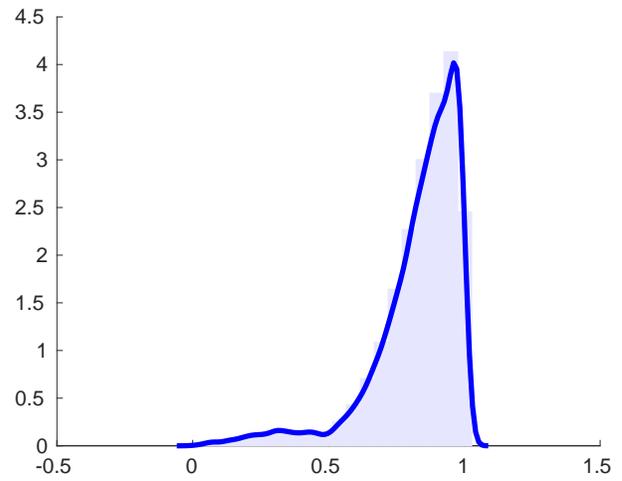


Figure A.3: ρ_g from a flexible AR(1) specification for g

NOTES: The posterior distributions for the parameters of the AR(1) process for g in a specification in which both g and z are allowed to be freely estimated.

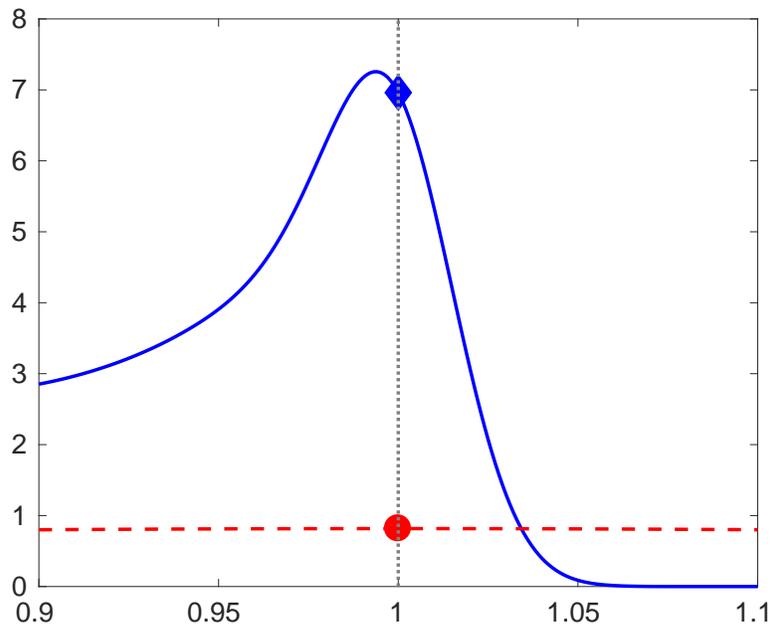


Figure A.4: SDDR for ρ_g

NOTES: The illustration of the Savage-Dickey density ratio for ρ_g in a specification in which g and z were both estimated AR(1) processes, accounting for the pileup problem via priors. Evidence suggests that the assumption that $\rho_g = 1$ is valid.

Fully Flexible Specificaton			
	10 th Perc.	Median	90 th Perc.
a_1	0.85	1.17	1.43
a_2	-0.53	-0.30	0.00
a_r	-0.14	-0.08	-0.04
b_1	0.59	0.68	0.76
b_Y	0.05	0.10	0.17
σ_1	0.05	0.24	0.49
σ_2	0.75	0.80	0.86
σ_3	0.62	2.23	5.24
σ_4	0.31	0.55	0.63
σ_5	0.06	0.17	0.35
ρ_g	0.64	0.87	0.98
ρ_z	0.15	0.53	0.81
μ_g	0.53	0.71	0.85
λ_g	0.11	0.33	1.01
λ_z	0.16	0.80	4.73

Table A.1: Information from posterior distributions of the parameters from the fully flexible specification

A.5 Prior Distributions of λ_g and λ_z and Pile-Up

The prior distributions for the σ_i 's were chosen to reflect the high degree of uncertainty about the volatility of the hidden processes. Using uniform distributions gave us a simple way to allow for significant mass across potentially larger values without significantly underweighting the region close to zero. We use restrictions on λ_g and λ_z , requiring that they have properties that limit the risk of pileup. Indeed, Figures A.5 and A.6 compare our implied prior distributions for λ_g and λ_z to those used by Holston et al. (2017) and Pescatori and Turunen (2016). Our priors on λ_g and λ_z are less informative than others used in the literature, this is especially true for the case of λ_z , where inverse gamma distributions with means near the HLW point estimates actually place more mass to the left of that estimate, very close to zero.

Figures A.7 and A.8 show that the unobserved volatility parameters display no signs of pileup. The signals from our analysis line up with a finding from Clark and Kozicki (2005) that λ_z and λ_g may be higher than estimated by Laubach and Williams (2003).

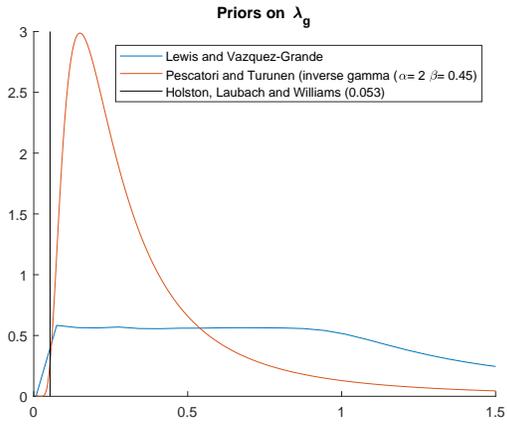


Figure A.5: Priors of λ_g

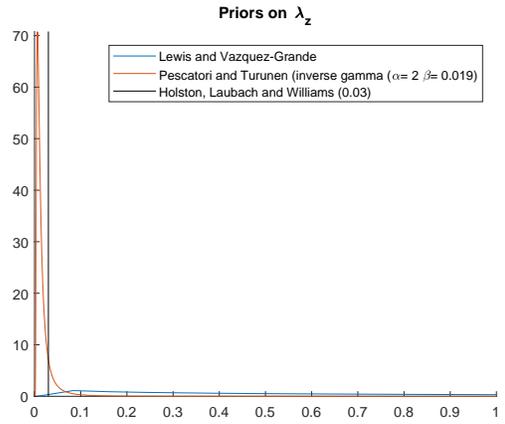


Figure A.6: Priors of λ_z

NOTES: These inverse gamma prior parameters are consistent with the moments reported on Pescatori and Turunen (2016).

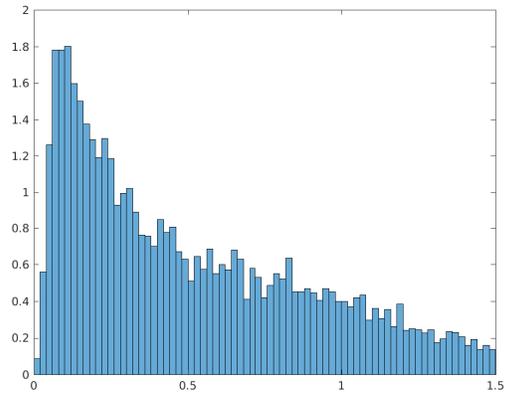


Figure A.7: Posterior Distribution of σ_3

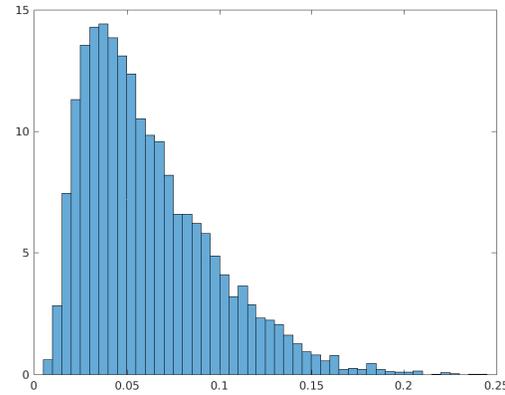


Figure A.8: Posterior Distribution of σ_5

NOTES: Posterior Distributions of unobserved volatilities where the “pileup problem” was a concern show no evidence of pileup.

A.6 The Long-run natural rate

An important difference between the baseline and alternative specifications is that in the baseline specification r^* is, by construction, a long-run object. Having introduced transitory shocks in the alternative specification, we will need to transform our new measure of r^* to align it better for a direct comparison. To do this, we extract the lower-frequency component of the new r^* measure. Following Del Negro, Giannone, Giannoni, and Tambalotti (2017), we use the medium term forecast (specifically the ten-year projection) of the rate as our long-run r^* :

$$r_t^{*LR} = E_t(r_{t+40}^*). \quad (\text{A.8})$$

Figure A.9 shows the path of long-run r^* under the alternative specification along with the median path from the baseline specification. The baseline specification remains in a relatively tight area around the alternative specification for much of the sample, then plummets during the financial crisis. While the median path of the baseline model drops about three percentage points to around -1, the dip in the alternative specification, driven more significantly by the growth rate, is significantly less. Thus, a major factor in determining the level of long-run r^* in 2017 would appear to be the assumption that all the shocks during the financial crisis are permanent.

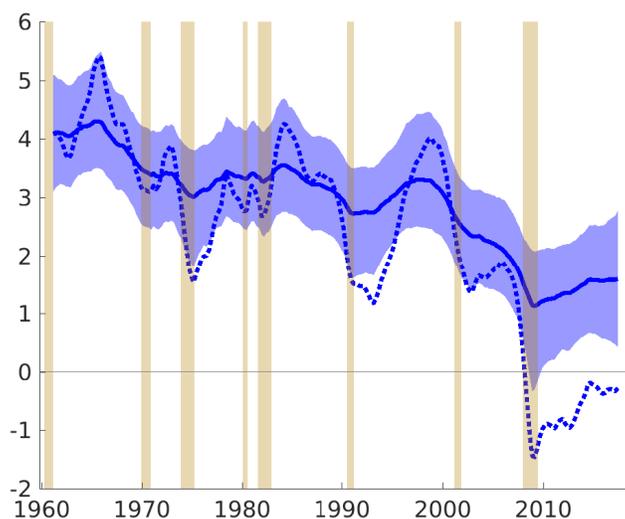


Figure A.9: r^* Path

NOTES: A comparison of the path of long-run r^* under the baseline and alternative models. The solid blue line shows the median path of the smoothed (two-sided) estimate of the alternative specification and the blue-shaded area is bounded by the 5th and 95th percentiles of this estimated path. The dotted blue line shows the median estimated path of the long-run r^* under the baseline specification. The vertical shaded bars represent NBER-dated recessions.