

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

Kurt F. Lewis Francisco Vazquez-Grande

January 3, 2019

A Appendix

A.1 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.2 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.¹ Following some algebraic manipulation, these equations are given as follows. First, the observation equations on real GDP and inflation.

¹Here, we retain a fully flexible structure for AR(1) properties in g and z , but $\mu_z = 0$ throughout the exercises in this paper.

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

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

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

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

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

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

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

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

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

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

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.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.3.1 Cross- and Serial-Correlation of Error Processes

A component of assessing model fit is to examine some of the properties of the error processes which result from the model estimates. Here, we compile the time series of errors for each draw of θ , the model parameters. Each draw of the parameters generates time series of the five errors and this allows us to view the correlation properties of the errors for that draw of the parameters. Each sample of errors is compared for cross- and serial-correlation, and distributions (based on all the draws from posterior distribution of θ) of those correlations are shown in Figures A.1 and A.2. The correlations across errors (shown in the “off-diagonal” distributions the array of distributions) and the serial correlations of each error (the distributions shown on the diagonal of the array, shown in blue) appear to be generally distributed in an area around zero.

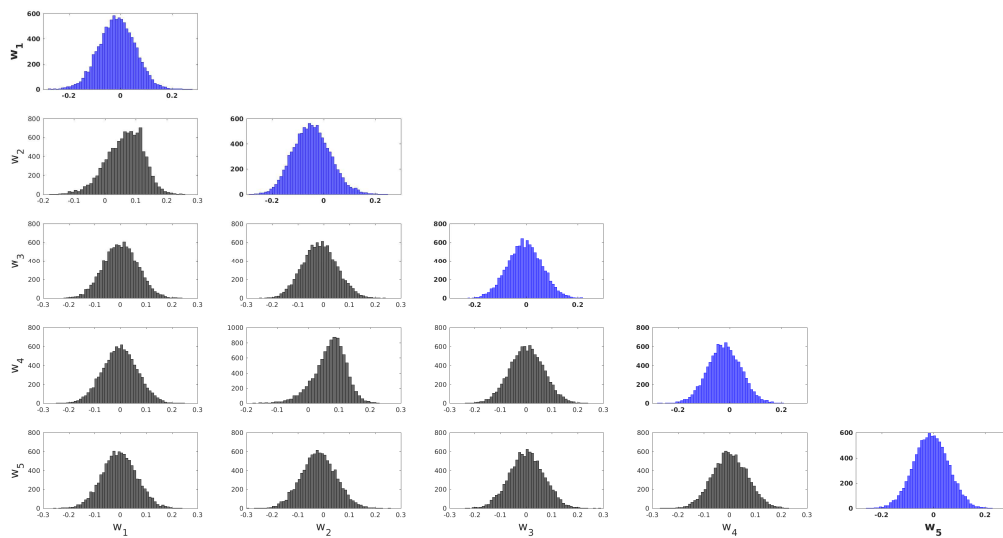


Figure A.1: Baseline: Distributions of Cross- and Serial-Correlation of the Model Errors

NOTES: The diagonal of this lower-triangular array of distributions shows the distributions of the serial correlation of each of the five error. These are shown in blue and have bold font in the axis. The “off-diagonal” distributions are the correlations across the errors, shown in black.

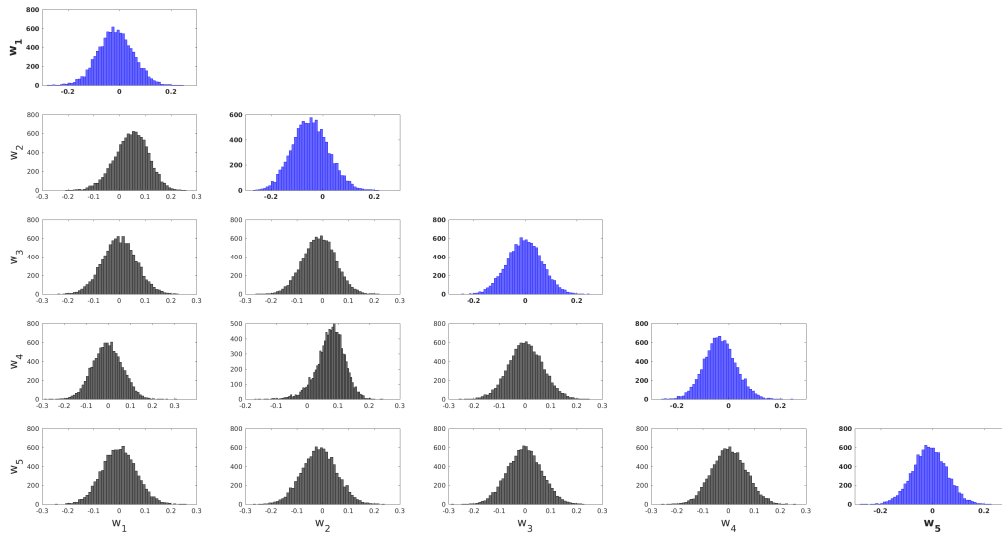


Figure A.2: Alternative: Distributions of Cross- and Serial-Correlation of the Model Errors

NOTES: The diagonal of this lower-triangular array of distributions shows the distributions of the serial correlation of each of the five error. These are shown in blue and have bold font in the axis. The “off-diagonal” distributions are the correlations across the errors, shown in black.

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.3. 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.4 and A.5. 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.6. 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.3: 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.4: μ_g from a flexible AR(1) specification for g




Figure A.5: 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.6: 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.7 and A.8 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.9 and A.10 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).

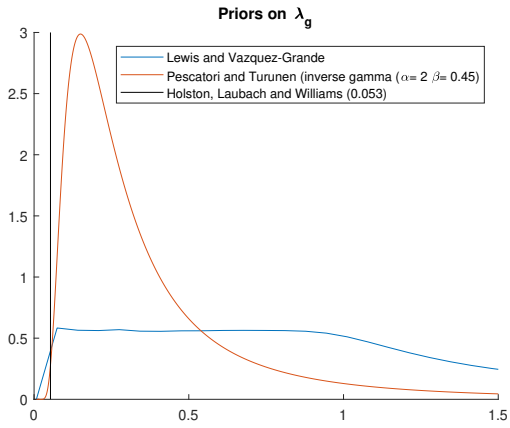


Figure A.7: Priors of λ_g

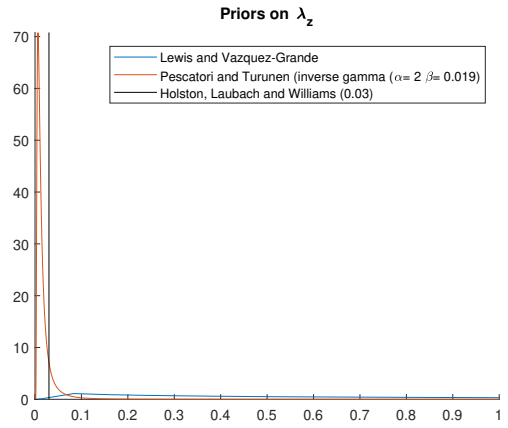


Figure A.8: Priors of λ_z

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




Figure A.9: Posterior Distribution of σ_3




Figure A.10: 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.11 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.11: 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.

A.7 An Alternative Interest Rate and Sample Period

This section of the appendix provides details on the examination of the alternative subsample and interest rate measure discussed in the result section of the main paper. We examine the effects of (1) using an alternative measure of the real rate constructed using the shadow rate of Wu and Xia (2016) and (2) of beginning the sample in 1983 at the onset of the period referred to as the Great Moderation. We begin by replicating the main results of the paper, shown in alternative versions of Table 2 from the main text for both of these changes individually. The remainder of the appendix goes into greater detail on the key implications.

	Bayesian		MLE	
	Baseline	Alternative	Baseline	Alternative
a_1	1.242 [0.93,1.51]	1.137 [0.79,1.50]	1.520 [1.34, 1.70]	1.518 [1.34, 1.70]
a_2	-0.327 [-0.59,-0.04]	-0.235 [-0.58,0.10]	-0.579 [-0.96, -0.48]	-0.576 [-0.97, -0.48]
a_r	-0.113 [-0.19,-0.05]	-0.108 [-0.19,-0.05]	-0.066 [-0.09, -0.04]	-0.063 [-0.09, -0.03]
b_1	0.681 [0.57,0.79]	0.673 [0.56,0.78]	0.674 [0.60, 0.74]	0.676 [0.61, 0.75]
b_Y	0.058 [0.03,0.12]	0.070 [0.03,0.14]	0.073 [0.03, 0.11]	0.072 [0.03, 0.12]
σ_1	0.404 [0.12,0.64]	0.326 [0.08,0.61]	0.366 [0.21, 0.52]	0.376 [0.21, 0.54]
σ_2	0.803 [0.74,0.87]	0.797 [0.74,0.86]	0.792 [0.75, 0.84]	0.793 [0.75, 0.84]
σ_3	0.650 [0.09,1.91]	2.165 [0.79,4.11]	0.180 [0.10, 0.26]	0.195 [0.10, 0.29]
σ_4	0.515 [0.15,0.64]	0.539 [0.10,0.65]	0.567 [0.47, 0.67]	0.562 [0.46, 0.67]
σ_5	0.048 [0.02,0.12]	0.047 [0.02,0.10]	0.030 [0.02, 0.03]	0.029 [0.02, 0.03]
ρ_z	1*	0.736 [0.36,0.91]	1*	0.942 [0.79, 1.09]

(a) Estimation of the Parameters – Wu-Xia Shadow Rate

	Bayesian			MLE	
	LL(θ_{med})	Log Marg. Like.	BF	Log. Like.	BIC
Baseline	-520	-530	0.15	-519	1091
Alternative	-517	-526	6.8	-518	1096

(b) Model Comparison Under Bayesian and MLE Methods – Xu-Wia Shadow Rate

Table A.2: Panel (a) shows the medians of the marginal posterior distributions for each of the model parameters from the baseline and alternative specifications, along with the MLE, all estimated using the Wu-Xia shadow rate—in place of the effective federal funds rate—to construct the real rate. 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 log-likelihood of the model evaluated at θ_{med} , the medians of the marginal posterior distributions from Panel (a) along with 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.

	Bayesian		MLE	
	Baseline	Alternative	Baseline	Alternative
a_1	1.188 [0.72,1.50]	1.181 [0.80,1.46]	1.691 [1.44, 1.95]	1.700 [1.46, 1.94]
a_2	-0.383 [-0.66,-0.06]	-0.356 [-0.62,0.01]	-0.722 [-0.96, -0.48]	-0.723 [-0.97, -0.48]
a_r	-0.018 [-0.065,-0.01]	-0.024 [-0.07,-0.01]	-0.009 [-0.05, 0.03]	-0.008 [-0.05, 0.03]
b_1	0.261 [0.12,0.40]	0.251 [0.56,0.78]	0.387 [0.25, 0.52]	0.388 [0.25, 0.53]
b_Y	0.153 [0.05,0.35]	0.159 [0.05,0.33]	0.025 [-0.01, 0.06]	0.025 [-0.04, 0.09]
σ_1	0.155 [0.03,0.36]	0.158 [0.03,0.36]	0.233 [0.09, 0.38]	0.240 [0.07, 0.41]
σ_2	0.628 [0.57,0.70]	0.629 [0.57,0.70]	0.577 [0.52, 0.64]	0.576 [0.51, 0.64]
σ_3	1.000 [0.12,4.01]	2.320 [0.38,4.57]	2.240 [-7.21, 11.69]	2.877 [-11.63, 17.39]
σ_4	0.40 [0.21,0.48]	0.397 [0.23,0.48]	0.412 [0.33, 0.50]	0.411 [0.32, 0.50]
σ_5	0.159 [0.08,0.25]	0.151 [0.07,0.24]	0.019 [0.01, 0.02]	0.019 [0.01, 0.02]
ρ_z	1*	0.647 [0.19,0.95]	1*	0.922 [0.62, 1.23]

(a) Estimation of the Parameters – Post-1983 Sample

	Bayesian			MLE	
	LL(θ_{med})	Log Marg. Like.	BF	Log. Like.	BIC
Baseline	-240	-249	0.6	-247	543
Alternative	-239	-247	1.6	-246	547

(b) Model Comparison Under Bayesian and MLE Methods – Post-1983 Sample

Table A.3: Panel (a) shows the medians of the marginal posterior distributions for each of the model parameters from the baseline and alternative specifications, along with the MLE, all estimated with a data sample that begins in 1983. 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 log-likelihood of the model evaluated at θ_{med} , the medians of the marginal posterior distributions from Panel (a) along with 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.

A.7.1 Replacing the Federal Funds Rate with the Shadow Rate

As the federal funds rate reached a lower bound in the aftermath of the Great Recession, we look at the effect of imputing additional monetary policy stimulus into the real rate data by using the shadow rate of Wu and Xia (2016). This series is identical to the effective federal funds rate except during the period of the zero lower bound; the two series are shown in Figure A.12.

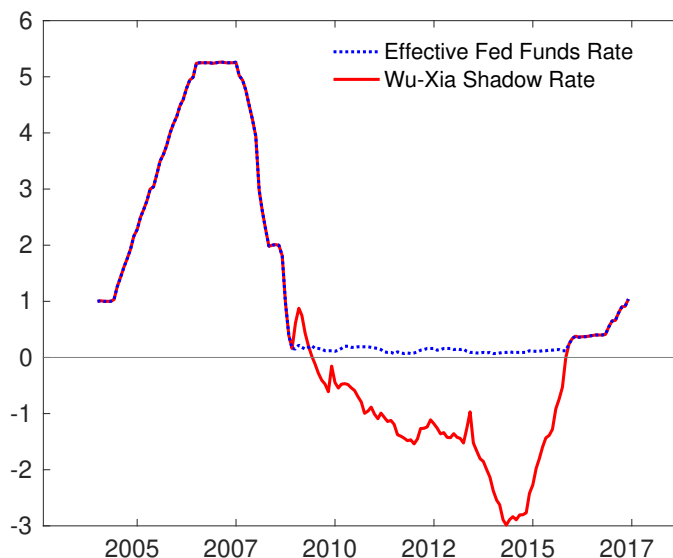


Figure A.12: The Effective Federal Funds Rate and the Shadow Rate from Wu and Xia (2016)

We can see from Table A.2 that fairly small changes to the parameter estimates occur as the result of this change. For example, and as noted in Pescatori and Turunen (2016), we get a small (insignificant) reduction in the a_R parameter estimate as a result of incorporating this different measure of monetary policy accommodation. As would be expected given the purpose of the shadow rate, the model results display considerably more accommodation than under the estimation that uses the effective federal funds rate to build the real rate data. While r^* estimates do decline slightly during the ZLB period when r^* is estimated using the shadow rate (more on this shortly), the resulting rate gap ($r_t - r_t^*$) during the ZLB episode is still considerably more negative, as shown in Figure A.13.

This more negative rate gap indicates to the model that, relative to the results when estimating the model with the effective federal funds rate, the output gap should open wider as a result of this increased accommodation. Accordingly, estimates of the trend growth rate during the ZLB episode are pushed down slightly as the IS equation takes the rate gap signal to assign more of the output growth during that period to the transitory gap component rather than to potential output, as seen in Figure A.14.

The path of r^* is estimated to be slightly lower, partly as a result of this lower relative estimate of g , but the majority of the relatively small decline in r^* comes from the non-growth component, z . Finally, the small changes to the model's estimates of r^* when using the shadow rate of Wu and Xia (2016) do not change any of the central results in the main paper regarding the likelihood of transitory shocks to z . The Bayes Factor in favor

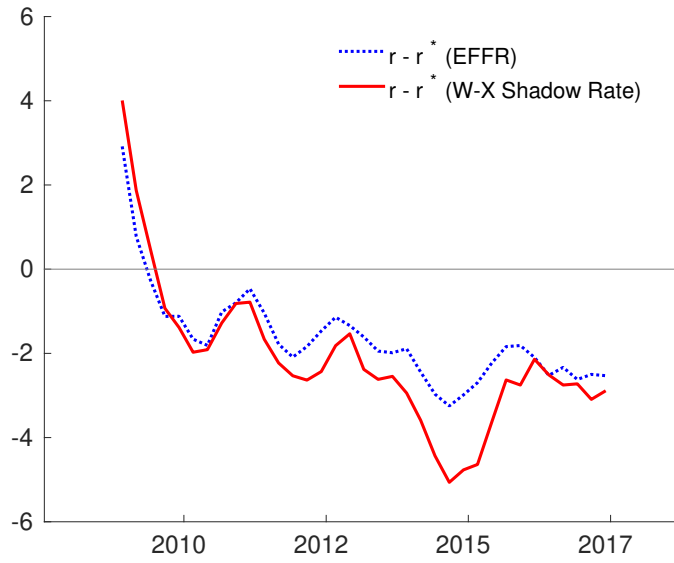


Figure A.13: The rate gap ($r_t - r_t^*$)

of transitory shocks is roughly 7, in the same statistical category (in the Kass and Raftery, 1995 categorization) as the result when using the effective federal funds rate to build the real rate data.

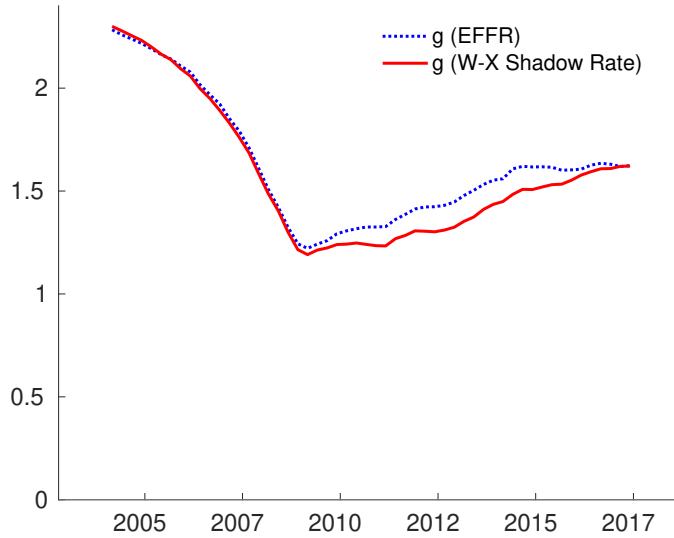


Figure A.14: The growth rate of potential output, g

A.7.2 Sub-Sample Beginning in 1983

The estimation using the sub-sample period shown in Table A.3 has results which are somewhat different from the full-sample results shown in Table 2 of the main text. We begin by looking at the estimate of the non-growth component of r^* , the principle object of interest in our study. Figure A.15 shows that while the directional pattern of fluctuations of z remains similar to that of the full sample estimated path, the magnitude of these movements is considerably lower.




Figure A.15: z Path with 1983-2017 Sub-Sample

NOTES: The blue-shaded area is bounded by the 5th and 95th percentiles of the full-sample estimated path.

Figure A.16 shows the marginal posterior distribution of ρ_z under the full sample and the sub-sample, along with the prior to demonstrate the decline in the Bayes Factor shown

in Table A.3 that occurs when we remove the first 40% of the data. The Bayes Factor still represents evidence in favor of the alternative specification with transitory shocks, but it has declined to 1.6.

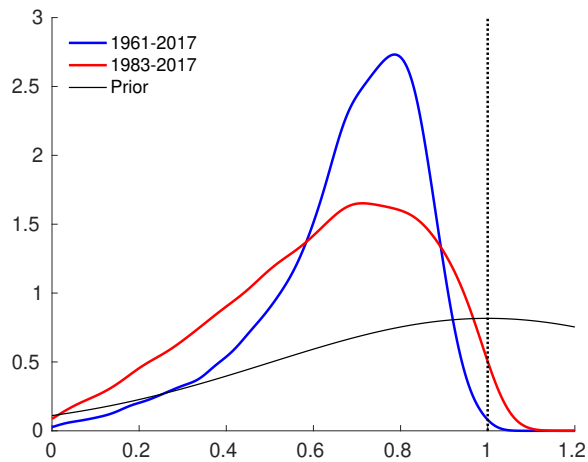


Figure A.16: Prior and Posterior of ρ_z

This reduction in the Bayes Factor is due to the flattening of the posterior distribution of ρ_z . While there is a decline in the Bayes Factor under the sub-sample, the alternative model is still preferred, the mode of the sub-sample marginal posterior distribution is actually below the full sample estimate, and the median declines due to the increase in posterior weight on values of ρ_z closer to zero more than those near one. The removal of 40% of a data sample is likely to reduce the precision of posterior estimates of a parameter, but the part of the sample which is excluded in this exercise magnifies the effect. The decline in the Bayes Factor is not so much the result of particularly compelling evidence during the Great Moderation that ρ_z is precisely equal to one so much as it is the result of the Great Moderation period emphasizing the challenges of estimating the properties of the non-growth component of r^* via its effects on the output gap. That is, the higher uncertainty of the posterior distribution of ρ_z is related to the “flatter” IS curve estimated during the subsample.

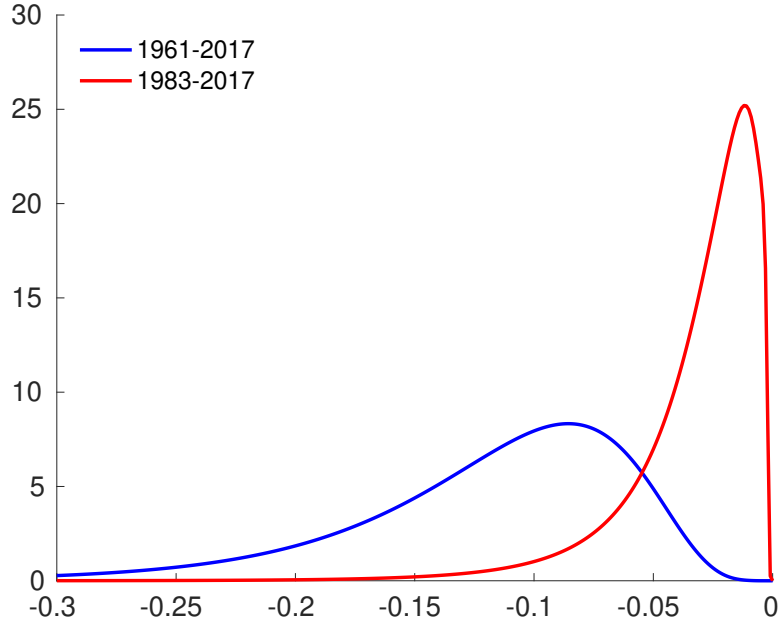


Figure A.17: Posterior of a_r estimated during the Great Moderation

Figure A.17 shows that the relationship between the rate gap and the output gap implied by a_r is estimated to be much weaker when using the post-1983 subsample in comparison to the full sample. Recall that in estimations of this equation in the literature (e.g. this study, Holston et al., 2017; Pescatori and Turunen, 2016) a prior restriction is imposed that the coefficient on the real rate gap in the IS equation, a_r , is negative. As the IS equation becomes “flatter,” Figure A.16 indicates that the autoregressive properties of z become less precisely estimated. The results in Table A.3 indicate that this problem extends to the MLE results, but that the broad relationships across the model comparison results are similar: Bayesian methods choose the alternative model while MLE chooses the baseline. Others have also dealt with this issue. Pescatori and Turunen (2016) make use of this shorter sample period and produce Bayesian estimates, dealing with the potential identification problem by use of priors. They use tight priors on the parameters which define the time-series properties of z to be similar to those of the original Laubach and Williams (2003) estimates. (Specifically, they set the mean of the prior distribution for their counterpart to ρ_z equal to 0.99 with a standard deviation of 0.0025. Similarly, the values in their table of prior parameters, combined with their prior on λ_z shown in Figure A.8 demonstrates the very tight priors for σ_z used for identification of z under the sub-sample.) These priors are able to overcome the identification problem when using data with less volatility in inflation, interest rates and output. In our study, however, because we are investigating the more-focused question of whether r^* is hit by transitory or only permanent shocks, strong prior assumptions about the properties of z would pre-determine the results of our investigation.

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.
- 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.
- 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.
- 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.
- Kass, R. E. and A. E. Raftery (1995). Bayes factors. *Journal of the American Statistical Association* 90(430), 773–795.
- 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.
- Wu, J. C. and F. D. Xia (2016, March). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit and Banking* 48(2-3), 253–291.