Measuring the Natural Rate of Interest: Alternative Specifications

Kurt F. Lewis and Francisco Vazquez-Grande

Measuring the Natural Rate of Interest:

Alternative Specifications[*]

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

Current version: May 22, 2017

First version: February 15, 2017

Abstract

We build on the work of Laubach and Williams (2003) and subsequent studies by analyzing the effect on the estimates of the natural rate of interest (r^*) of accounting for full parameter uncertainty and alternative specifications for the underlying components of the natural rate. Our estimation technique delivers richer time-series dynamics for the median estimate of r^* within the Laubach and Williams model. Additionally, we find that models with transitory shocks to the non-growth component of the natural rate have a higher marginal likelihood and produce an upward-sloping post-crisis trajectory of the r^* path and thus a higher recent median point estimate (1.8% in 2016:Q3).

[†]Federal Reserve Board. kurt.f.lewis@frb.gov

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

1   Introduction

For central banks that use a short-term interest rate as their main policy tool, the "natural" or "equilibrium" rate of interest is of crucial importance for the conduct of monetary policy. This is because the natural rate provides a measure of whether the stance of policy is contractionary or expansionary, which is still informative even after reaching the effective lower bound of short-term interest rates. We follow Laubach and Williams (2003) and use the definition for the natural rate (r^*) from Bomfim (1997): the real short-term interest rate consistent with output converging to potential, where potential is the level of output consistent with stable inflation and is therefore the medium-term real rate for the conduct of monetary policy and the intercept in Taylor-type rules.[1]

r^* is unobservable and it is well-documented that the range of r^* estimates for the U.S. economy across models is fairly wide and that within individual models the uncertainty about the level of r^* is high. However, in recent years there has been a public discussion among policymakers centering around the question of whether r^* has declined to new depths, and whether it may ever rebound (see, for example Yellen, 2015 and Williams, 2016, 2017). In several cases financial "headwinds" and the role of the non-growth components of r^* have been cited as important contributors to its low level since the great recession (e.g. Williams, 2015). Our analysis demonstrates that within a well-known central bank model of r^*, changes in estimation strategy and model specification can change the median path of r^* estimates in economically significant ways.

In this paper we build on the work of Laubach and Williams (2003, 2016) and provide alternative estimates of the natural rate by varying the estimation technique to account for full parameter uncertainty, as well as by making small changes to the model specification driving the dynamics of the unobserved components of the natural rate. Our estimation technique employs Bayesian methods to incorporate the uncertainty of all the model param-

---

[1]Recent literature (Kiley and Roberts, 2017, for example) show that these simple rules can work badly in the vicinity of the effective lower bound.

eters jointly into the estimate of r^* under loose priors. In contrast to much of the extant literature we do this in a single step, rather than the three-step process of Stock and Watson (1998) used by Laubach and Williams (2003) and subsequent work. Within this framework we also explore alternative specifications for the latent processes that drive the natural rate. As in Laubach and Williams (2003, 2016) we consider r^* to be driven by two hidden process, one that controls the rate of growth of potential output and another that represents the persistent non-growth component of interest rates. In this paper we revisit whether the two components of the natural rate follow random walks and find that, in contrast to the first investigations, there is some evidence that the data do not indicate a unit root in non-growth component of r^*. Following our analysis, we briefly discuss how the non-growth component of the estimated natural rate may account for changes in the intratemporal ratio of marginal utilities, in excess of expected consumption growth, and how these changes can be attributed to financial shocks.

By incorporating the uncertainty of estimating all the parameters jointly in a single step under uninformative priors—but using the Holston, Laubach, and Williams (2016) model specification—we obtain richer time-series dynamics of r^*. Our estimation shows deeper drops during recessions, with subsequently larger increases during recoveries, than the estimates obtained from maximum likelihood methods. Our median path of the natural rate also shows a different trajectory since the end of the Great Recession, and we obtain an increase since the trough in 2008, in contrast to the "non-recovery" displayed of the U.S. estimate in Holston, Laubach, and Williams (2016).

We explore alternative specifications without permanent shocks on the non-growth component of r^* and find an elevated level of the median estimate after the great recession, about 1.75% higher than that of Holston, Laubach, and Williams (2016) in the third quarter of 2016. The dynamics of the non-growth component are hard to estimate, a finding which mirrors the original results of Laubach and Williams (2003) as well as Kiley (2015). When this process is not assumed to be a random walk, we estimate a greater recovery of the

2

natural rate since the lows of the Great Recession, reaching 1.8% at the end of the third quarter of 2016. We therefore infer that permanent shocks to the non-growth component of the natural rate are needed to produce a persistent low level of r^* after the Great Recession.

This paper follows on a strain of the literature of empirical macroeconomics started by the work of Laubach and Williams (2003), which has been extended in different ways. Lubik and Matthes (2015) use a time-varying parameter vector-autoregression model, which allows for more nonlinear relationships than the structure of the original Laubach and Williams (2003) study, and find a similar secular decline in the level of r^*. Some of the more recent literature has also pursued a Bayesian approach to the investigation of r^*. Kiley (2015) removes the dependence of r^* on potential output growth and finds that the remaining non-growth-based natural rate is very difficult to identify, a finding with which our analysis of the non-growth r^* component concurs. Johannsen and Mertens (2016) include stochastic volatility and an effective lower bound of interest rates in the measurement and state equation, but maintain the random walk assumption of potential output growth and find a level of r^* slightly lower than that reported in Holston, Laubach, and Williams (2016). Pescatori and Turunen (2016) follows an approach similar to our own in starting with the standard Laubach-Williams definition for r^*, but makes use of tight priors on the parameters of some of the unobserved components and addresses other issues by bringing in additional information to the model such as outside estimates of the output gap, "shadow" interest rates to help account for the zero lower-bound, and data on savings and economic uncertainty to aid the non-growth component. Also related, Del Negro et al. (2017) use macroeconomic series as well as financial and survey data to estimate a VAR with shifting endpoints and a DSGE model with highly persistent financial shocks to conclude that the low level of r^* is mostly due to an increase in the premium for safety and liquidity, what in our model would be attributed to an persistent decline in the non-growth component of the natural rate.

An important element of the recent empirical work is evident in Kiley (2015) and the primary focus of Hamilton, Harris, Hatzius, and West (2016), who call into question the rel-

evance of the growth of potential GDP in the dynamics of r^*.[2] Almost every macroeconomic models postulate a relation between expected growth of potential output and r^*, in their analysis Hamilton et al. suggests that the connection between ex post real interest rates and realized real GDP growth is statistically limited. We take this insight to modify the $r*$ equation of Laubach and Williams and relate the level of the natural rate to the *expected* level of future output growth and the expected level of a non-growth component.[3]

Our paper shows that Bayesian analysis of the original model Laubach and Williams under uninformative priors calls into question the finding that the level of r^* has been in a secular decline since the early 2000s. Further, a simple (and data-supported) change to the dynamics of the non-growth component of r^* can give rise to a median path with an upward trend over the past 6 years that has already returned r^* to roughly its pre-crisis level.

This paper is organized as follows, section 2 describes the model and specifications to be estimated. Section 3 describes our estimation procedure. Section 4 comments on the results and 5 concludes.

2   The r^* Model

2.1   Model structure

The fundamental model closely follows that of Laubach and Williams (2003, 2016) (henceforth referred to as LW) and features an output gap equation, an inflation equation and the evolutions of unobserved variables such as the level and growth rate of potential GDP and

---

[2]Clark and Kozicki (2005) also note the challenge of using any estimate of potential output growth in "real-time", but in the present paper we set aside real-time concerns to look at the broader trends implied by the path of r^* over the past 50 years.

[3]Note that this does not limit our ability to compare our results to those of Laubach and Williams, as under their dynamics for the hidden processes, both models are isomorphic.

4

the natural rate of interest. The six equations at the heart of our exercise are:

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

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

$$r_t^* = E_t \left(g_{t+1} + z_{t+1} \right) \tag{2.3}$$

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

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

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

where y_t is 100 times the natural log of real GDP in period t, y_t^* is the potential GDP analog of this value and $\tilde{y}_t = y_t - y_t^*$. Similarly, $\tilde{r}_t = r_t - r_t^*$ where r_t is the real short-term interest rate in period t, defined as the nominal short-term rate less inflation expectations for that period, and r^* is the "natural rate of interest" which is the focus of this study.

Note that our specification of r^* coincides with that of LW when using their dynamics for g_t and z_t, that is, when ρ_g and ρ_z are assumed to be equal to 1.

Substituting the formula for r_{t-1}^* in equation (2.3) into the output gap specified by equation (2.1), and rearranging to construct the traditional observation/transition equation structure, we can write the system of equations (2.1) to (2.6) in state space form.

$$s_t = A s_{t-1} + B u_t + C w_t \tag{2.7}$$

$$x_t = D s_t + F u_t + G w_t \tag{2.8}$$

where the state vector (s_t) contains the level of potential GDP (as well as two lags), in addition to two lags each of g_t and z_t. The observed variables (x_t) are real GDP and inflation, the lags of which are treated as exogenous variables (u_t) along with the lags of the real rate. We specify each of the shocks in w_t to be an *i.i.d.* $N(0, 1)$ and estimate the variance of the shock processes in the system using the σ_i parameters. Further specifics regarding the state space formulation are available in Appendix A.

Table 1: Model Reference

| Parameter | Model | | | |
	I	II	III	IV
ρ_g	1	Estimated	1	Estimated
μ_g	Unidentified	Estimated	Unidentified	Estimated
ρ_z	1	1	Estimated	Estimated

NOTES: The four model specifications examined in the paper vary the dynamics of the processes for the g_t and z_t processes by restricting their autoregressive coefficients to be 1 or to be estimated. When g_t is allowed to be stationary, the unconditional mean is also estimated but when g_t is a random walk, the unconditional mean is unidentified.

2.2 Alternative specifications

We examine different combinations of dynamics for the g_t and z_t process and the effects of these alternative specifications on the median path of the r^* estimate.

The first specification tested will be the one that most closely replicates the results of LW. That is, Model I sets $\rho_g = \rho_z = 1$ as in LW but uses the standard Bayesian methods discussed below to estimate the model.In this specification, shocks to both the growth rate of potential GDP and to the z process are considered to be permanent.

The second specification estimates the persistence and mean of the process for potential GDP growth. That is, in Model II, both ρ_g and μ_g are estimated while $\rho_z = 1$. The third specification returns the permanence of shocks to the path of potential GDP growth but estimates the persistence of the shocks in the z_t process. Finally, Model IV estimates the parameters of both processes jointly. See Table 1 for a concise reference on the model differences.

Note that—unlike the process for potential GDP growth—we do not estimate an unconditional mean for the z_t process, or rather we assume that mean is zero. We interpret z_t to be playing more of a "head-winds" roll in this reduced form model and as such expect its unconditional mean to be zero. A more systematic investigation of the possible economic drivers of the z_t process is the subject of ongoing work by the authors.

3   Estimation Procedure

3.1   Data

To facilitate result comparison, the data used in this analysis is the same as the data used in Holston, Laubach, and Williams (2016), and it is transformed in the same way.[4] 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 Holston, Laubach, and Williams (2016) 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.

3.2   Bayesian Estimation

Because we implement a fully Bayesian methodology we do not employ the three-step process developed in Stock and Watson (1998) and used in Laubach and Williams (2003, 2016). That process was put in place as a way to deal with the so-called "pile-up" problem when estimating models of this type with maximum likelihood. As demonstrated by DeJong and Whiteman (1993) and Kim and Kim (2013), Bayesian methods do not suffer from the "pile-up" problem and thus we proceed with the one-step estimation discussed above. In order to confirm that our structure nests that of the literature that has followed Laubach and Williams (2003), we have estimated a version of the model that imposes the restrictions implied in the three-step process and recover the median path of r^* for the U.S. reported in Holston, Laubach, and Williams (2016).[5]

---

[4]See the data appendix in Holston, Laubach, and Williams (2016) for additional specifics on obtaining the data.

[5]Specifically, we impose the values for $\lambda_g = \sigma_g/\sigma_{y^*}$ and $\lambda_z = \frac{a_r\sigma_z}{\sigma_{\tilde{y}}}$ which are estimated by the HLW model in their first and second steps and which are then imposed during the final step.

Each version of the state space model outlined in section 2 is estimated using standard Bayesian methods.[6] 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 filtered path for the unobserved variables, including r^*.

3.3   Prior Distributions

From equations (2.1) to (2.6), we construct the vector of model parameters:

$$\theta = \begin{bmatrix} a_1 & a_2 & a_r & b_1 & b_Y & \rho_g & \mu_g & \rho_z & \sigma_1 & \sigma_2 & \sigma_3 & \sigma_4 & \sigma_5 \end{bmatrix}.$$

The prior distributions for the model parameters were chosen with a mind toward minimal informativeness and to enforce a few restrictions, most of which were inherited from the Laubach and Williams (2003) structure. For consistency with their paper, we enforce that a_r be negative and b_Y positive just as was done in Holston, Laubach, and Williams (2016).[7] Similarly, 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 priors on the standard deviations of the shock processes (σ_i's) are assumed to be uniform, inclusive of zero, with a maximum we set to 5. The summary of the prior distributions used is given in Table 2. Regarding implementation of the priors for ρ_g and ρ_z, when those parameters are set to one in the models the priors are changed to be degenerate at one.

---

[6]Popularized by Tierney (1994) and Chib and Greenberg (1995, 1996), examples of textbook treatments of this approach can be found in Geweke (2005) and Herbst and Schorfheide (2015).

[7]We follow their practice of using $a_r < -0.0025$ and $b_Y > 0.025$ as the actual restrictions.

Table 2: Prior Distributions

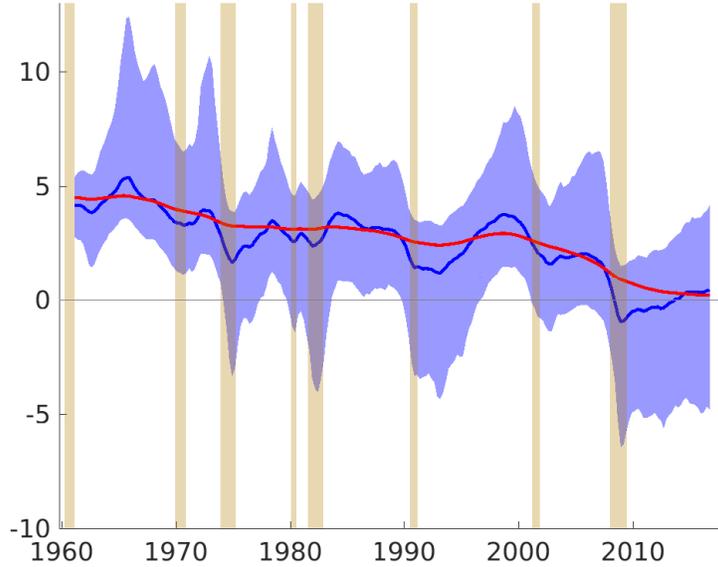| Name	Domain	Density	Parameter 1	Parameter 2
a_1	$\mathbb{R}$	Normal	0	2
a_2	$\mathbb{R}$	Normal	0	2
a_r	$\mathbb{R}^-$	Normal	0	2
b_1	$[0,1]$	Uniform	0	1
b_Y	$\mathbb{R}^+$	Normal	0	2
ρ_g	$\mathbb{R}^+$	Normal	0	2
μ_g	$\mathbb{R}$	Normal	0	2
ρ_z	$\mathbb{R}^+$	Normal	0	2
σ_1	$[0,5]$	Uniform	0	5
σ_2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

NOTES: The table presents the marginal prior distributions for the individual model parameters. 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.

4   Results

The median path of r^* under all of the model specifications studied in this paper displays richer time-series dynamics than that which is estimated using the Laubach and Williams (2003)-style techniques. Said another way, it appears that the MLE strategy results in fairly small estimates of σ_3 and σ_5 which feeds into a smaller estimate of the time series standard deviation of r^*. This is because, in the multi-step strategy some unobservable parameters are fixed to their MLE estimate in the last step of the estimation.

Figure 1 shows the path of r^* estimated under the assumption that both g_t and z_t are random walks (Model I). The blue line in Figure 1 (and all subsequent plots of r^*) represents the median estimated path of r^* while the red line provides the median path from Holston, Laubach, and Williams (2016), which will be included for reference whenever
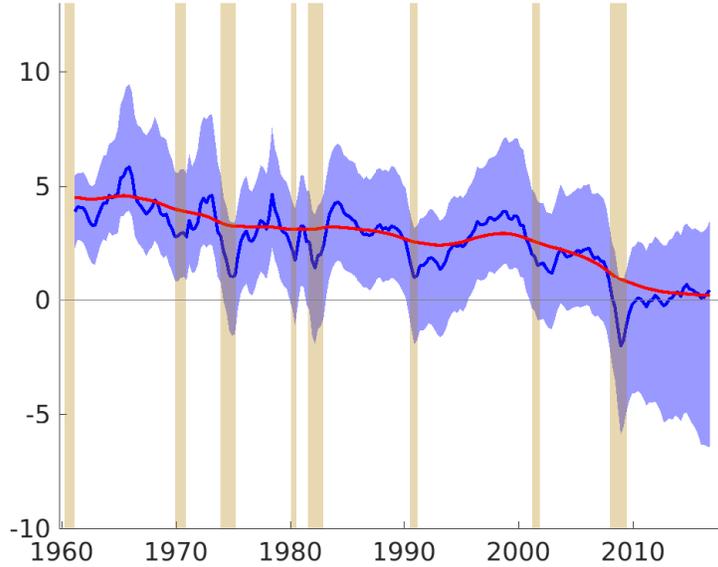
Figure 1: r^* Path (Model I)

NOTES: The path of r^* under Model I when $\rho_g = \rho_z = 1$. The solid blue line shows the median path of the smoothed (two-sided) estimate and the blue-shaded area is bounded by the 10^{th} and 90^{th} percentiles of the estimated path. The red line is the two-sided r^* path for the U.S. from Holston, Laubach, and Williams (2016). The vertical shaded bars represent NBER-dated recessions.

alternative models of r^* are shown. The change in estimation technique generates a non-trivial difference in the median path insofar as the natural rate of interest seems to respond much more quickly to adverse shocks and to have some bounce-back following recessions. While there is a still somewhat of a secular decline in the median r^* path estimated in Model I, the recent trajectory of that path is markedly different. Both the estimates fall during the Great Recession, but Model I's estimate plummets and begins a recovery immediately following the recession. That recovery is slow and ultimately ends up in roughly the same place as the last estimated value in the LW series, but the last 6 years of the two paths tell a different story. The path estimated in LW drops during the great recession, then the rate of decline slows dramatically, but no recovery in r^* is apparent.

Before moving on to the path of r^* under additional models, it is important to note that
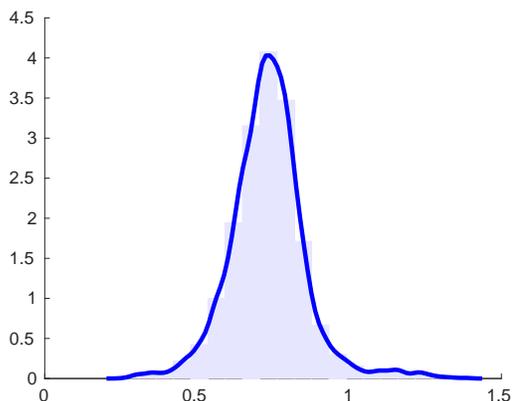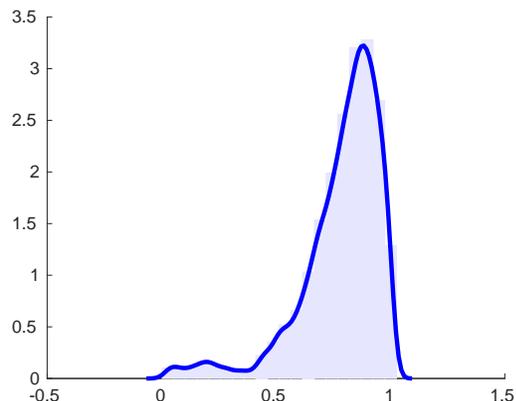
10

Figure 2: r^* Path (Model II)

NOTES: The path of r^* under Model II when $\rho_z = 1$ but ρ_g and μ_g are estimated. The solid blue line shows the median path of the smoothed (two-sided) estimate and the blue-shaded area is bounded by the 10th and 90th percentiles of the estimated path. The red line is the two-sided r^* path for the U.S. from Holston, Laubach, and Williams (2016). The vertical shaded bars represent NBER-dated recessions.

all of the paths for r^* estimated in this paper are not statistically significantly different from the LW estimated path. As seen in Figure 1, the level of uncertainty about the path of r^* is fairly high and we cannot reject the LW path based on our analysis. Similarly, our paths fall within the confidence bands given by LW. That said, the differences in the dynamics of the two paths may lead to important policy implications. For example, aside from the recent trajectory just discussed, our estimates cannot reject an r^* path which is constant at around 2%.

Allowing for the growth rate of potential GDP to be an estimated process which could be stationary (Model II) changes the median estimated path of r^* only slightly relative to Model I. As can be seen in Figure 2, the volatility of the median path is slightly higher than

Figure 3: Posterior of μ_g Figure 4: Posterior of ρ_g

NOTES: The growth process parameter posterior distributions for Model II. In the panel to the left, μ_g represents the unconditional mean in the AR(1) process and therefore the average quarterly growth rate of potential real GDP. In the panel to the right, ρ_g is the persistence parameter for the AR(1) process.
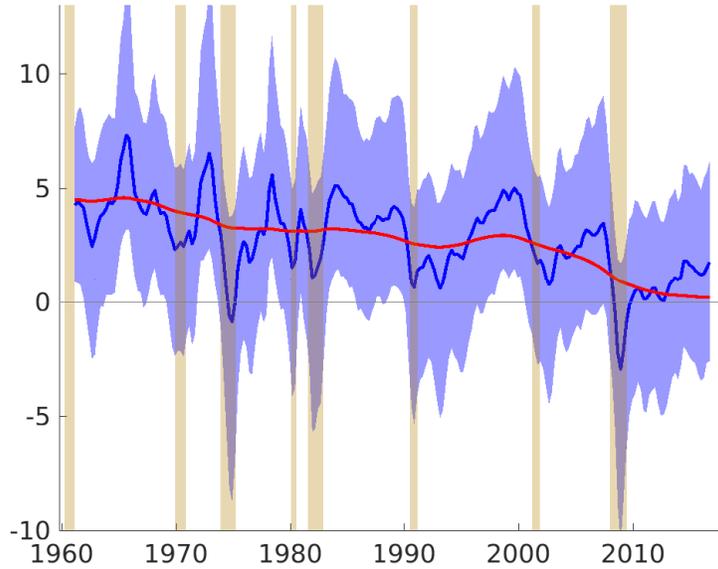
in Model I, and the decline followed by bounce-back around recessions is somewhat stronger. The post-crisis trajectory of r^* under Model II is similarly one of a bounce that is tempered over time and concludes near the end of the Holston, Laubach, and Williams (2016) path for r^*, as it did in Model I.

Estimating the growth process in a manner that allows for the possible stationarity allows us to get estimates of the average quarterly growth rate of potential GDP and to investigate the value of modeling the process as a unit root. Figure 3 shows that the posterior distribution of μ_g peaks around 0.7, giving us a reasonable average quarterly growth rate of potential GDP. Figure 4 shows that while the posterior distribution of ρ_g peaks below 1, there is significant mass on 1, giving some support to the unit root model of potential GDP.[8]

Model III goes back to modeling the growth process for potential real GDP as a random walk. However, it permits the non-growth component of r^* (z_t) to be modeled as a zero-mean AR(1) process. This allows the flexibility for non-permanent shocks that could resemble something more transitory that would appear to be so-called "headwinds." The resulting

---

[8]Whenever ρ_g or ρ_z are estimated, note that they are *not* restricted to be less than 1. The maximum value in the posterior distribution of g_t in Model II is around 1.04.

Figure 5: r^* Path (Model III)

NOTES: The path of r^* under Model III when ρ_z is estimated but $\rho_g = 1$. The solid blue line shows the median path of the smoothed (two-sided) estimate and the blue-shaded area is bounded by the 10^{th} and 90^{th} percentiles of the estimated path. The red line is the two-sided r^* path for the U.S. from Holston, Laubach, and Williams (2016). The vertical shaded bars represent NBER-dated recessions.

path for r^* is shown in Figure 5.

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 in Model I or Model II. Following the sharp dip and bounce-back in the aftermath of the Great Recession, the median path of r^* has generally trended in a positive direction. Under Model II, the recovery from negative values of r^* was also fairly swift, but has been followed by a median path of r^* which was essentially flat over the past five or six years. In Model III however, while the pace of recovery in the level of r^* following the Great Recession slowed by late-2010, it has persisted. The trend over the past five or six years has remained positive, meaning that the small change to the model of the non-growth component had an important effect on the median path of r^*.
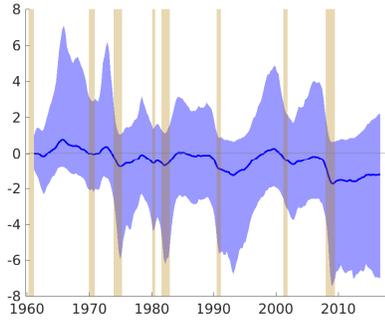
Figure 6: z_t (Model I)
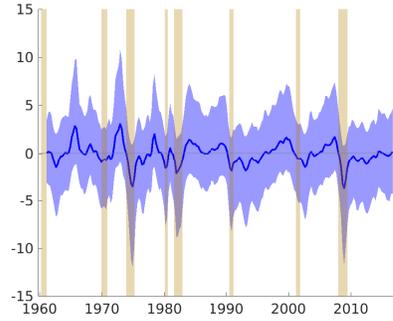

Figure 7: z_t (Model III)
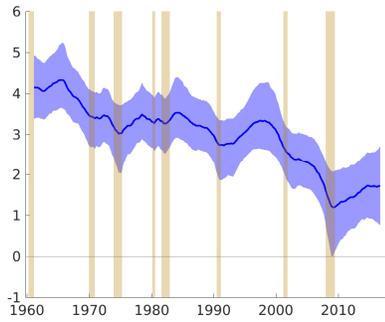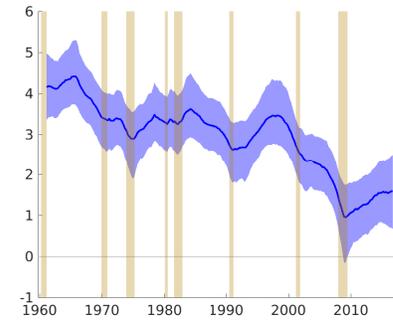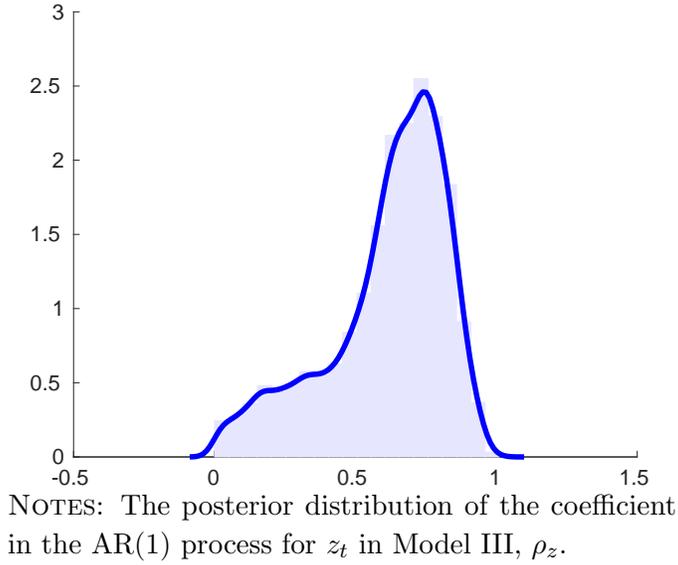

Figure 8: g_t (Model I)

Figure 9: g_t (Model III)

NOTES: The paths of the non-growth component of r^* under Model I and Model III are shown in the top row (Figures 6 and 7). The paths of the quarterly growth rate of potential real GDP under Model I and Model III are shown in the bottom row (Figures 8 and 9. Model I assumes that z_t is a unit root process while Model III estimates an AR(1) process with a zero mean.

Indeed, the path of the z_t process is significantly different under the assumption of permanent shocks. Figures 6 and 7 show the contrast in behavior of the median z_t path in the 21$^{\text{st}}$ century. Though the path recovers a little in the years following the Great Recession, but in Model I the median path of z_t appears to be grinding steadily lower. Under this specification, this more-noticeable secular trend to z_t makes it clear how the differences between the paths of r^* in Models I and III is driven primarily by the behavior of z_t. This is apart from the reality that Model III allows for permanent shocks to growth. Figures 8 and 9 show that the change to stationary z_t has very little effect on the path of the growth rate

14

Figure 10: Posterior of ρ_z

NOTES: The posterior distribution of the coefficient in the AR(1) process for z_t in Model III, ρ_z.

and that growth in Model III continues to be weak, as in Model I. That is, even modeling growth in such a way as to allow for secular stagnation (e.g. Summers, 2014 and Eggertsson, Mehrotra, and Summers, 2016), the median path of r^* shown in Figure 5 still shows a slow and steady recovery when z_t has transitory shocks.

Model III attempts to estimate the autoregressive dynamics of the z_t process rather than assume all shocks are permanent. In doing so, we confirm that the z_t process is fairly difficult to estimate. By looking at equations (2.1) to (2.6) we can see that z_t is going to be more difficult to estimate than the processes associated with potential GDP. The observables allow for a more-or-less traditional decomposition of real GDP into trend and cycle, which gives more information about potential GDP and its growth rate than is obtained about z_t through the rate-gap element of the output gap in equation (2.1). This leads to a fairly imprecise estimate of the AR(1) process for z_t, which can be seen in Figure 10.

The estimate of ρ_z is less precise than the estimate of ρ_g in Model II, which is expected given that there is fairly little information used to filter its value. It has been acknowledged that z_t is a hard object to estimate well by Laubach and Williams (2003) and subsequent papers in the Laubach and Williams mold. Indeed, Kiley (2015) doesn't use the growth rate

15

of potential GDP in calculating r^* (leaving r^*, in LW parlance, as essentially just z_t) and finds the z_t analog to very challenging to estimate, with many of its properties derived mostly from the priors chosen by the modeler. But, we can see from the posterior distribution in Figure 10 that—unlike the coefficient in the AR(1) process for growth in Model II—there appears to be less evidence that shocks to the z_t process are permanent. We return to this shortly.

The final specification, Model IV, allows for joint estimation of AR(1) dynamics in both the g_t and z_t processes. The resulting path of r^*, shown in Figure 11 has similar contours to that of Model III, but is shifted up somewhat as a result of the potential for mean reversion in the growth of potential GDP. When both processes are estimated, the median level of r^* at the end of 2016 is almost 2 percentage points higher than that estimated in Holston, Laubach, and Williams (2016).
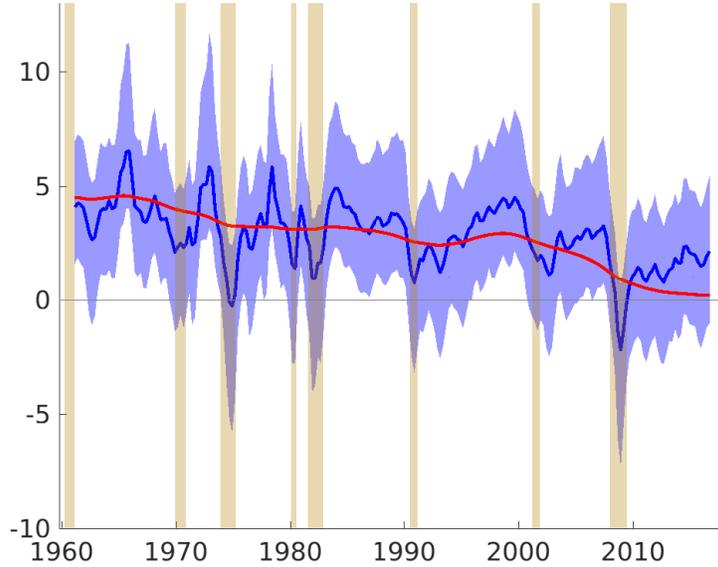
4.1   Model Comparison

It is important to recall that the paths of r^* shown here are not statistically significantly different from HLW. This can be seen by the fact that (by and large) the HLW path of r^* lies inside the shaded region for our path of r^* in each of the four models shown in Figures 1, 2, 5, and 11. However, there is some evidence that the data prefers not to assume a unit root in z_t. To demonstrate this, we find the Bayes factor in favor of Model III over Model I.

We construct this Bayes factor using the Savage-Dickey density ratio introduced by Dickey (1971), which allows us to easily build a Bayes factor for nested models. Specifically, the Savage-Dickey density ratio can be used to construct the Bayes factor when one model can be converted to the other model by setting a parameter to a specific value. In the case of Models I and III, we can see that Model III reduces to Model I when $\rho_z \equiv 1$.[9] When

---

[9]An additional assumption that must be satisfied to use the Savage-Dickey density ratio is that the priors be "separable." A sufficient condition for this (see, for example, Verdinelli and Wasserman, 1995) in the case of models I and III is that $p_{III}(\rho_z, \hat{\theta}) = p_{III}(\rho_z)p_I(\hat{\theta})$ where $\hat{\theta}$ is the vector of estimated parameters excluding ρ_z and $p_i(\cdot)$ is the prior distribution under model i.

Figure 11: r^* Path (Model IV)

NOTES: The estimated path of r^* under Model IV when ρ_z, μ_g and ρ_g are all estimated. The solid blue line shows the median path of the smoothed (two-sided) estimate and the blue-shaded area is bounded by the 10$^{\text{th}}$ and 90$^{\text{th}}$ percentiles of the estimated path. The red line is the two-sided r^* path for the U.S. from Holston, Laubach, and Williams (2016). The vertical shaded bars represent NBER-dated recessions.
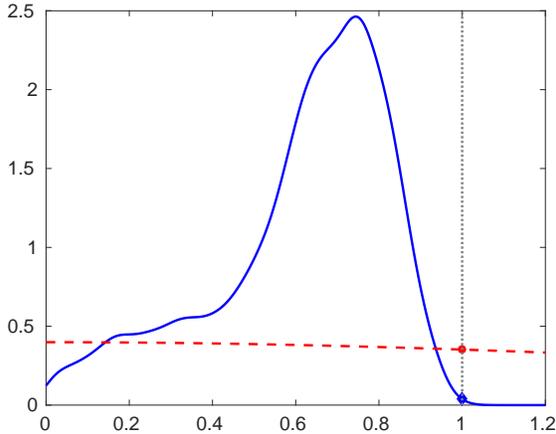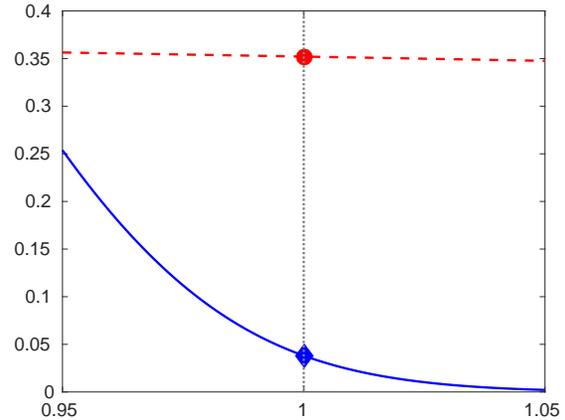
this is the case, the Bayes factor can be written in terms of the output of the estimation process for the *unrestricted* model, in our case model III:

$$B_{III,I} = \frac{p_{III}(\rho_z = 1)}{p_{III}(\rho_z = 1|Y)}$$

where $p_{III}(\rho_z = 1|Y)$ is the value of the pdf of the marginal posterior distribution for ρ_z under Model III at $\rho_z = 1$, and $p_{III}(\rho_z = 1)$ is the value of the pdf of the prior on ρ_z evaluated at 1, also under Model III. Both of these objects are readily available.[10] The concept is illustrated in Figures 12 and 13.

We find that $B_{III,I} = 9.2$, which according to the table in Appendix B of Jefferys (1961),

---

[10]The draws from the marginal posterior of ρ_z are used to build a smoothed marginal density from which we can determine the pdf value at 1. The prior is a truncated normal distribution so the value is easily calculated.

Figure 12: Prior and Posterior of ρ_z Figure 13: Area around $\rho_z = 1$

NOTES: An illustration of the Savage-Dickey density ratio. Figure 12 shows the marginal posterior distribution of ρ_z under Model III (the solid blue line) and the prior distribution over the same interval (the dashed red line). The vertical gray dashed line indicated where $\rho_z = 1$. Figure 13 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.

is "substantial" evidence in favor of model III. Kass and Raftery (1995), who develop their own scale for Bayes factors label this as "positive" evidence in favor of Model III.[11] We also considered an alternative prior distribution, $\rho_z \sim N(1, 0.5^2)$. Under this prior, we are placing much more weight on the possibility that shocks to the non-growth component of r^* are permanent, but it had very little effect on the marginal posterior distribution of ρ_z while increasing the prior probability significantly. From a model comparison standpoint using Bayes factors, that means:

$$B_{III,I} = \frac{p_{III}(\rho_z = 1)}{p_{III}(\rho_z = 1|Y)} \approx \frac{0.352}{0.038} = 9.2$$

$$\hat{B}_{III,I} = \frac{\hat{p}_{III}(\rho_z = 1)}{\hat{p}_{III}(\rho_z = 1|Y)} \approx \frac{0.816}{0.048} = 17.1,$$

where $B_{III,I}$ is the factor under the original prior and $\hat{B}_{III,I}$ is the factor under the alternative prior. That is, when we rewrite the prior to give more weight to the hypothesis of permanent

---

[11]In 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."

shocks to the non-growth component of r^*, the changes to the marginal posterior of ρ_z are almost imperceptible. As a result, the alternative Bayes factor increases as the data continue to prefer stationary processes for z_t despite the tighter prior around the unit root hypothesis.

While the evidence is not conclusive, it favors a model in which z_t may be stationary. Similar tests are insignificant for g_t (particularly when ρ_z and ρ_g are jointly estimated in Model IV), as one might expect based on the posterior weight on $\rho_g = 1$ in Figure 4. We note that these tests are only tests against a certain type of permanent shocks. We are not testing, for example, against the possibility that z_t or g_t have structural breaks.

4.2   Discussion of the nature of z_t

The remainder of the section discusses some implications of the results of our study. These results raise the importance of the non-growth component of r^*, z_t, the "special sauce" of Laubach and Williams (2003) (as referred to by Williams, 2015). Given the importance of the permanent shocks in z_t to the secular decline in r^* since 2000, it becomes crucial to understand what economic phenomena is driving the z_t process.

One connection to economic theory is to return to the original Laubach and Williams (2003) motivation of deriving the r^* formulation from a linearized Euler equation. Thus, for a stochastic discount factor (SDF) S_t, we could write:

$$e^{-r_t^*} = E_t\left[S_{t+1}\right] \tag{4.1}$$

and then consider an SDF that diverges from that of log-utility by an extra term Z_t, similar to the methodology of Campbell and Cochrane (1999) and others. For example:

$$r^* = \log E_t\left[\frac{C_{t+1}}{C_t}Z_{t+1}\right] = \log E_t\left[e^{g_{t+1}+z_{t+1}}\right] \approx E_t\left[g_{t+1} + z_{t+1}\right], \tag{4.2}$$

where z_t can be interpreted as an asset pricing term that measures the separation from log utility of the SDF. We interpret this as a theoretically-founded way of using z_t to talk about headwinds. In this case z_t can represent any state-variable that modifies the expected ratio of marginal utilities of the representative agent, in excess of expected output growth. Many

19

possibilities could be entertained, for example: changes in aggregate risk-aversion, changes in aggregate wealth, changes in the probability of a large shock ("disaster"), or the lingering effects of a financial crisis.

Still, our analysis reveals that despite the nature of the z_t process, the existence of a unit-root is necessary to obtain a subdued level of the estimate of the natural rate of interest since the great recession. Also the common interpretation of z_t as financial "headwinds", together with our analysis, would imply that these headwinds would need to be "permanent" in nature to justify a low level of the estimated r^*. Further study is needed to better understand the best way to micro-found and analyze the properties of z_t. Asset prices would seem to be a reasonable avenue of research and are the subject of ongoing research by the authors.

5   Conclusion

This paper provides empirical estimates of the natural rate of interest based on the well-known Laubach and Williams (2003). We extend their work in two ways: we estimate all the model parameters jointly and we explore alternative specifications for the unobserved processes that compose the natural rate. By using Bayesian methods to fully incorporate parameter uncertainty in a single-step estimation procedure, we obtain richer time-series dynamics of r^* under the Holston, Laubach, and Williams (2016) model specification. Our estimation shows deeper drops during recessions—and subsequently larger increases during recoveries—than the estimates obtained from the standard 3-step MLE procedure of Stock and Watson (1998). Our median path of the natural rate also shows a different trajectory since the end of the great recession and we obtain an increase since the trough in 2008, in contrast to the "non-recovery" displayed of the U.S. estimate in Holston, Laubach, and Williams (2016). We explore alternative specifications without permanent shocks on the non-growth component of r^* and find an elevated level of the median estimate after the great recession, roughly 1.5% higher than that of Holston, Laubach, and Williams (2016) in

the third quarter of 2016. The dynamics of the non-growth component are hard to estimate, a finding which mirrors the original results of Laubach and Williams (2003) as well as Kiley (2015). When this process is stationary, we estimate a greater recovery of the natural rate since the lows of the great recession, reaching 1.8% at the end of the third quarter of 2016. We therefore infer that permanent shocks to the non-growth component of the natural rate are needed to produce a persistent low level of r^* after the great recession.

References

Bomfim, A. N. (1997, October). The Equilibrium Fed Funds Rate and the Indicator Properties of Term-Structure Spreads. *Economic Inquiry 35*(4), 830–846.

Campbell, J. Y. and J. Cochrane (1999, April). Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy 107*(2), 205–251.

Chib, S. and E. Greenberg (1995). Understanding the metropolis-hastings algorithm. *The American Statistician 49*(4), 327–335.

Chib, S. and E. Greenberg (1996, August). Markov Chain Monte Carlo Simulation Methods in Econometrics. *Econometric Theory 12*(03), 409–431.

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.

Geweke, J. (2005). *Contemporary Bayesian Econometrics and Statistics*. Wiley Series in Probability and Statistics. Wiley.

Hamilton, J. D., E. S. Harris, J. Hatzius, and K. D. West (2016, November). The Equilibrium Real Funds Rate: Past, Present, and Future. *IMF Economic Review 64*(4), 660–707.

Herbst, E. and F. Schorfheide (2015). *Bayesian Estimation of DSGE Models*. Princeton: Princeton University Press.

Holston, K., T. Laubach, and J. C. Williams (2016, June). Measuring the natural rate of interest: International trends and determinants. Working Paper Series 2016-11, Federal Reserve Bank of San Francisco.

Jefferys, H. (1961). *Theory of Probability* (3 ed.). Oxford University Press.

Johannsen, B. K. and E. Mertens (2016, April). A Time Series Model of Interest Rates With the Effective Lower Bound. Finance and Economics Discussion Series 2016-033, Board of Governors of the Federal Reserve System (U.S.).

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.).

Kiley, M. T. and J. M. Roberts (2017). Monetary policy in a low interest rate world. Brookings papers on economic activity, Brookings.

Kim, C.-J. and J. Kim (2013, October). The 'Pile-up Problem' in Trend-Cycle Decomposition of Real GDP: Classical and Bayesian Perspectives. MPRA Paper 51118, University Library of Munich, Germany.

Laubach, T. and J. C. Williams (2003). Measuring the natural rate of interest. *Review of Economics and Statistics 85*(4), 1063–1070.

Laubach, T. and J. C. Williams (2016). Measuring the natural rate of interest: Redux. *Business Economics 51*, 257–267.

Lubik, T. A. and C. Matthes (2015). Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches. *Richmond Fed Economic Brief 10*(Oct), 1–6.

Pescatori, A. and J. Turunen (2016). Lower for longer: Neutral rate in the u.s. *IMF Economic Review 64*(4), 708–731.

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.

Tierney, L. (1994, 12). Markov chains for exploring posterior distributions. *The Annals of Statistics 22*(4), 1701–1728.

Verdinelli, I. and L. Wasserman (1995). Computing bayes factors using a generalization of the savage-dickey density ratio. *Journal of the American Statistical Association 90*(430), 614–618.

Williams, J. C. (2015). Measuring the natural rate of interest. Presentation delivered at The Brookings Institute, Washington, D.C., October 30.

Williams, J. C. (2016, August). Monetary Policy in a Low R-star World. *FRBSF Economic Letter 23*, 1–6.

Williams, J. C. (2017, February). Three Questions on R-star. *FRBSF Economic Letter 5*, 1–5.

Yellen, J. L. (2015). The economic outlook and monetary policy. Speech delivered at the Economic Club of Washngton, Washington, D.C., December 2.

A State Space Structure

Based on the system of equations (2.1) to (2.6), substituting the formula for r_t^* into the output gap equation (2.1) 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.

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

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

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

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

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

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

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:

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

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.