The Effective Debt Limit in the United States

Parker Sheppard

The Heritage Foundation

July 26, 2022

Abstract

The Congressional Budget Office's Long-Term Budget Outlook forecasts federal debt held by the public to rise to about 200 percent of GDP by 2050. Such a large increase raises the possibility that the federal government may default on a portion of its debt. This paper uses a non-linear DSGE model to estimate the effective debt limit – the point where markets anticipate default is likely in the near future and interest rates rise rapidly. The results indicate that there is likely plenty of fiscal space before the debt reaches the effective limit. However, if the debt continues to grow as projected, it may reach the effective limit in about twenty years.

JEL codes: H63 Debt, Debt Management, Sovereign Debt; G01 Financial Crises, E13 General Aggregative Models: Neoclassical.

Nothing written here is to be construed as necessarily reflecting the views of The Heritage Foundation or as an attempt to aid or hinder the passage of any bill before Congress.

1 Introduction

The Congressional Budget Office's (CBO) most recent Long-Term Budget Outlook expects federal debt held by the public to rise to 194 percent of Gross Domestic Product by 2050. Debt held by the public is currently near its historical high from World War II around 100 percent of GDP. To practically double the debt in thirty years would be well outside historical experience.

Despite the CBO forecast, Congress does not appear to have any hesitancy about adding to the debt, passing several major bills funded by deficits in recent years. In 2017, Congress passed the Tax Cuts and Jobs Act, which the CBO projected to add $1.4 trillion to the debt over 10 years. In 2020, Congress passed the Coronavirus Aid, Relief, and Economic Security Act, which the CBO projected to add $1.7 trillion to the debt in the 10-year budget window. In 2021, Congress passed the American Rescue Plan Act, which CBO estimated would add $1.9 trillion to the debt over the budget window.

The prospect of rapidly rising debt and the lack of a political consensus to stabilize the deficit raises the question of the possibility of a default by the federal government. Because the federal government has historically paid its debts and maintained a low level of debt relative to income, little is known about how high the U.S. federal debt might climb before either an outright default or a financial crisis precipitated by the prospect of default occurs.

Looking to international experiences, the debt-to-GDP ratio that corresponds to default can vary by country. Reinhart et al. (2003) document episodes of default occurring when debt-to-GDP levels range from 0.31 to 2.14. They conclude that differing countries have differing willingness to commit to the fiscal retrenchment that is required to pay down debt. Thus, while some countries may default at low values of debt, others can allow debt to rise and remain committed to paying it down.

CBO provides Congress with a long-term forecast of the debt, but does not try to forecast the prospects for a debt crisis. Their forecast assumes that the debt is always repaid and can always be rolled over, though the text does mention the possibility of a default crisis rising with the debt.

CBO does anticipate that interest rates on federal debt will rise as the debt increases (Gamber, 2020). They estimate an increase in interest rates of a quarter percentage point for every 10 percentage point increase in the debt-to-GDP ratio (Gamber and Seliski, 2019). However, such a linear relationship will underestimate the speed with which a debt market can switch from standard behavior to crisis. If markets anticipate that the federal government will not commit to repaying the real value of debt, they will require substantially higher interest rates to continue to roll over outstanding debt. International experience shows that the onset of a crisis can occur quickly. A linear model may provide policy makers with a false sense of complacency.

A non-linear model captures important dynamics as debt reaches crisis levels. Bi and Traum (2012, 2014) show how to incorporate sovereign default into a non-linear dynamic, stochastic general equilibrium model (DSGE). They fit their model to cases in Greece and Italy, who both experienced stress in the market for their sovereign debt, particularly with Greece undergoing a restructuring. The non-linear model showed improved forecasting of interest rates during the crisis periods relative to a linearized model.

This paper applies the model of Bi and Traum to the United States to answer the above questions. It also extends the model to account for the decline in interest rates, both nominal and real, that has occurred in recent history. Bauer and Rudebusch (2020) show that a trend is necessary to describe the behavior of real interest rates, which poses a challenge for adapting the data to a stationary model. This paper uses external estimates of the trend real interest rate to account for the apparent non-stationary behavior. Using external estimates is a parsimonious way to incorporate a phenomenon caused by many potential factors while retaining an economic interpretation that allows for identification of default probabilities.

The rest of the paper proceeds as follows: Section 2 describes the real DGSE model used. Section 3 describes the estimation procedure and econometric choices made in estimation. Section 4 describes the results. Section 5 puts the results in context and provides a concluding summary.

2 Model

The economy consists of a household sector with a representative agent and a government. Production of a single composite good, y_t depends on the amount of labor supplied by households, n_t, and productivity, a_t. The stochastic process for productivity is

$$a_t = (1 - \rho^a)\bar{a} + \rho^a a_{t-1} + \epsilon_t^a, \quad \epsilon_t \sim N(0, \sigma_a^2), \tag{1}$$

where $\bar{a}$ is the steady-state value of productivity, ρ^a is the AR(1) coefficient, and ϵ_t^a is an i.i.d. productivity shock.

Output is devoted to either private consumption, c_t or government purchases, g_t. Therefore, the economy's resource constraint satisfies

$$c_t + g_t = a_t n_t \equiv y_t, \tag{2}$$

where y_t is another way to write output.

The government's expenditures are purchases of the consumption good and lump-sum transfers to the household sector, z_t. It finances its expenditures by taxing labor income at rate τ_t and by issuing one-period bonds, b_t. Bonds issued at period t are sold for q_t units of the consumption good and pay one unit of the consumption good at time $t+1$. Additionally, at time t the government may default on a fraction of its outstanding liability b_{t-1}. The fraction defaulted is represented by the random variable Δ_t. The total amount of debt outstanding after default is denoted $b_t^d \equiv$

$(1-\Delta_t)b_{t-1}$. The government's budget is

$$(1-\Delta_t)b_{t-1} + g_t + z_t = \tau_t a_t n_t + q_t b_t, \tag{3}$$

where sources of revenue from taxes and new debt issuance cover expenditures on repaying debt, purchases of the final good, and transfers to the representative household.

Whether the government defaults in each period depends on on the debt-to-GDP ratio from the prior period, $s_t \equiv b_t/y_t$. If the debt-to-GDP ratio exceeds the debt limit, $\tilde{s}_t$, then the government defaults. That is,

$$\Delta_t = \begin{cases} 0 & \text{if } s_{t-1} < \tilde{s}_t \\ \delta & \text{if } s_{t-1} \geq \tilde{s}_t. \end{cases} \tag{4}$$

To capture the idea that the agents in the model do not know the exact debt limit, $\tilde{s}_t$ is taken to be stochastic. Its distribution has a cumulative distribution function that is a logistic function with shape parameters η_1 and η_2.

$$p_{t-1} \equiv P(s_{t-1} \leq s_t^*) = \frac{\exp(\eta_1 + \eta_2 s_{t-1})}{1 + \exp(\eta_1 + \eta_2 s_{t-1})}, \tag{5}$$

where p_{t-1} is the default probability associated with s_{t-1}.

The government adjusts fiscal policy to target a desired level of debt, $\bar{b}$. Adjustments depend on the current level of outstanding debt, previous levels of taxation and spending, and stochastic shocks,

$$\tau_t = (1-\rho^\tau)\bar{\tau} + \rho^\tau \tau_{t-1} + \epsilon_t^\tau + \gamma^\tau(b_t^d - \bar{b}), \quad \epsilon^\tau \sim N(0, \sigma_\tau^2) \tag{6}$$

$$g_t = (1-\rho^g)\bar{g} + \rho^g g_{t-1} + \epsilon_t^g - \gamma^g(b_t^d - \bar{b}), \quad \epsilon^g \sim N(0, \sigma_g^2), \tag{7}$$

$$z_t = (1-\rho^z)\bar{z} + \rho^g z_{t-1} + \epsilon_t^z, \quad \epsilon^z \sim N(0, \sigma_z^2), \tag{8}$$

where $\gamma^\tau > 0$ and $\gamma^g > 0$ are coefficients that describe the speed of adjustment to deviations from the target debt level. The process for transfers does not contain a term related to debt, nor are transfers used as an observable in estimation, so the series can be thought of as a residual capturing movements in debt not explained by taxes or spending.

A representative household chooses consumption, c_t, labor hours n_t, and bond purchases b_t each period to solve

$$\max_{c_t,n_t,b_t} E_t \sum_{j=t}^{\infty} \beta_t^j \left(\ln(c_j - hc_{j-1}) + \phi \ln(1 - n_t) \right), \tag{9}$$

$$\text{subject to } c_t + q_t b_t = (1 - \tau_t) a_t n_t + (1 - \Delta_t) b_t + z_t, \tag{10}$$

where h is a habit parameter, ϕ is a parameter determining the weight placed on leisure, and β_t is the time t discount factor. The inverse of β_t corresponds to the interest rate in the model that would prevail once the effects of all shocks have dissipated. The time subscript indicates that the household's discount factor may change over time. However, changes in the discount factor are unanticipated, and the household solves its optimization problem as if it will not change in the future. Therefore, the process for how β_t changes over time is not needed to solve for equilibrium in any given period and is not specified. More detail about how changes in β_t are accounted for is in the section on estimation.

A solution to the household's problem satisfies the first-order conditions

$$\phi \frac{c_t - hc_{t-1}}{1 - n_t} = a_t (1 - \tau_t), \tag{11}$$

$$q_t = \beta_t E_t \left((1 - \Delta_{t+1}) \frac{c_t - hc_{t-1}}{c_{t+1} - hc_t} \right), \tag{12}$$

and the transversality condition

$$\lim_{j\to\infty} E_t\left(\beta_t^{j+1}\frac{u_{c,t+j+1}}{u_{c,t}}(1-\Delta_{t+j+1})b_{t+j}\right)=0, \tag{13}$$

where $u_{c,t}$ is the derivative of the utility function in a single period. The transversality condition requires the value of the households' debt position to be zero at the infinite horizon, preventing it from infinitely delaying receiving payment.

Model Solution

A solution to the model consists of decision rules for c_t, n_t, and b_t that that satisfy the household's first-order conditions, the resource constraint, the government's budget equation, and the stochastic processes for productivity and fiscal policy given the values of the states $\psi_t = [b_t^d, c_{t-1}, a_t, g_t, \tau_t, z_t, \beta_t]$. Given that there is only one true market in the model, I solve for the debt policy rule that clears the debt market and use the system of equations to find the remaining variables.

The solution procedure follows Bi and Traum. The solution is obtained by solving the model over a discrete grid of points in the state space. The process starts by guessing that debt remains constant. Productivity, labor supply, and the resource constraint determine the household's current consumption. The bond price is derived by evaluating the household's Euler equation using Gauss-Hermite quadrature over shocks to productivity, the tax rate, government spending, and default. The bond price and the government budget equation imply a new value for debt. The process is repeated until the difference in debt values is within a chosen tolerance.

3 Estimation

Data and Calibrated Values

The data sample used for estimation is from 1984:1-2019:4. The start date is chosen to be within the Great Moderation, while the end date is just before the disruptions from the outbreak of Covid-19. The sample covers 144 quarters and includes four expansions and three recessions.

The data are converted to stationary versions before being used in estimation. In making conversions, I deviate somewhat from Bi and Traum, who convert their data by expressing it as percentage deviations from the sample mean, or from a linear trend in the case of real GDP. I choose to do as few calculations on the raw data as possible in order to keep modeling decisions clear. Centering the data on mean values expresses a view about the central tendency, so I explicitly report when calibrated steady state values differ from the sample mean.

The time series for debt relative to GDP illustrates the issues motivating this choice very well. I take as the measure of debt federal debt held by the public and subtract federal debt held by Federal Reserve Banks, which treats the Federal Reserve as part of the government. The series shows substantial variation over the sample period, spending very little time near the sample mean. It starts at 0.27 in 1984, rises to a peak of 0.43 in 1994, then falls to a trough of 0.26 in 2001, rising slowly before a marked jump from 0.29 to 0.60 between 2007 and 2012, and rising slowly thereafter. Though excluded from the sample used for estimation, the debt jumps from 0.63 to 0.81 times GDP between the first and second quarters of 2020 with the brief recession and stimulus spending accompanying the outbreak of Covid-19. With recovery following the end of the most restrictive measures, the debt ratio had fallen to 0.70 by the third quarter of 2021.

Using the average debt level from 1984-2019 would imply that debt was below target from 1984-2007, or even from the end of the Second World War until 2007. By inferring a target that no observer would have chosen prior to 2008, it would skew estimation of the fiscal policy parameters to smaller values as the debt fell away in the '90s from a target that did not exist.

I choose to set the debt-to-GDP target $\bar{b}$ at 0.35. That level of debt fits the early part of the sample by the eyeball test and is consistent with the post-WWII average. Moreover, I define the average residual z_t by first defining the average spending and tax targets and solving for the residual consistent with maintaining the debt target. The fact that the resulting residual matches the average sample residual supports the choice of 0.35.

In doing so, I interpret the sharp increases in debt following the American Recovery and Reinvestment Act (ARRA) and the Coronavirus Aid, Relief, and Economic Security Act (CARES Act) as large, persistent deviations from the debt target. An alternative view could be that the debt target itself has shifted. Modeling that scenario is significantly more complicated and is left for future work.

For government spending and tax rates, I use the values relative to GDP as observable data. I set the mean for government spending at 5.9 percent of GDP and the mean for tax revenue at 18 percent of GDP, which are both near the sample averages.

For the output series, I use the log deviations of real GDP from potential GDP as measured by the CBO. Fitting a log-linear time trend to real GDP over the sample, which includes the sharp drop in output from the 2007-09 recession produced a stationary series that was either significantly above or below the linear trend. Using potential GDP as an approximation of the steady state instead of a linear time trend reflects variation at the business-cycle frequency better.

Additionally, I calibrate the steady state of hours worked, $\bar{n}$, to be 0.2. This is consistent with average weekly hours around 34 out of 168. Bi and Traum use 0.25, but also report that their estimates favor lower values of h. I tested both 0.2 and 0.25 and found that the higher calibration for hours favored the lower values for h as well. The model fixes the Frisch elasticity of labor supply, which removes a dimension for labor supply to vary independently of the habits parameter (ϕ has to correspond to the steady state value for labor in the labor supply equation in equilibrium). The lower value for $\bar{n}$ is more consistent with the prior for h chosen from the literature.

Finally, the values within the model for the discount factor are based on exogenous estimates of the trend real rate with some observation error,

$$\beta_t = \frac{1}{r^*t} + \epsilon_t^{\beta}, \quad \epsilon_t^{\beta} \sim N(0, \sigma_{\beta}^2) \tag{14}$$

The values for the trend real interest rate are taken from the model of Laubach and Williams (2016), as reported by Bauer and Rudebusch (2020). Bauer and Rudebusch produce a number of estimates of the trend real rate, but that series appeared to be the best fit to the present model in concept. The discount factor determines the interest rate in the model's steady state, where the system rests after all transitory shocks die out. The Laubach and Williams estimates are calculated to be the neutral real rate that occurs after all transitory shocks have died out. The value for σ_{β}^2 is calibrate to match the error bounds reported by Bauer and Rudebusch.

The model is cast in the state-space form

$$x_t = f(x_{t-1}, \epsilon_t, \theta), \tag{15}$$

$$v_t = x_t + \xi_t, \tag{16}$$

where x_t is the unobserved vector of state variables, ϵ_t is the vector of model shocks, θ is the vector of model parameters, v_t is the vector of model observables, and ξ_t is a vector of normally-distributed model and measurement errors with mean 0 and variance matrix Σ. I calibrate the standard deviation on measurement errors through an iterative process. A common value in the literature is to use a standard deviation that is 20% of the standard deviation of the data (for instance, Smets and Wouters (2007)). However, that rule of thumb is often used in models where the observables are expressed as growth rates, where the observables here are in levels. Thus, I start with standard deviations at 20% and run a test estimation on a small sample. If the model errors implied by the data and the filtered estimates at the posterior mean are greater than the calibrated values, I increase the value and repeat the process. Correspondingly, if the estimates are smaller, then I tighten the values to aid in estimation.

Identification

The key equation for estimation is the household's Euler equation, which determines the price of government bonds. The equation can be written as

$$
\begin{aligned}
q_t = (1-p_t)\beta_t E_t & \left(\frac{u'_{c,t+1}}{u'_{c,t}} | \text{no default} \right) \\
+ p_t \beta_t E_t & \left((1-\delta) \frac{u'_{c,t+1}}{u'_{c,t}} | \text{default} \right),
\end{aligned} \tag{17}
$$

where $u'_{c,t}$ is shorthand for the household's marginal utility with respect to consumption. This equation shows how the observed interest rate can be broken into a trend component, a cyclical component, and a default premium. The strategy for estimation will be to account for the trend rate and the cyclical component, then fit the probability distribution implied by the debt level to the residual.

Having an observable factor is necessary for identifying default probabilities. Without an observable factor to identify movements in the discount factor, it is possible to conflate those movements with changes in default probabilities.

Within the model, the discount factor determines the interest rate in the steady state. Laubach and Williams (2016) estimate a trend real rate that corresponds to a similar definition – the interest rate that prevails when all the cyclical shocks have dissipated. Thus, an estimate of the trend real interest rate would provide an observable that could be used for identification here.

For estimation I use

$$\beta_t = \frac{1}{{r^*}_t} + \epsilon_t^\beta, \quad \epsilon_t^\beta \sim N(0, \sigma_\beta^2), \tag{18}$$

where r_t^* corresponds to the observed trend real rate and ϵ_t^β is measurement or estimation error. The effect this has on estimation is explained more in the next subsection.

An alternative approach would be to incorporate the trend estimates in the process of preparing the data to conform to the model. However, doing so treats the point estimates of the trend rate as observations rather than estimates, which could potentially introduce measurement error into the model that would bias the results. Using the process for the discount factor specified above treats the trend rate as a hidden state of the model. Thus the model incorporates any measurement or estimation error for the trend real rate. Given the error bounds and continuing debate over the nature of trend real rates, The model makes no attempt to explain the behavior of trend real rates, only to control for them in the process of pricing government debt.

Parameter Estimation

I use the tempered particle filter of Herbst and Schorfheide (2019) to evaluate the likelihood of each parameter. Like other particle filters, the tempered particle filter uses simulated series to approximate the model state at each time period. It adds a series of tempering steps to the evaluation of one-step-ahead forecasts, starting from an inflated model/observation error variance, and tempering down to the calibrated value. With each tempering step, the particles are resampled based on their closeness to the observed value, which improves the accuracy of forecasts. I use 10,000 particles and a target inefficiency ratio of 5. The values were chosen after a series of tests evaluating the trade-off between run-time and Monte Carlo error.

The estimation for the posterior distribution uses the simulated tempering sequential Monte Carlo method of Herbst and Schorfheide (2014). The target of the estimation is the posterior distribution of the model parameters, which is the product of the prior distribution and the likelihood. SMC starts with a draw from the prior distribution and produces a sequence of bridge distributions, gradually placing more weight on the likelihood with each iteration.

I use a λ of 3, which is greater than the choice of 2 by Herbst and Schofheide. This further delays introducing information from the likelihood too quickly and allows the estimation to explore the parameter space more during the initial sequences. I use $N_\phi = 150$ tempering steps. I use $N_b = 1$ block at each step, given that obtaining acceptances did not appear to be a problem. The total sample size consists of 5,000 parameter vectors.

The proposal distribution is just a normal centered on the current parameter value at each step, and the tuning parameter in proposals was chosen to be 0.25. Using the mixture distribution in Herbst and Schorfheide that favors proposals in the direction of the sample mean tended to produce posteriors that shrunk toward the sample mean. The proposal used here is not biased toward any particular value, so movement in that direction can be judged to be due to improving likelihood.

Resampling occurred when the effective sample share fell below 0.80 of the total sample. Herbst and Schorfheide use a value of 0.5, but I found that value would often produce importance weights that were quite skewed, an indication of high variance in the sampling weights that could lead to a biased and unstable estimator (Koopman et al., 2009). Resampling tended to capture about 60% of the sample with unique values and 40% as duplicates.

The code was implemented using Matlab's Parallel Computing Toolbox. Running on a server with 24 threads and 25 gigabytes of memory, the total run time was just over 767 hours.

Prior Distributions

The prior means for ρ_a, ρ_g, γ_g, ρ_τ, γ_τ, and ρ_z are taken from AR(1) regressions on the sample data, and the prior standard deviations are set to match the associated standard errors. The means for σ_a, σ_g, σ_τ, and σ_z are from the associated standard error of the residuals, and their prior standard deviations are set to about 40 percent of the mean values.

To facilitate comparison to other models, I report the variances scaled relative to the steady state values of the corresponding variables,

$$\sigma_{k,p} = \frac{\sigma_x}{\bar{k}}, \quad \text{for } k \in \{a, g, \tau, z\}.$$

The parameters for policy adjustment are reported as elasticities,

$$\gamma_{k,e} = \gamma_k \frac{\bar{b}}{\bar{k}}, \quad \text{for } k \in \{g, \tau\}.$$

For example, the prior median of $\gamma_{\tau,e}$ of roughly 0.005 means that for a one percent deviation of debt from its target, the government would raise the deviation of the tax rate from its steady state by one-half percent above the policy it would otherwise set.

The prior for s^{low} is taken to be uniform on the interval from 0.5 to 1.5. The uniform distribution was chosen to be uninformative. The lower bound is in the middle of the observed range, and should be well below realistic values for the limit. The upper end is about twice the maximum observed value, which is sufficiently far away that it becomes difficult to distinguish between higher values for the effective debt limit. Thus, prior for the effective debt limit is not a reflection of the state of knowledge derived from the literature, and the posterior should not necessarily be interpreted as reflecting the best combination of the information in the literature and the information in the data. The prior is an uninformative distribution over a range where the model is likely to be informative, and the posterior shows the extent to which values in that range are compatible with the data.

4 Results

Table 1 shows the medians and the 90% confidence intervals for the model parameters of interest. Figure 1 plots the priors and the posteriors. The posterior for the effective debt limit shows a sizable shift relative to the prior. Most of the lower values are eliminated as inconsistent with the data, raising the fifth percentile from 0.55 to 0.81. The median rises from 1.00 to 1.17 as the bulk of the distribution shift to higher values.

However, while the data is somewhat informative about the distribution of the effective debt limit, the distribution still covers a wide range. The range between the 5th and 95 percentile is still 0.65, or approximately the value of debt held by the public less debt held by the Federal Reserve in the first quarter of 2020. Using the most recent GDP estimate (2021Q3), that corresponds to approximately $15.1 trillion.

The wide range of the posterior is not surprising given the nature of the problem. The model is trying to estimate the behavior of interest rates for debt-to-GDP ratios well outside those observed in the sample. The non-linear form of the default risk used in the model means that default risk premia are very low for debt values far below the effective debt limit, and therefore can only provide limited information. Still, the model is able to rule out very low values for the effective debt limit, since those values would have produced noticeable risk premia increases not observed in the data.

Table 2 shows the mean and standard error of the model observation error using the posterior mean. The mean errors are generally close to zero, indicating that the model on average fits the data. All of the estimated standard deviations are less than the calibrated standard deviations, suggesting that the calibrated values are not a constraint on the estimates. Observed data and the filtered series produced by the model are show in Figure 2.

Figure 3 plots the distribution of default probabilities estimated by the model over the sample data. The distribution of the probability of default is derived from the posterior distribution of the effective debt limit and the observed history of the debt-to-GDP ratio. Overall, the median probability of default is quite low, peaking at less than 0.1 percent at the end of the sample. The non-linear model of default implies that the distribution of default probability at any given time is skewed to the right. The 75th percentile peaks at around 0.6 percent, while the 95th percentile (not plotted) peaks around 6 percent. Thus, the historical estimates suggest that the chance of a default has been and remains quite low, though the tail of the effective debt limit distribution implies a higher risk.

Figure 4 shows the decomposition of the interest rate into a trend rate, a cyclical risk premium, and a default risk premium averaged across the posterior distribution. The top panel shows that observed interest rates track closely with the long-run interest rate from the model. The cyclical component of the risk-free rate fluctuates around zero, spiking during recessions when forecast income growth increases. Most notable is that the average default premium is very small, less than 0.015 percent per year throughout the sample. This estimate is consistent with the widespread notion of Treasury debt as a nearly risk-free asset.

The results for the effective debt limit reported here are consistent with others in the literature. Bi and Traum's estimate of the effective debt limit for Greece using essentially the same model found a 90% confidence interval of 1.50 to 1.56. However, they use gross debt rather than debt held by the public outside the central bank, so a comparable range based on external debt is 1.11 to 1.15.[1] Ostry et al. (2010) estimate a comparable idea (termed maximum sustainable debt) for the United States using reduced form data on the difference between interest rates and growth rates. They estimate a maximum sustainable debt ratio for the U.S. of 1.61 based on gross debt, which corresponds to about 1.11 using debt held by the public less debt held by the Federal Reserve.

[1]According to Greece's Public Debt Management Agency, the private sector held 74 percent of Greek debt in 2011, the last year data is reported before the debt was restructured.

Finally, Table 3 shows the implications of the model estimates for the CBO's long-term forecast. Assuming that debt held by the Federal Reserve remains constant, the debt reaches the effective limit around 2040. Consistent with the definition of the effective limit, the median probability of default is around 30 percent, which also implies a 6 percent default risk premium. If the debt were to continue to grow past that level as in the CBO forecast, the chances of default and associated risk premium only increase futher.

5 Conclusion

This paper has estimated a non-linear DSGE model with sovereign default using data from the United States. The non-linear solution captures the sharp rise in interest rates consistent with historical observations in other countries, which would be lost by solving the model using standard log-linear approximations. Using a DSGE model allows the implications of the model to be extended outside of the observed data range, which is important given the lack of a sovereign default crisis in U.S. history. The novel contribution of this paper is an extension of existing models to allow for long-run interest rates to change over time, which is necessary to fit the model to U.S. data.

The key parameter of interest in the model is the effective debt limit – a level of debt where interest rates are increasing rapidly and default in the near future is likely. Technically, the effective debt limit is defined to be the level of debt that corresponds to a 30% probability of default. The results of paper return a median value for debt held by the public less debt held by the Federal Reserve to GDP of about 115 percent. The actual value in the third quarter of 2020 was 67 percent, well below the median estimate. However, the possibility that the effective debt limit is closer to current debt levels cannot be ruled inconsistent with the data.

Looking to the future, the CBO projects federal debt held by the public as a share of GDP to pass 140 percent in 2040. Assuming that federal debt held by the Federal Reserve remains around 25 percent of GDP, the current trajectory for federal debt would likely reach the effective debt limit in about 20 years.

While this result is informative for policymakers, plenty of work remains to refine the estimates. The model here is quite simple. Though using external estimates of the trend real rate is an effective way to reduce computation, additional features that relate the trend rate to economic factors could help provide insight for forecasts. Adding an open economy is likely necessary to fully capture factors driving the downward trend in interest rates. Accounting for the interaction between fiscal and monetary policy is also likely important.

References

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

Huixin Bi and Nora Traum. Estimating Sovereign Default Risk. *American Economic Review*, 102(3):161–66, May 2012.

Huixin Bi and Nora Traum. Estimating Fiscal Limits: The Case of Greece. *Journal of Applied Econometrics*, 29(7):1053–1072, 2014.

Edward Gamber. The Historical Decline in Real Interest Rates and Its Implications for CBO's Projections. Working papers, Congressional Budget Office, 2020.

Edward Gamber and John Seliski. The Effect of Government Debt on Interest Rates: Working Paper 2019-01. Working Papers 55018, Congressional Budget Office, March 2019.

Edward Herbst and Frank Schorfheide. Sequential Monte Carlo Methods for DSGE Models. *Journal of Applied Econometrics*, 29(7):1073–1098, 2014.

Edward Herbst and Frank Schorfheide. Tempered particle filtering. *Journal of Econometrics*, 210(1):26–44, 2019.

Siem Jan Koopman, Neil Shephard, and Drew Creal. Testing the assumptions behind importance sampling. *Journal of Econometrics*, 149(1):2–11, 2009.

Thomas Laubach and John C. Williams. Measuring the natural rate of interest redux. *Business Economics*, 51(2):57–67, 2016.

Congressional Budget Office. The 2021 Long-Term Budget Outlook, 2021. URL `https://www.cbo.gov/publication/56977`.

Jonathan David Ostry, Atish R. Ghosh, Jun I. Kim, and Mahvash S. Qureshi. Fiscal Space. IMF Staff Position Notes 2010/011, International Monetary Fund, September 2010.

Carmen M. Reinhart, Kenneth S. Rogoff, and Miguel A. Savastano. Debt Intolerance. *Brookings Papers on Economic Activity*, 34(1):1–74, 2003.

Frank Smets and Rafael Wouters. Shocks and frictions in U.S. business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606, 2007.

Appendix

Data Series

The data sample is from 1984:1 to 2019:4. There are six series used as observables.

1. *Real GDP*. Nominal GDP from the Bureau of Economic Analysis (BEA) is deflated by the GDP deflator, also from BEA, to get real GDP. Potential GDP from the Congressional Budget Office, also deflated, is used as an approximation for steady-state output. The series used as an observable is the deviations of real GDP from steady state, calculated as the log of real GDP minus the log of real potential GDP, plus 0.005, because CBO projects actual output to be below potential output over the long-run by one-half percent.

2. *Government spending*. Federal government consumption expenditures is taken from the BEA. The observable series is federal government consumption expenditures divided by nominal GDP. Note that federal government consumption expenditures is significantly smaller than total federal expenditures, which also includes transfer payments and other items.

3. *Tax revenue*. Federal current receipts is taken from the BEA. The observable series is federal current receipts divided by nominal GDP.

4. *Debt*. Federal debt held by the public and federal debt held by Federal Reserve Banks are reported by the U.S. Treasury. The observable series is nominal debt held by the public minus nominal debt held by Federal Reserve Banks, with the difference divided by nominal GDP.

5. *Interest rate*. The nominal interest rate is the rate on 10-year Treasury notes reported by the Board of Governors of the Federal Reserve System. A real interest rate is calculated by adjusting for the average expected inflation as measured by the Federal Reserve's Perceived Target Rate.

6. *Trend Real Interest Rate.* Bauer and Rudebusch (2020) report several estimates of the trend real interest rate. I use their reported values for the filtered series from the Laubach and Williams model as the observable.

Simulated Tempering SMC Algorithm

1. Draw a sample of N vectors θ_1^n from the prior distribution $p(\theta)$ and evaluate the likelihood for each one. Assign particle weights $W_1^n = 1$ for $n = 1, \ldots, N$.

2. For $i = 2, \ldots N_\phi$:

   (a) Calculate the incremental weights, $\tilde{w}_i$, and update the normalized weights, $\tilde{W}_i$, so that the average value remains 1.

$$\tilde{w}_i^n = [p(Y|\theta_{i-1}^n)]^{\phi_i - \phi_{i-1}},$$
$$\tilde{W}_i^n = \frac{\tilde{w}_i^n \tilde{W}_{i-1}^n}{\frac{1}{N}\sum_{n=1}^N \tilde{w}_i^n \tilde{W}_{i-1}^n}.$$

   (b) Calculate the effective sample size,

$$ESS_i = \frac{N}{\frac{1}{N}\sum_n \tilde{w}_i^2}.$$

      i. If $ESS_i < 0.8N$, draw a new sample $\{\hat{\theta}_i^n\}_{n=1}^N$ with replacement according to the sample weights and reset the weights $W_i^n = 1$ after re-sampling.

      ii. Otherwise, set $\hat{\theta}_i^n = \theta_{i-1}^n$ and $W_i^n = W_{i-1}^n$.

   (c) For each candidate parameter vector in the sample, apply a Metropolis-Hastings algorithm with stationary distribution $\pi_n(\theta)$.

      i. Draw a sample of mutated vectors $\hat{\theta}_i^{n,*} \sim N(\hat{\theta}_i^{n,*}, c^2\Sigma_i)$.

      ii. Calculate the probability of acceptance,

$$\alpha(\hat{\theta}_i^{n,*}, \hat{\theta}_i^n) = \min\left\{1, \frac{p^{\phi_i}(Y|\hat{\theta}_i^{n,*})p(\hat{\theta}_i^{n,*})}{p^{\phi_i}(Y|\hat{\theta}_i^n)p(\hat{\theta}_i^n)}\right\}$$

iii. Let

$$\theta_i^n = \begin{cases} \hat{\theta}_i^{n,*} & \text{with probability } \alpha(\hat{\theta}_i^{n,*}, \hat{\theta}_i^n) \\ \hat{\theta}_i^n & \text{otherwise.} \end{cases}$$

3. The final approximation of $E_\pi(h(\theta))$ is

$$\bar{h}_{N_\phi} = \sum_{n=1}^{N} h(\theta_{N_\phi}^n) W_{N_\phi}^i$$

Tempered Particle Filtering Algorithm

1. Draw a swarm of M particles in the state space, x_0.

2. For $t = 1, \ldots, T$:

   (a) Estimate the likelihood for time t with an inflated error variance.

   i. Draw a sample of shocks ϵ_t and apply the state transition equations to forecast the particle swarm,

$$\tilde{x}_t = f(x_{t-1}, \epsilon_t, \theta).$$

   ii. Set $n = 1$. Find the intitial scale ϕ_1 for the error variance the maintains the target inefficiency ratio, r^*. Start with $\phi_1 = 1$ and calculate:

   A. The associated incremental weights,

$$\tilde{w}_t^{i,1}(\phi_{1,t}) = \exp\left(-\phi_{1,t}\frac{1}{2}(y_t - \tilde{x}_t)'\Sigma_\epsilon^{-1}(y_t - \tilde{x}_t)\right)$$

   B. The normalized sampling weights in the swarm of particles,

$$\tilde{W}_t^{i,1}(\phi_{1,t}) = \frac{\tilde{w}_t^{i,1}}{\frac{1}{N}\sum_{i=1}^{N}\tilde{w}_t^{i,1}}.$$

   C. The inefficiency ratio

$$\text{Ineff}(\phi_{1,t}) = \frac{1}{M}\sum_{i=1}^{M}(\tilde{W}_t^i(\phi_{1,t}))^2.$$

   If $\text{Ineff}(\phi_{1,1}) < r^*$, set $\phi_1 = 1$ and $N_t^\phi = 1$. Otherwise, let $\phi_{1,t}$ be the solution to $\text{Ineff}(\phi_{1,t}) = r^*$ and $\tilde{W}_t^{i,1} = \tilde{W}_t^{i,1}(\phi_{1,t})$.

iii. The initial estimate of the likelihood at time t is

$$p_1(y_t|\tilde{x}_t) = \frac{1}{M}\sum_{i=1}^{M}\tilde{W}_t^{i,1}$$

iv. Resample the particles according to the normalized sampling weights,

$$\left\{\tilde{x}_t^{i,1}, \tilde{\epsilon}_t^{i,1}, x_{t-1}^{i,N_{t-1}^{\phi}}\right\} \rightarrow \left\{x_t^{i,1}, \epsilon_t^{i,1}, x_{t-1}^{i,N_{t-1}^{\phi}}\right\}$$

(b) Apply tempering steps to revise the estimate of the likelihood while reducing the variance of the errors. Do until $\phi_{n,t} = 1$:

i. Let $n = n + 1$. Update the tempering parameter ϕ_n by calculating

A. The incremental weights,

$$\begin{aligned}\tilde{w}_t^{i,n}(\phi_{n,t}) &= \frac{p_n(y_t|x_t^{i,n-1})}{p_{n-1}(y_t|x_t^{i,n-1})} \\ &= \left(\frac{\phi_{n,t}}{\phi_{n-1,t}}\right)^{\frac{n_y}{2}} \exp\left(-\left(\phi_{n,t} - \phi_{n-1,t}\right)\frac{1}{2}(y_t - \tilde{x}_t)'\Sigma_\epsilon^{-1}(y_t - \tilde{x}_t)\right).\end{aligned}$$

B. The normalized weights,

$$\tilde{W}_j(\phi_{n,t}) = \frac{\tilde{w}_t^i}{\frac{1}{M}\sum_{i=1}^{M}\tilde{w}_t^i}.$$

C. The inefficiency ratio,

$$\text{Ineff}(\phi_{n,t}) = \frac{1}{M}\sum_{i=1}^{M}(\tilde{W}_t^i(\phi_{n,t}))^2.$$

If $\text{Ineff}(1) < r^*$, let $\phi_{n,t} = 1$, $N_t^{\phi} = n$, and $\tilde{W}_j = \tilde{W}_j(1)$. Otherwise, let $\phi_{n,t}$ be the solution to $\text{Ineff}(\phi_{n,t}) = r^*$ and $\tilde{W}_j = \tilde{W}_j(\phi_{n,t})$.

ii. Resample the particles according to the normalized sampling weights

$$\{\tilde{x}_t, \epsilon_t, x_{t-1}\} \rightarrow \{\hat{x}_t, \hat{\epsilon}_t, x_{t-1}\}$$

Resampling must ensure that the state transition equation $\hat{x}_t = f(x_{t-1}, \hat{\epsilon}_t, \theta)$ is satisfied for every particle. Reset sampling weights to 1.

iii. Draw a proposed vector of shocks $\check{\epsilon}_t \sim N(0, c^2 I)$. Apply the state transition equation

$$\check{x}_t = f(x_{t-1}, \check{\epsilon}_t, \theta).$$

Calculate the ratio

$$\alpha = \min\left(1, \frac{p_n(\check{x}_t)}{p_n(\hat{x}_t)}\right).$$

Set ϵ_t according to

$$\epsilon_t = \begin{cases} \check{\epsilon}_t & \text{with probability } \alpha. \\ \hat{\epsilon}_t & \text{with probability } 1 - \alpha. \end{cases}$$

(c) Approximate the contribution to the likelihood for period t with

$$\hat{p}(y_t|Y_{1:t-1}) = p_{N_t^\phi}(y_t|Y_{1:t-1}) = \prod_{n=1}^{N_t^\phi} \left(\frac{1}{M} \sum_{i=1}^{M} \tilde{w}_t^{i,n} \right).$$

3. The likelihood of the sample is

$$\hat{p}(Y_{1:T}) = \prod_{t=1}^{T} \hat{p}(y_t|Y_{1:t-1}).$$

Table 1: Posterior Distributions

	Prior			Posterior		
	Median	5%	95%	Median	5%	95%
s^{low}	1.0046	0.5501	1.4525	1.1685	0.8085	1.4609
h	0.7083	0.5192	0.8550	0.7137	0.6295	0.7916
$\gamma_{\tau,e}$	0.0078	0.0021	0.0192	0.0064	0.0012	0.0136
$\gamma_{g,e}$	0.0150	0.0086	0.0235	0.0128	0.0065	0.0196
ρ_a	0.9653	0.9141	0.9914	0.9782	0.9494	0.9967
ρ_g	0.9826	0.9621	0.9940	0.9798	0.9624	0.9945
ρ_t	0.9269	0.8603	0.9692	0.9376	0.8891	0.9823
ρ_z	0.7635	0.6648	0.8477	0.8267	0.7435	0.9035
$\sigma_{a,p}$	0.0040	0.0018	0.0074	0.0138	0.0113	0.0166
$\sigma_{g,p}$	0.0154	0.0077	0.0271	0.0159	0.0129	0.0194
$\sigma_{t,p}$	0.0211	0.0092	0.0411	0.0214	0.0179	0.0248
$\sigma_{z,p}$	0.0709	0.0329	0.1313	0.0577	0.0456	0.0721
$\sigma_{\beta,p}$	0.0100	0.0100	0.0100	0.0100	0.0100	0.0100

The main parameter of interest, the effective debt limit, is s^{low}. The habit parameter is h. The parameters governing policy rules expressed as elasticities are γs. The first autocorrelation parameters of variables subject to shocks are ρs. The standard deviation of the shocks expressed as shares of the steady state are σs.

Table 2: Estimates of Model/Observation Error

	Output	Spending	Receipts	Debt	Real Rate
Mean of Filtered Error	0.0013	0.0067	-0.0060	-0.0436	-0.0008
SD of Filtered Error	0.1073	0.0697	0.1415	0.3005	0.0552
SD of Model Error	0.2000	0.1500	0.2000	0.5000	0.2000
SD of Observations	1.8083	0.8362	1.0677	12.0996	1.7553
Relative SD	0.1106	0.1794	0.1876	0.0413	0.1139

The observable series are the output gap, government spending relative to output, government receipts relative to output, debt relative to output, and the real interest rate. Units are percentage points, i.e. the standard deviation of the output gap is 1.80 percent. The relative SD is the calibrated value for the SD of model errors to the SD of the observations, where 0.2 is a commonly used value in the literature. This table shows that the filtered estimates are close to observed data and that the standard deviation of the errors is consistent with the calibrated values used in estimation.

Table 3: Implications of the CBO Long-Term Forecast

Year	Debt/GDP	Prob. Def.	Def. Prem.
2025	82.21	1.85	0.10
2030	82.69	2.11	0.12
2035	97.09	6.77	0.67
2040	116.68	30.04	6.27
2045	141.05	79.66	20.56
2050	169.42	98.09	22.80

This table shows the implications of debt growing according to the CBO long-term forecast, assuming that debt held by the Federal Reserve is held at 25 percent of GDP. The Debt/GDP column reports the debt held by the public, less debt held by Federal Reserve banks relative to GDP as a percentage. The next column shows the median probability of default as a percentage given the posterior distribution of the effective debt limit. The last column reports the median risk premium as an annual rate in percentage points across the posterior distribution for the effective debt limit.

Figure 1: Prior and Posterior Distributions

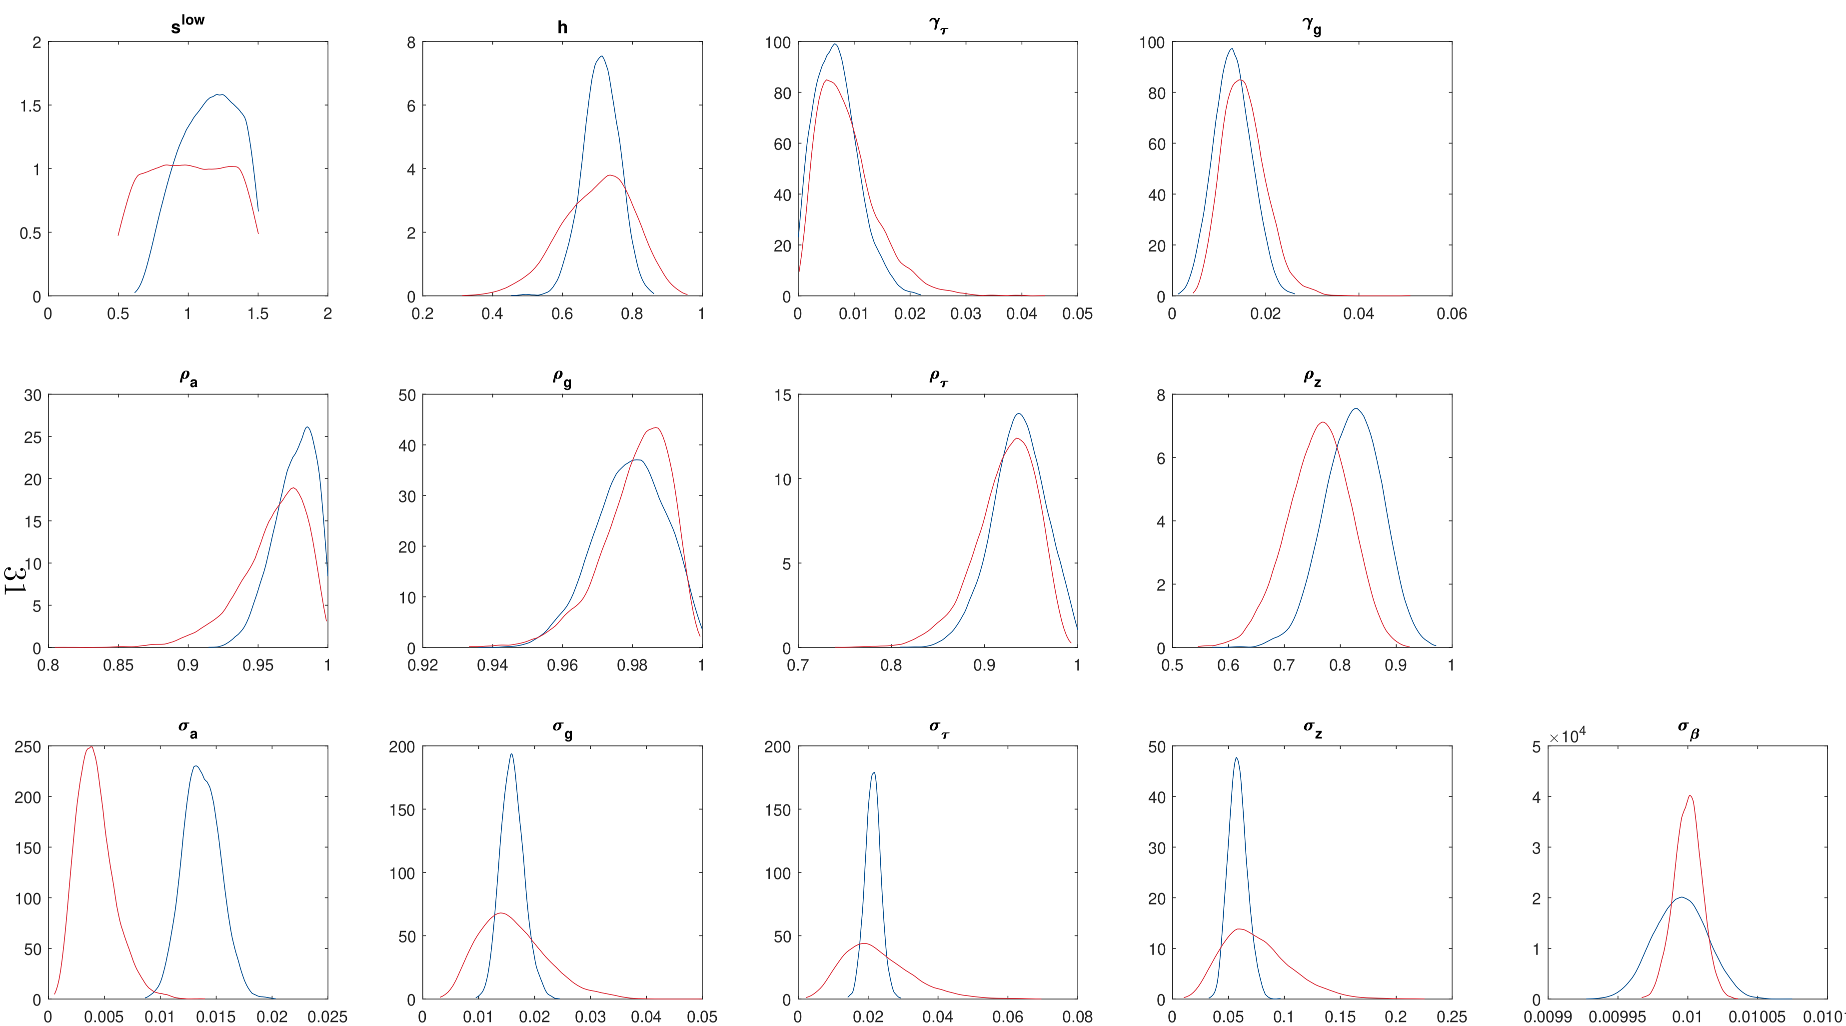

The blue line shows the posterior distribution, while the red line shows the prior distribution. For the main variable of interest, s^{low}, which corresponds to the effective debt limit, the posterior places lower weight on the lowest debt-to-GDP values, and generally shifts to the right.

Figure 2: Filtered Time Series Using Posterior Mean

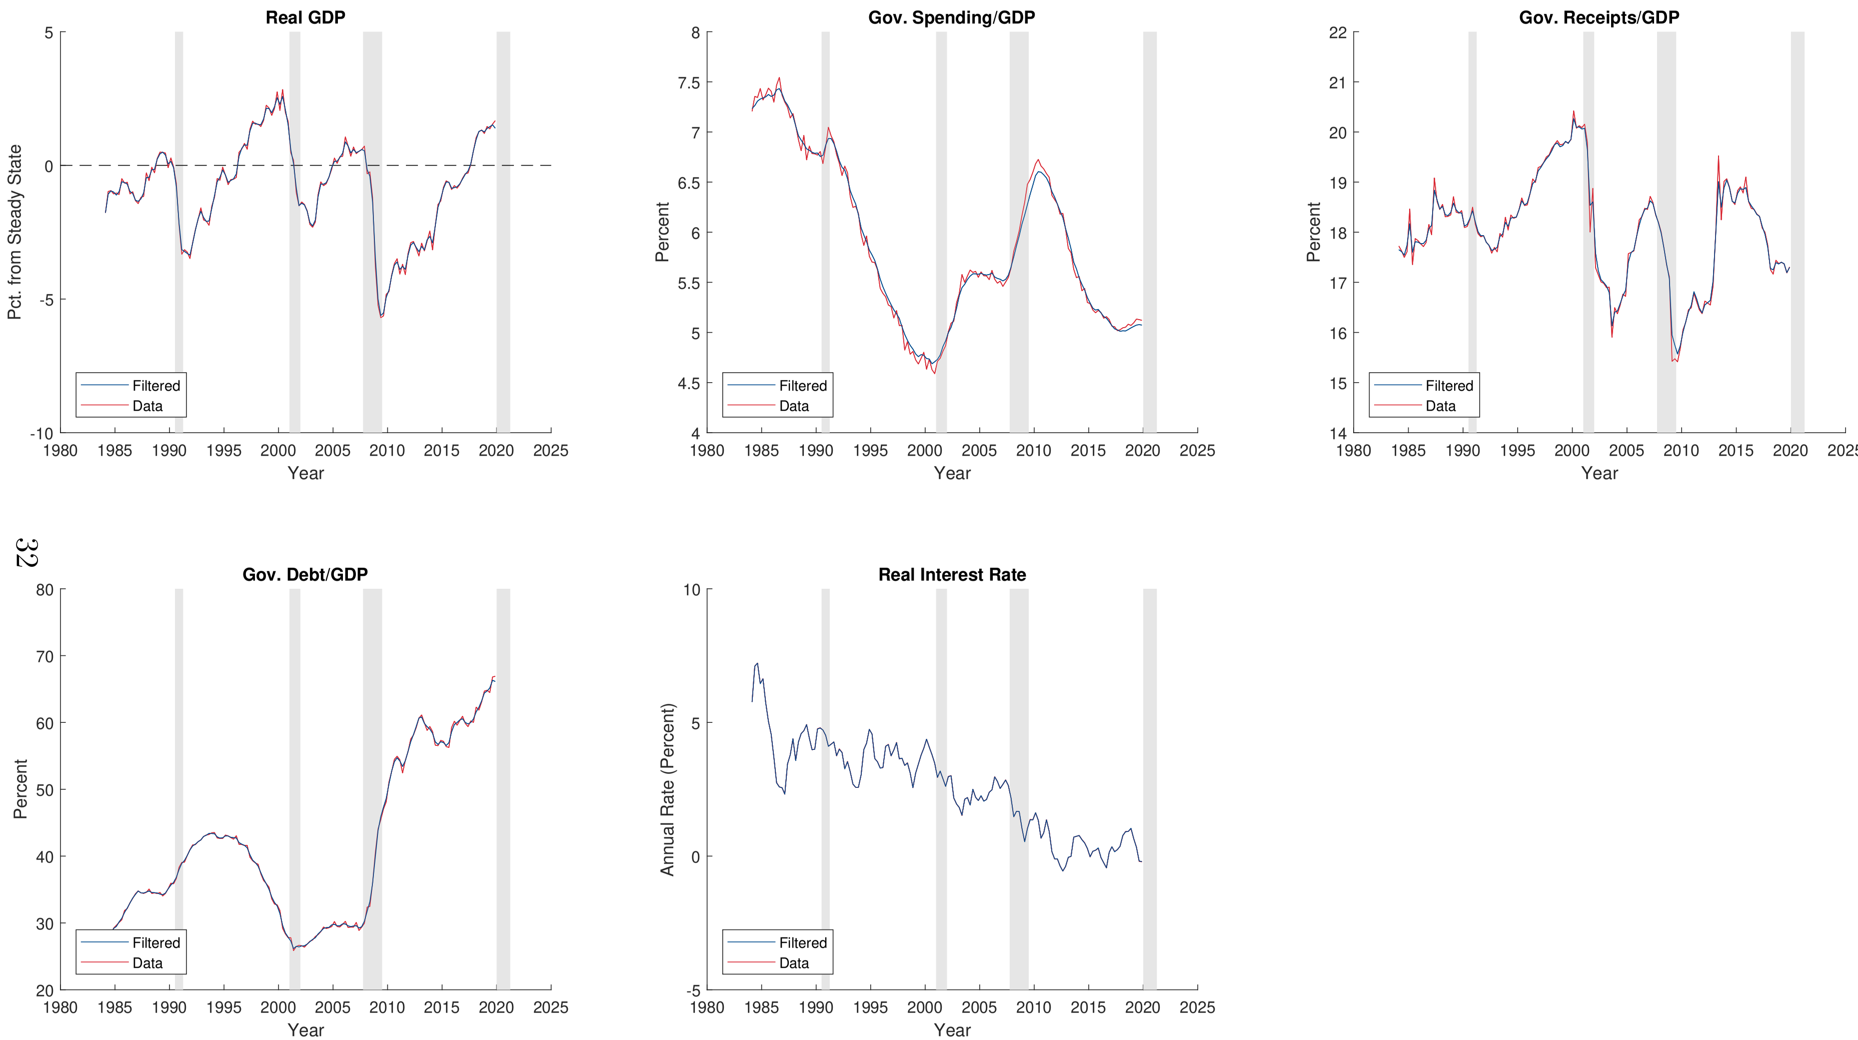

Observed data is plotted in red, while filtered series from the model are in blue. The filtered series generally show close agreement with the observed data. Shaded areas are NBER recession dates.

Figure 3: Distribution of the Probability of Default

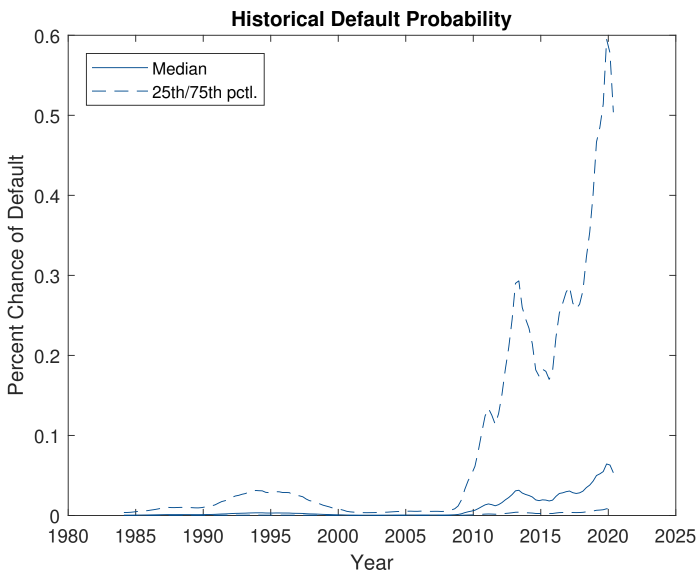

The solid line plots the median probability of a default estimated by the model over the sample. The dotted lines plot the values at the 25th and 75th percentiles. Due to the non-linear model of default probabilities, the distribution is skewed so that right tail of the distribution implies substantially higher probabilities than the left tail, which is bounded at zero.

Figure 4: Decomposition of Interest Rates

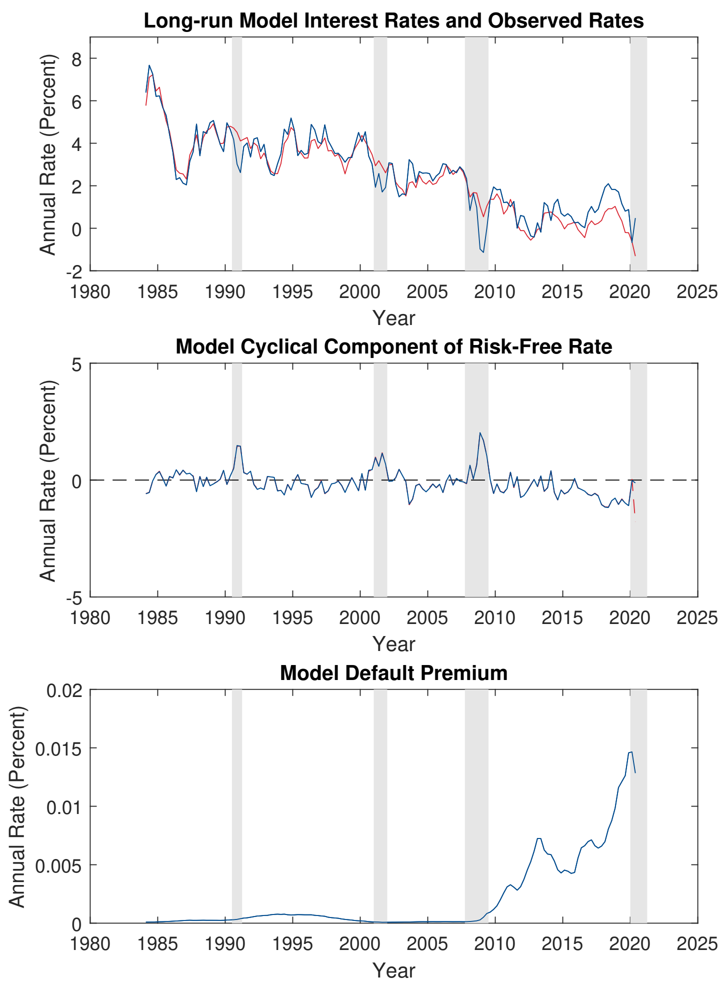

In the top panel, the red line is the observed real interest rate, and the blue line is the model's long-run interest rate. In the middle panel, the red dashed line is observed rates minus the discount factor (the portion unexplained by the trend), and the blue line is the model's cyclical adjustment (difference between the risk-free rate and the trend rate). The third panel shows the model's filtered default premium (difference between the risky debt and the risk-free rate). Shaded areas are NBER recession dates.