The Value of Government Debt

John H. Cochrane[*]

Economics Working Paper 20107

HOOVER INSTITUTION
434 GALVEZ MALL
STANFORD UNIVERSITY
STANFORD, CA 94305-6010

February 25, 2020

The market value of government debt equals the present discounted value of primary surpluses. Applying present value decompositions from asset pricing to this valuation equation, I find that half of the variation in the market value of debt to GDP ratio corresponds to varying forecasts of future primary surpluses, and half to varying discount rates. Variation in expected growth rates is unimportant.

1   Introduction

This paper measures the sources of variation in the value of government debt, expressed as a ratio to GDP. I start with the identity that the value of debt equals future primary surpluses discounted by ex-post returns. After linearizing and accounting for GDP, it follows that the variance of the debt to GDP ratio equals the sum of its covariance with future surpluses plus its covariance with future growth and less its covariance with future real returns.

In this paper's central calculation, I find that about half of the postwar variation in debt to GDP ratios corresponds to variation in expected future surpluses, and about half to variation in expected returns, i.e. discount rates. Essentially none of the variation in debt to GDP ratios corresponds to changing growth forecasts. Growth helped to resolve the WWII debt, but high growth does not resolve large debt to GDP ratios on average, nor does low debt presage low growth on average. Large debts do correspond on average to low subsequent returns. From an ex-post point of view low returns bring the value of debt back down. From an ex-ante point of view, a low discount rate rationalizes a high value of debt.

To focus on business cycle variation, I develop a variant of the variance decomposition for filtered data, and for innovations. Deficits in recessions are followed by surpluses in booms, so business-cycle frequency debt variation largely corresponds to variation in expected surpluses. Variation of innovations in the value of debt, however, is dominated by discount rate movements, which change bond prices.

Did people in 1947 know that an unprecedented era of economic growth would bring debt to GDP back in line? Or perhaps they expected more taxes, less spending, lower economic growth, and lower bond returns? We cannot know from time-series data what expectations were at each point in time, but we can learn what events follow a given debt on average, and thus, with rational expectations, what people expected on average in such events. In this paper's calculation, on average in the configuration of data as of 1947, people expect large surpluses and low returns to resolve a large debt to GDP ratio, not high growth.

What expectations of growth, surpluses, and rates of return underlie the large value of today's debt? With other state variables as well as debt to GDP at similar values to those just after WWII, this paper's calculation offers a similar analysis: A high value of debt is sustained about half by expectations of future surpluses and about half by low expected bond returns.

I also consider a backward-looking decomposition, that answers what combination of initial debt, past surpluses and deficits, inflation, bond returns, and growth produce each date's debt to GDP ratio. This calculation forms a set of stylized facts that help us to digest the variance

decomposition estimates, and neatly summarizes the lessons of existing literature.

The analysis of government finances, how debt is paid off, grown out of, or inflated away, is a long literature. Hall and Sargent (2011) and Hall and Sargent (1997) are the most important immediate precursors. Hall and Sargent emphasize accounting for the market value of debt, not the face value reported by Treasury, and consequent proper accounting for interest costs as the rate of return on the government debt portfolio. Bohn (2008) examines a long history of US debt, notes its value is stationary arguing that present values are finite, and shows that primary surpluses are higher following large debts, though growth also brings down debt to GDP ratios. A number of prominent economists, most notably Blanchard (2019), now raise the possibility that via permanent $r < g$, or other means, debts do not have to be paid by future surpluses, even at the margin if debts are expanded. The expectations underlying the large value of debt are of obvious interest to this debate.

The variance decompositions and identities are straightforward adaptations of techniques used in asset pricing to study how price-dividend ratios vary in response to news about future returns and dividend growth rates. Cochrane (2011) provides a short summary of this literature. The identities I use here are adaptations of the Campbell and Shiller (1988) identities. I linearize in the level of the surplus to GDP ratio, or surplus to value ratio, rather than the its logarithm, as surpluses are often negative. I adopt the decomposition approach in Cochrane (1992): With $x = y + z$, I explore $var(x) = cov(x, y) + cov(x, z)$ rather than $var(x) = var(y) + var(z) + 2cov(y, z)$.

2 Data

I use data on the market value of US Federal debt held by the public and the nominal rate of return of the government debt portfolio from Hall, Payne, and Sargent (2018).

I measure the debt to GDP ratio by the ratio of the market value of debt to personal consumption expenditures, times the average GDP to consumption ratio. In this time-series application a debt to GDP ratio introduces cyclical variation in GDP to the measure of the value of debt. We want only a detrending divisor, an indicator of the economy's long-run level of tax revenue and spending, a measure of debt to cyclically-adjusted GDP. Consumption is a good stochastic trend for GDP, and does not suffer the look-ahead bias of potential GDP or Hodrick-Prescott filtered GDP. Since its units are debt to GDP and the latter is a more familiar term, I refer to it as debt to GDP, and more frequently as just "debt," but it is really a debt to consumption ratio.

I infer the surplus from a linearized version of the government debt flow identity,

$$v_{t+1} = v_t + r_{t+1}^n - \pi_{t+1} - g_{t+1} - s_{t+1}. \tag{1}$$

The log debt to GDP ratio at the end of period $t + 1$, v_{t+1}, is its value at the end of period t, v_t, increased by the log nominal return r_{t+1}^n on the portfolio of all government bonds, less log inflation π_{t+1}, less log real GDP growth g_{t+1} , and less the primary surplus s_{t+1}. The quantity s_{t+1} can represent either the ratio of real primary surplus to the previous period's value of debt, or the real primary surplus to GDP ratio scaled by the average debt to GDP ratio. Both definitions lead to the same linearized identity. I refer to s_{t+1} for brevity simply as the "surplus." This linearized identity is derived in the Appendix to Cochrane (2019), which also shows that the identity is reasonably accurate for these purposes – the difference between the surplus imputed by linear and nonlinear versions of the identity is small in the postwar period relative to sampling error and other uncertainties of the calculations. The identity holds for any detrending divisor. Again, I use the debt to consumption ratio and consumption growth.

By using a surplus s_{t+1} inferred from the identity (1), we measure how much money the government actually borrows, which NIPA surpluses do not measure – they do not obey the linear or nonlinear version of (1) – and we produce time series that obey the linearized identity exactly, so that VAR variance decompositions add up exactly without adding a confusing approximation error term.

Figure 1 presents the surplus. The vertical dashed line indicates 1947 where I begin the sample for analysis. There *are* primary surpluses. One's impression of endless deficits comes from the full deficit including interest payments on the debt. After the great depression and WWII there is an era of steady primary surpluses, which, as we will see, helped to pay off WWII debt. From 1975 there is a new era of primary deficits, but also interrupted by the strong surpluses of the late 1990s. Outside of WWII, primary surpluses have a clear cyclical pattern, revealed by the correlation of declining surpluses with recession bars and the strong correlation of surpluses with the cyclical and longer-term variation in unemployment.

3   The history of debt

I start by iterating (1) backward, calculating at each date t the terms of

$$v_t = v_0 + \sum_{j=1}^{t} \left(r_j^n - \pi_j - g_j - s_j \right). \tag{2}$$
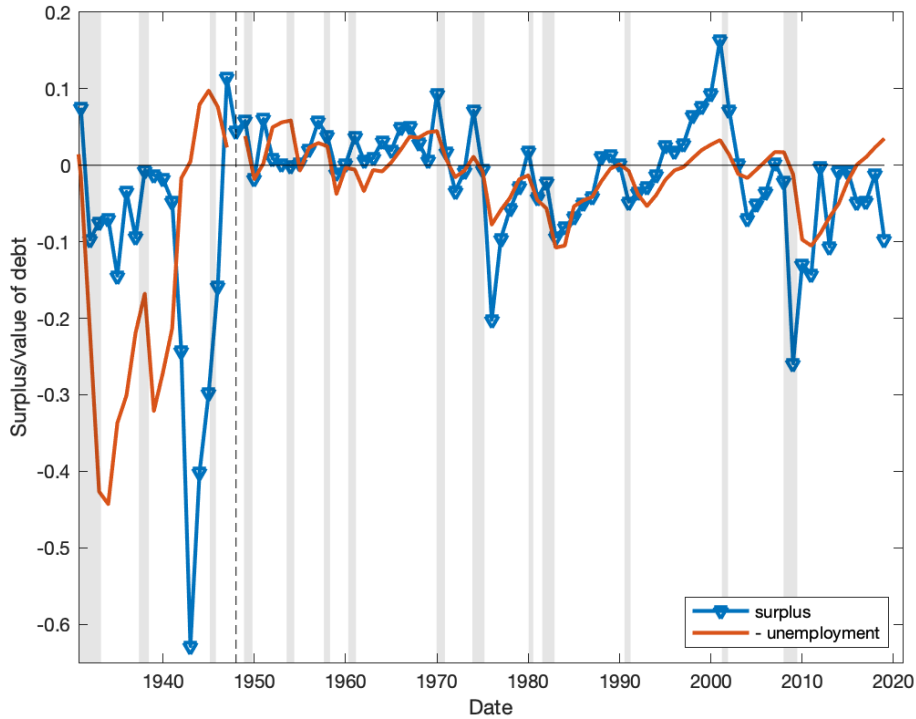
Figure 1: Surplus and negative of the unemployment rate. The unemployment rate is rescaled to fit on the same graph.

This identity tells us where the debt to GDP ratio came from at each date – what combination of prior surpluses and deficits, yes, but also prior bond returns, inflation, and growth produce the debt to GDP ratio at each date t. It provides us with a set of stylized facts that help to interpret the forward-looking variance decomposition calculations that follow.

Figure 2 presents the terms of the decomposition (2). Each line presents what debt to GDP ratio v_t would be at date t, starting from v_0, if that were the only term, setting the other terms to zero: "$-s$" plots $-\sum_{j=1}^{t} s_j + v_0$, "$r-g$" plots $\sum_{j=1}^{t} \left(r_j^n - \pi_j - g_j \right) + v_0$. The dashed lines give individual components, nominal return r^n, inflation π, real return $r = r^n - \pi$, and growth g. Though the method is somewhat different, the results are similar to Hall and Sargent (2011) Figure 7.

In the v line we see the evolution of the market value of debt to GDP ratio since 1930. There are three distinct periods, with different behavior. First, debt to GDP rises in the great depression and in WWII, to a value greater than 100% , i.e. $v_t = 0 = \log(1)$. The $-s$ line is the cumulative primary deficit. This rise in debt to GDP ratio is driven by the deficits of the 1930s and WWII. Returns and growth matter as well, however. About half of the rise in debt to GDP
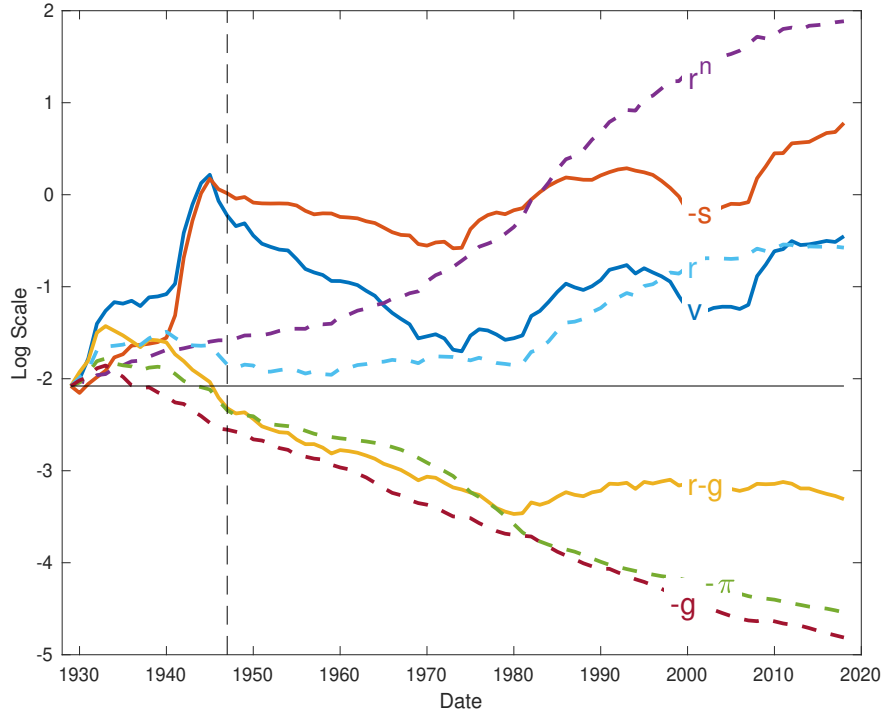
Figure 2: Backwards decomposition of the value of debt. Terms of the identity
$v_t - \bar{v} = v_0 - \sum_{j=1}^{t} s_j + \sum_{j=1}^{t} g_j + \sum_{j=1}^{t} \left(r_j^n - \pi_j \right)$.

ratio by 1940 comes from the r term, the high real bond returns coming from deflation in the great depression. Likewise, the rise in debt to GDP ratio in WWII was only about two-thirds as large as the cumulative deficits. Here low returns and the rise in GDP kept the debt to GDP ratio from rising as much as it otherwise would have done.

Second, from the end of WWII to about 1975, the debt to GDP ratio fell steadily. About a third of this fall came from cumulative primary surpluses, as shown by the fall in the $-s$ line. The rest came from returns and growth, $r - g$. Nominal returns match inflation for a nearly zero contribution of real returns. But that is an unusually low contribution – we expect a positive real return. Growth, more GDP, accounts for 2/3 of the reduction in the debt to GDP ratio.

The inflation of the 1970s is visible, in a speeding up of the fall in the $-\pi$ line. But nominal returns in the r^n line rise as well, so real returns r are pretty much constant and small until 1980. To devalue outstanding debt, inflation must come swiftly, relative to the maturity structure of the debt. The short maturity structure of debt in the 1970s means that even that large and unexpected inflation did not really do that much to lower the value of debt.

A sharp break in the behavior of surpluses to GDP ratios occurs around 1975. First, there are two waves of large deficits, with an interlude of surpluses in the 1990s. (See also Figure 1.)

The rise and fall of the value of debt largely reflects these surpluses and deficits – the s and v lines move in parallel. The cumulative effect of returns and growth $r - g$ is small. Growth continues with a barely-visible slowdown. But starting in 1980, cumulative real returns rise, and raise the value of debt. The era of $r < g$ ended, so now the net effect of returns and growth is basically flat.

Since 2000 bond returns have declined – the r line does not rise as fast – but growth also declined. Has the US has entered a long-lasting period with the average return on government bonds less than growth, $r < g$, so that the debt to GDP ratio will decline even with zero primary surpluses? The slow decline in $r - g$ starts only in 2010, and it is declining more slowly than in the 1945-1975 era, so the graph does not yet offer a lot of comfort to such a view. Whether r or g would change in the face of a large fiscal expansion is another important question on which the graph offers no evidence. Marginal $r - g$ may not be the same as average $r - g$.

4   Decompositions

Iterating (1) forward, we have a present value identity,

$$v_t = \sum_{j=1}^{T} s_{t+j} + \sum_{j=1}^{T} g_{t+j} - \sum_{j=1}^{T} \left(r_{t+j}^n - \pi_{t+j} \right) + v_{t+T}. \tag{3}$$

The value of debt, divided by GDP, is the present value of future surpluses, discounted at the ex-post real return, and adjusted by GDP growth.

The identity (3) holds ex-post. It is, essentially, the definition of ex-post return, rearranged. Taking conditional expectation, with any information set that contains v_t, or any set of probabilities, it therefore also holds ex-ante. We can write

$$v_t - \bar{v} = E_t \sum_{j=1}^{\infty} s_{t+j} + E_t \sum_{j=1}^{\infty} g_{t+j} - E_t \sum_{j=1}^{\infty} \left(r_{t+j}^n - \pi_{t+j} \right). \tag{4}$$

Here I use the fact that the debt to GDP ratio is stationary, to set the limiting debt/GDP term to its unconditional mean $\bar{v}$.

To calculate expectations, I use a first-order VAR with annual data, consisting of the nominal return on the government bond portfolio r^n, the consumption growth rate g, inflation π, surplus s, the value of debt to GDP ratio v, the three-month Treasury bill rate i and the 10 year bond yield y. The latter are important forecasting variables for growth, inflation, and long-term bond returns.

Appendix Table 2 presents the OLS estimates of the VAR coefficients. Some important

points: The interest rates help to forecast the government bond portfolio return. Debt is not the only important state variable for the conditional expectations in (4). Debt predicts higher surpluses. This coefficient drives the finding that debts are partially resolved by surpluses. Debt does not predict growth, nor do surpluses which are forecast by debt. This fact drives the result that growth forecasts do not help to explain debt variation. The debt is persistent, with an 0.98 coefficient on its lag, and thereby constitutes the dominant root for long-term forecasts. Higher surpluses and growth lower the debt. Interest rates forecast the debt, largely mirroring their forecast of returns on the debt.

I focus on the sample since 1947, since WWII is obviously a different phenomenon than the subsequent cyclical variation, its deficits are so huge they dominate any estimates, and the linear approximation leaves a substantial error in the huge deficit of 1944. The Appendix includes full-sample results which in the end are not much different.
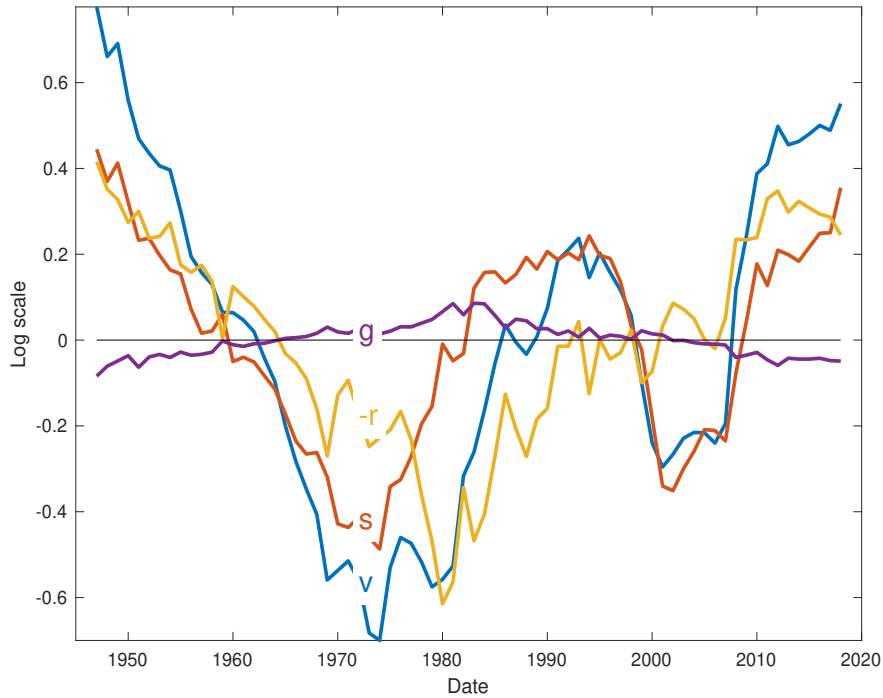
Figure 3: Decomposition of the value of debt to GDP ratio v_t. Each line presents the VAR estimate of the components of the value of debt at each date, i.e. $E_t \sum_{j=1}^{\infty} s_{t+j}$.

Figure 3 presents VAR estimates of the decomposition of the value of debt at each date calculated from (4), e.g. $E_t \sum_{j=1}^{\infty} s_{t+j}$, labeled "$s$" and so forth. The value v line is equal to the sum of the surplus s, discount rate $-r$, and growth g lines. The figure shows visually that variation in value v is composed about equally of variation in expected surpluses s and discount rates $-r$ with little contribution from expected growth g. There are three great waves of expected surplus,

while discount rates trend down and then back up again.

For example, in this analysis, the value of debt to GDP ratio was high at the end of WWII about equally because surpluses were expected to be large, and because discount rates – real returns on government bonds – were expected to be low. Though ex-post growth resolved as much as 2/3 of the WWII debt to GDP ratio, the VAR does not find that high debts are followed by large growth on average, so does not assign much force to expected growth in 1947. The subsequent episodes of high and low debt to GDP ratio did not resolve by high and low growth. For example, the low debt to GDP 1970s followed by good, not bad, growth in the 1980s. Debt does not forecast growth directly or indirectly in the VAR. The postwar growth that helped to pay off the WWII debt was, by this measure, unexpected, which is reasonable given pervasive fear at the time of a return to depression, and that the world had never before seen three decades of growth such as occurred following WWII.

The value of debt bottomed out in the late 1970s by a reversal of both surpluses and returns. The swings in value of debt from 1980 to today reflect large swings in surpluses, as debt and surplus lines move in parallel, with a trend to lower real returns contributing an upward trend to the value of debt.

Not every data point is half surplus and half return. For example, low debt in 1980 is largely attributed to high expected bond returns, and the higher debt in 1990 is largely supported by expected future surpluses. Both expectations proved correct ex-post.

The 2018 value of the debt is, by this analysis, again supported half by higher than normal expected surpluses, and half by lower than normal expected bond returns. The VAR has not read, or does not trust, the CBO's apocalyptic long-term fiscal policy analysis. (To be fair, the CBO's forecast is not a conditional mean, but a warning of what will happen if the US does not raise surpluses.) The VAR also does not signal an extended period of $r < g$ that will painlessly pay off debts without a return to primary surpluses. Of course, the VAR does not, by definition, see structural shifts. It only sees what today's state variables imply for the future given historic correlations. The VAR does include today's bond yields along with all the other state variables in making this forecast.

Agents have more information than we do, so we must be clear about what these VAR expectations mean. Since it stems from an ex-post identity, (4) holds with the information set of a VAR, even when agents know more than we do. We can condition down to the VAR information set and obtain a valid equation.

However, the resulting estimates only measure people's expectations on average, conditioned down to the VAR variables, not their expectations in each historical episode. We do not

know that the 1947 value of debt was supported by the plotted VAR expectations of surpluses, growth and discount rates. We do measure that the constellation of debt, interest rates and other VAR state variables that held in 1947, corresponded *on average* to the plotted ex-post realizations of surpluses growth and returns, and thus, if to the average of fully-informed rational expectations.

The calculations are based on an identity, so there is nothing to test. Which component of an identity matters is an interesting measurement, however, and that is the point of this paper.

It is a tempting but classic mistake to calculate at each date the right hand side of (4) and then compare the result to v_t as a test of the present value relation. The present value relation is an identity, so if one correctly includes v_t in the VAR and data obey identities, the answer comes out to v_t exactly, as it does here. Leaving v_t out of the VAR to attempt a test, though common, is an error. Doing so assumes that agents have no more information than the VAR and that they use the estimated VAR coefficients to forecast. Leaving v_t out of the VAR also violates the conditioning-down assumption that leaves v_t and not $E(v_t|x_t)$, a lesser information set, on the left-hand side. One can test models of discount rates that differ from the ex-post return, but one cannot test identities. (For a recent example, see Jiang et al. (2019). Hansen, Roberds, and Sargent (1992) analyze the issue.)

4.1   Variance decomposition

Here, I calculate a decomposition of the variance of the debt. In part, this calculation can summarize the lessons of Figure 3. Multiplying (3) by $v_t - E(v_t)$ and taking expectations, we obtain

$$var\left(v_t\right) = \sum_{j=1}^{\infty} cov(v_t, s_{t+j}) + \sum_{j=1}^{\infty} cov(v_t, g_{t+j}) - \sum_{j=1}^{\infty} cov\left[v_t, \left(r_{t+j}^n - \pi_{t+j}\right)\right]. \tag{5}$$

Dividing by $var(v_t)$, we obtain the decomposition. The terms add up to one and represent fractions of variance accounted for by each component. The components may be negative. The Appendix gives the formulas for calculating each variance and covariance from the VAR.

These calculations answer a simple question – *on average*, in the past, when the debt to GDP ratio has been high, was that event followed by a long string of positive surpluses? Or was it followed by a long string of high growth rates or low real returns? If the latter, were the low real returns the result of low nominal returns or high inflation?

Such components are also the coefficients of single regressions by which v_t forecasts the other terms. Though the variance decomposition is equivalent to univariate long-run regressions, given our limited sample size and the slow variation of the debt to GDP ratio v_t, calculating

the implied long-run regression coefficients from the VAR improves the precision of estimates, at the cost of imposing the VAR structure.

Component	Plain	Innovation	Filtered
		$(\cdot)/var(\tilde{v}_t)$	
$var(\tilde{v}_t)$	1.00	1.00	1.00
$cov(\tilde{v}_t, \Sigma\tilde{s}_{t+j})$	0.55	0.28	0.85
$-cov[\tilde{v}_t, \Sigma(\tilde{r}^n_{t+j} - \tilde{\pi}_{t+j})]$	0.52	0.85	0.09
$cov[\tilde{v}_t, \Sigma\tilde{g}_{t+j}]$	-0.07	-0.13	0.07
$-cov(\tilde{v}_t, \Sigma\tilde{r}^n_{t+j})$	0.90	1.63	0.14
$cov(\tilde{v}_t, \Sigma\tilde{\pi}_{t+j})$	-0.38	-0.78	-0.05
$\sigma(\tilde{v}_t)$	0.37	0.07	0.14

Component	Plain		Innovation		Filtered	
			$(\cdot)/var(\tilde{v}_t)$			
	25%	75%	25%	75%	25%	75%
$cov(\tilde{v}_t, \Sigma\tilde{s}_{t+j})$	0.51	0.75	0.23	0.54	0.71	0.81
$cov[\tilde{v}_t, \Sigma(\tilde{r}^n - \tilde{\pi})_{t+j}]$	-0.56	-0.24	-0.86	-0.46	-0.21	-0.09
$cov(\tilde{v}_t, \Sigma\tilde{g}_{t+j})$	-0.08	0.03	-0.13	0.03	0.07	0.11
$-cov(\tilde{v}_t, \Sigma\tilde{r}^n_{t+j})$	0.91	0.48	1.54	0.91	0.27	0.16
$cov(\tilde{v}_t, \Sigma\tilde{\pi}_{t+j})$	-0.39	-0.19	-0.75	-0.38	-0.10	-0.03
$\sigma(\tilde{v}_t)$	0.29	0.40	0.06	0.07	0.12	0.15

Table 1: Decomposition of the variance of the value of debt. Sample 1947-2018. Bottom panel: Monte Carlo quantiles.

The first column of Table 1 presents the decomposition of the variance of the value of debt to GDP ratio v_t. Looking down the first column,

- *Half (55%) of the variance of the debt* $var(v_t)$ *corresponds to variation in forecasts of future surpluses. Half (52%) comes from discount rate variation. Essentially none (-7%) comes from variation in forecasted growth.*

Expected real returns come from a 90% variation in nominal returns tempered by 38% variation in expected inflation.

That growth g is not important here illustrates the difference between this variance accounting, the average experience following a large debt, and the previous ex-post history from 1947-1975, that focuses on one historical event.

The low returns following high debt of WWII, which helped to bring down that debt, and the high returns following the low debt of 1980, which helped to raise the value of the debt, both

count in the VAR to say that on average, resolution of large debts – mean-reversion to the historic debt/GDP ratio – also comes from low returns.

Though discount rate variation (looking forward) or rate of return variation (looking backward) is not prominent in the literature that studies the value of the debt, the opposite is true in asset pricing, where typically all variation in market price-dividend ratios corresponds to variation in discount rates and none to expected dividend growth. An asset pricer wonders why even more debt variation is not due to expected returns, and less due to expected surpluses and growth. One important difference helps to account for that fact: Here I account for variation in the *total market value* of debt, including issuances and redemptions, i.e. surpluses and deficits. The standard asset pricing calculation accounts for the variation of the value of a particular security. Variation in individual bond nominal values is 100% due to variation in nominal discount rates and none to variation in nominal cash flows, since the US has not defaulted in this sample, and real bond values are also dominated by real expected returns. But the value of debt also fluctuates due to bond issuances in times of deficit and bond redemptions in times of surplus. So there is no necessary puzzle reconciling these results with the usual asset pricing results. In fact, it is a bit surprising that surpluses do not account for more variation in debt, given the waves of surpluses and deficits that we have seen.

The bottom panel of Table 1 presents 25% and 75% Monte Carlo quantiles of the terms of the variance decomposition, to assess the accuracy of these measurements. The signs and broad magnitudes are reasonably measured. For example, the surplus contribution 0.55 in the top left corner has quantiles 0.51 to 0.75. But measuring long-run regressions is always imprecise. In this parametric implementation, it is sensitive to the largest root in the transition matrix A, via expressions $(I - A)^{-1}$. The sampling distributions are also non-normal, in part due to the non-normality and bias of OLS estimates of very large (0.98) autoregressive roots. We see these facts by point estimates that are not at the midpoints of the quartiles.

4.2   Filtered data

The debt, cumulative surpluses, and cumulative returns all vary very slowly. This fact means that our estimates are really based on only a few data points. Also, we may suspect that cyclical variation in debt has different foundations from the multi-decade swings that dominate the variance of the debt. For these reasons, I investigate the variance of the filtered debt.

Since the present value identity (3) is an identity, we can filter each side, applying $\theta(L)$ or $E_t - E_{t-1}$ before taking expectations, to obtain a present value relation (4) and variance decom-

position (5) that apply to filtered data. The same formulas thus apply to

$$\tilde{v}_t = \theta(L)v_t = v_t - \frac{1}{3}\left(v_{t-1} + v_{t-2} + v_{t-3}\right),$$

and to the VAR innovation in the value of debt,

$$\tilde{v}_t = v_t - E_{t-1}(v_t).$$

Starting with an ex-post identity is important to this transformation. If one starts with $p_t = E_t(m_{t+1}x_{t+1})$ one obtains $\theta(L)p_t = \theta(L)E_t[(m_{t+1}x_{t+1})]$, not $\theta(L)p_t = E_t[\theta(L)(m_{t+1}x_{t+1})]$. The Appendix includes discussion and an example.

   The last row of Table 1 gives the standard deviation of each measure of the value of debt. "Plain," the standard deviation of the log market value of debt to GDP v_t is 0.37. The debt to GDP ratio varies by 37 percentage points, a lot. The standard deviation of the debt innovation, 7% and filtered debt, 14 % are lower, as they eliminate the large low-frequency variation.

   The second and third columns of Table 1 present the variance decomposition applied to filtered data and innovations. In the "Filtered" column of Table 1, we see important differences:

- *Variation in the filtered value of debt corresponds almost entirely, 85%, to future surpluses, with only 9% to discount rates and 7% to growth.*

   Recession-related, and somewhat longer (Vietnam, 1980s deficits, 1990s surpluses) variation in the value of debt is driven by variation in expected surpluses. The cyclical pattern of surpluses seen in Figure 1 means that cyclical deficits are followed, on average, by cyclical surpluses. The calculation verifies that the surpluses are enough, on average, to repay the recession's debts. It's initially surprising that growth is less important at cyclical frequencies, but remember that I detrend debt with consumption not GDP. Debt to GDP ratios would rise more in recessions, as GDP falls relative to consumption, and fall more in surpluses, as GDP rises, and variation in GDP growth would account for that greater variation. Again, though, I am interested in the source of variation of the debt, not of GDP, so I intentionally leave out this source of cyclical time-series variation in true debt to GDP ratios.

   When we look at higher frequencies still, by considering VAR innovations in the variables, the situation turns around:

- *Annual variation in the unexpected value of the debt to GDP ratio corresponds less, 28%, to revisions in expected future surpluses, and more, 85%, to revisions in expected future discount rates. Growth contributes little, and in the wrong direction (-13%).*

The innovation in the value of debt naturally correlates highly with *current* bond returns, as a positive bond return raises bond prices and hence the value of debt. That is not, directly, the mechanism here, but it is an important indirect mechanism. The present value identity (4) and variance decomposition (5) correlate the innovation in the value of debt $(E_t - E_{t-1})\, v_t$ with *subsequent* surpluses and discount rates only, starting at $(E_t - E_{t-1})\, r^n_{t+1}$. But a persistent rise in expected bond returns results in a lower bond price today and a lower current return.

The reason, then, that innovations in the value of government debt correlate so strongly with expected returns at high frequency is simply that variation in the nominal price of government debt is more important at high frequency than variation in the quantity of debt, induced by surprise surpluses and deficits, or variation in the real value of debt due to surprise inflation. At an overnight frequency, variation in expected returns, and hence current bond returns, likely accounts for all variation in unexpected value of the debt.

The value of debt thus inherits the negative autocorrelation inherent to the mark-to-market value of any long-term bond portfolio. Marking to market is always pleasant in the abstract, but many economists and policy makers distrust it in practice, as temporary price variation means that higher values correspond to lower expected returns, and that component of value will melt away. But predictable variation in bond prices is no more ipso-facto irrational or non-fundamental than predictable variation in strawberry prices. As it is for banks and stocks, so it is for government debt. Whatever judgement one renders, predictable variation in the value of debt is a feature of mark-to-market accounting.

Appendix Figure 5 graphs filtered debt and the debt innovation. Both measures pick up familiar cyclical movements, abstracting from long-term trends including WWII and the post 2000 rise in debt. Appendix Table 3 presents the OLS estimates of the filtered VAR.

Appendix Figure 6 presents the time series decomposition of the filtered value of debt analogous to Figure 3, providing additional intuition for its variance decompositions. In 2000 and 2010 we see the strong correlation between the value of the debt and expected surpluses. 1980 is a different story, in which the sharp decline in value of debt lines up naturally with the sharp rise in expected returns of that era.

5   Concluding comments

This analysis evidently just scratches the surface. Its contribution is as much methodological, to show that asset pricing value decompositions can be productively applied to government debt questions, as it is empirical. Debt data go back centuries, and across many countries, allowing

much more comprehensive experience of how debt is resolved, but requiring us to think what is the same and different across different periods of history, different sources of fiscal shock (recessions, wars, revolutions, etc.) and different monetary/fiscal regimes. A parallel investigation of the foreign exchange value of the debt beckons. More variables can and should be included in the VAR or other technique for forecasting, including official forecasts and survey expectations, facing the usual limits of overfitting. State and local debts, commitments such as pensions and medical care, the overall size of government relative to Laffer limits and the structure and efficiency of the tax system are all potentially important forecasters of how debt will be resolved, along with better forecasts of returns and growth rates. VARs raised to powers can impose structure on long-term forecasts that is not there in direct forecasts of long-term variables, but the latter require longer time series. The size of the huge and ongoing variance decomposition literature in finance gives some sense of the many avenues left open for exploration.

References

Blanchard, Olivier. 2019. "Public debt and low interest rates." *American Economic Review* 109:1197–1229.

Bohn, Henning. 2008. "The Sustainability of Fiscal Policy in the United States." In *Sustainability of Public Debt*, edited by Reinhard Neck and Jan-Egbert Sturm. Cambridge, MA: MIT Press, 15–49.

Campbell, John Y. and Robert J. Shiller. 1988. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *The Review of Financial Studies* 1 (3):195–228. URL +http://dx.doi.org/10.1093/rfs/1.3.195.

Cochrane, John H. 1992. "Explaining the Variance of Price-Dividend Ratios." *Review of Financial Studies* 5:243–280.

———. 2011. "Discount Rates: American Finance Association Presidential Address." *Journal of Finance* 66:1047–1108.

———. 2019. "The Fiscal Roots of Inflation." *Manuscript* URL https://faculty.chicagobooth.edu/john.cochrane/research/papers/decomposition.pdf.

Hall, George J., Jonathan Payne, and Thomas Sargent. 2018. "US Federal Debt 1776 -1960: Quantities and Prices." *Manuscript* URL http://people.brandeis.edu/~ghall/.

Hall, George J. and Thomas J. Sargent. 1997. "Accounting for the Federal Government's Cost of Funds." *Federal Reserve Bank of Chicago Economic Perspectives* :18–28.

———. 2011. "Interest Rate Risk and Other Determinants of Post WWII U.S. Government Debt/GDP Dynamics." *American Economic Journal: Macroeconomics* 3:192–214.

Hansen, Lars Peter, William Roberds, and Thomas J. Sargent. 1992. "Implications of present value budget balance and of martingale models of consumption and taxes." In *Rational Expectations Econometrics*, edited by Lars Peter Hansen and Thomas J. Sargent. London and Boulder, Co.: Westview Press, 121–161. URL http://www.tomsargent.com/books.html.

Jiang, Zhengyang, Hanno N. Lustig, Stijn Van Nieuwerburgh, and Mindy Z. Xiaolan. 2019. "The U.S. Public Debt Valuation Puzzle." *Manuscript* URL https://ssrn.com/abstract=3333517.

Shiller, Robert J. 1981. "Do stock prices move too much to be justified by subsequent changes in dividends?" *American Economic Review* 71 (3):421–436. URL http://www.jstor.org/stable/1802789.

Online Appendix to "The Value of Government Debt"

1 VAR

Table 2 presents OLS estimates of the VAR coefficients. This is the same VAR as in Cochrane (2019). I present it here to keep the analysis self-contained. See the text for a description of these coefficients.

Each column is a separate regression. I do not orthogonalize shocks, so the order of variables is not important. I compute standard errors from a Monte Carlo. The stars in Table 2 represent one or two standard errors above zero, and are included just to save a table of standard errors, not as tests of economic or statistical significance.

	r^n_{t+1}	g_{t+1}	π_{t+1}	s_{t+1}	v_{t+1}	i_{t+1}	y_{t+1}
r^n_t	-0.17**	-0.02	-0.10**	-0.32*	0.28*	-0.08*	0.04*
g_t	-0.27*	0.20*	0.16*	1.37**	-2.00**	0.28**	0.06
π_t	-0.15	-0.14*	0.53**	-0.25	-0.29	0.09	0.04
s_t	0.12**	0.03	-0.03*	0.35**	-0.24*	-0.04*	-0.04**
v_t	0.01	-0.00	-0.02**	0.04*	0.98**	-0.01	-0.00
i_t	-0.32*	-0.40*	0.29*	0.50	-0.72	0.73**	0.36**
y_t	1.93**	0.54**	-0.17	-0.04	1.60*	0.11	0.46**
$100 \times std(\varepsilon_{t+1})$	2.18	1.53	1.12	4.75	6.55	1.27	0.82
Corr $\varepsilon, \varepsilon_\pi$	-0.29	-0.24	1.00	-0.14	-0.11	0.21	0.31
R^2	0.71*	0.17*	0.73*	0.48*	0.97*	0.82*	0.90*
$100 \times std(x)$	4.08	1.68	2.16	6.61	37.00	2.96	2.63

Table 2: VAR estimate. Sample 1947-2018. One (two) stars means the estimate is one (two) Monte Carlo standard errors away from zero.

2 VAR calculations

Denote the VAR

$$x_{t+1} = Ax_t + \varepsilon_t,$$

and let a_i denote a selection vector which picks out element i, e.g.

$$v_t = a'_v x_t.$$

To estimate the decomposition in Figure 3, write the present value identity (4) as

$$v_t = \left(a'_s + a'_g - a'_r\right)(I - A)^{-1}Ax_t, \tag{6}$$

where $a_r \equiv a_{r^n} - a_\pi$. The variance decomposition (5) is then

$$a'_v V a_v = \left(a'_s + a'_g - a'_r\right)(I - A)^{-1}AV a_v \tag{7}$$

where $V \equiv cov(x_t, x'_t)$ is the covariance matrix.

To estimate the variance decomposition for the innovation case, $\tilde{v}_t = (E_t - E_{t-1})\,v_t$, write the time-$t$ innovation of the present value identity (4) as

$$\Delta E_t v_t = \Delta E_t \sum_{j=1}^{\infty} s_{t+j} + \Delta E_t \sum_{j=1}^{\infty} g_{t+j} - \Delta E_t \sum_{j=1}^{\infty} \left(r^n_{t+j} - \pi_{t+j}\right) \tag{8}$$

where $\Delta E_t \equiv E_t - E_{t-1}$. In VAR notation

$$a'_v \varepsilon_t = \left(a'_s + a'_g - a'_r\right)(I - A)^{-1}A\varepsilon_t$$

so the variance decomposition of debt innovations, with $\tilde{v}_t = (E_t - E_{t-1})\,v_t$, is

$$a'_v \Sigma a_v = \left(a'_s + a'_g - a'_r\right)(I - A)^{-1}A\Sigma a_v$$

where $\Sigma \equiv cov(\varepsilon_t, \varepsilon'_t)$ is the covariance matrix of the VAR shocks. This formula is the same as (7), with the innovation covariance matrix Σ in the place of the covariance matrix V. It will therefore de-emphasize a persistent variable such as v_t whose variance is much larger than its shock variance.

To estimate the filtered variance decomposition $\tilde{v}_t = \theta(L)v_t$, we can most simply rerun the VAR with filtered data,

$$\theta(L)x_{t+1} = \tilde{A}\theta(L)x_t + \tilde{\varepsilon}_{t+1}$$

$$\tilde{x}_{t+1} = \tilde{A}\tilde{x}_t + \tilde{\varepsilon}_{t+1}$$

and use the same formulas (6) and (7) with tildes.

Filtering valuation equations is delicate. If we start with a valuation equation

$$p_t = E_t\left(m_{t+1}x_{t+1}\right),$$

we can write

$$(E_t - E_{t-1}) \, p_t = (E_t - E_{t-1}) \, (m_{t+1} x_{t+1}) \, ,$$

and

$$\theta(L) p_t = \theta(L) E_t \, (m_{t+1} x_{t+1}) \, ,$$

but we cannot write

$$\theta(L) p_t = E_t \, [\theta(L) \, (m_{t+1} x_{t+1})] \, .$$

The lag operator must also apply to the expectation on the outside. For example, if y_t is i.i.d. with $E_t(y_{t+1}) = 0$, then applying $(1 - L)$ correctly outside the expectation yields $E_t y_{t+1} - E_{t-1} y_t = 0$. Instead, applying $1 - L$ inside the expectation yields $E_t(y_{t+1} - y_t) = -y_t \neq 0$. But since we start here with an ex-post identity, we can first filter and then take expectations, so we do not make this mistake.

Even when one starts with the ex-post present value identity, filtering is less useful for asset pricing applications, as the hypothesis that filtered returns are unpredictable is not particularly interesting.

3   A forward-looking value decomposition

Figure 4 presents the terms of the forward-looking value decomposition (3),

$$v_t = \sum_{j=1}^{T} s_{t+j} + \sum_{j=1}^{T} g_{t+j} - \sum_{j=1}^{T} \left(r_{t+j}^n - \pi_{t+j} \right) + v_{t+T}. \tag{9}$$

This decomposition asks, "Suppose people knew the future, and discounted the actual ex-post surpluses using actual ex-post real returns. How would those clairvoyant expectations account for the value of the debt at each date?" Shiller (1981) pioneered the calculation of such "ex–post rational" valuations.

After 1975-1980, variation in the value of the debt to GDP ratio largely follows the variation in the sum of future surpluses, with little contribution from discount rates r or growth g. Why, for example, is the value of debt to GDP ratio low in 1980? A big part of that value is foreknowledge that the value of the debt to GDP ratio will be high in 2018, v_{t+T} – there will not be a default, an inflation, a growth disaster, or another big hit to r or g. Beyond that, the value of debt is low because on net, there will be a string of primary deficits in the 1980s and after 2008, so the s line is negative throughout.

From WWII to 1975, we see a different picture. We tell a version of the previous story
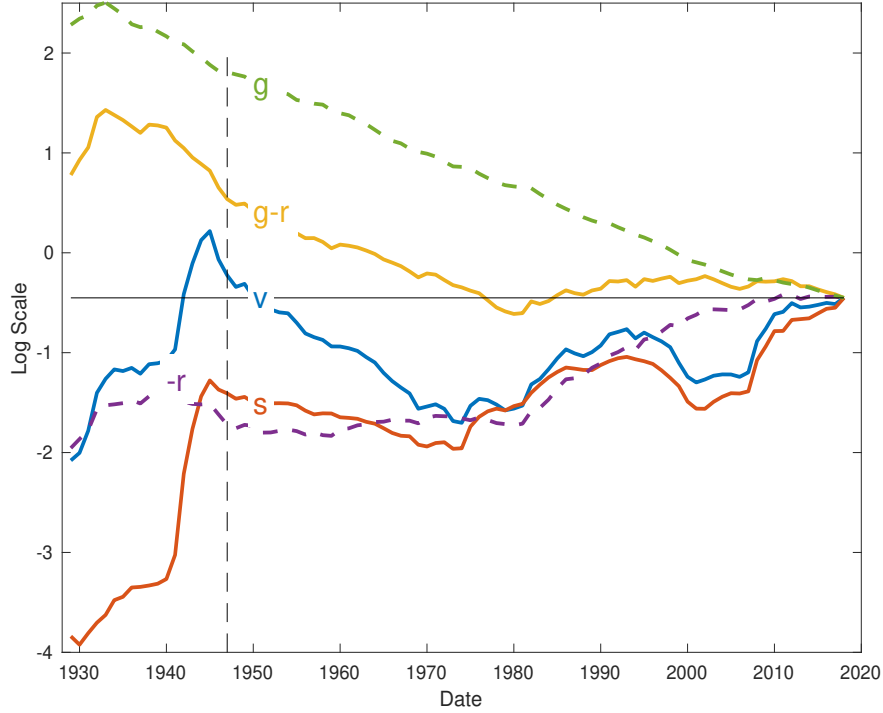
Figure 4: Forward decomposition of the market value of debt to GDP ratio.
$v_t = \sum_{j=1}^{T} s_{t+j} + \sum_{j=1}^{T} g_{t+j} - \sum_{j=1}^{T} \left(r_{t+j}^n - \pi_{t+j} \right) + v_{t+T}.$

backwards. Why was the value of debt to GDP ratio so high at the end of WWII relative to 1975? About 2/3 of it was high growth, with about 1/3 expectation of primary surplus to GDP ratios to come.

It is likely, however, that these data contain events people did not expect. One may reasonably speculate that in 1947 people expected lower growth, positive real interest rates, and therefore a longer period of larger primary surpluses than in fact occurred. Similarly, that low ex-post real returns and high ex-post growth helped to bring down the debt to GDP ratio once in the past does not imply that people will hold large debts today expecting to repeat that benign outcome. To answer the central question, what is expected, we use the central calculation using *expected* present values.

Figure 4 gives additional intuition for the variance decompositions from Table 1. You can see directly how the value of debt v_t at each date corresponds to the combination of the sum of subsequent surpluses s and the cumulative discount rate $r - g$. The variation in surpluses s includes both the surpluses until 1970 and the trend to deficits since then, plus the decade-long and cyclical variation since the 1970s that correlates so well with the value of debt. The discount rate variation is mostly longer-term, a switch from negative growth adjusted discount

rates in the first half of the sample to a very slightly positive discount rate in the second half. But variation is variation.
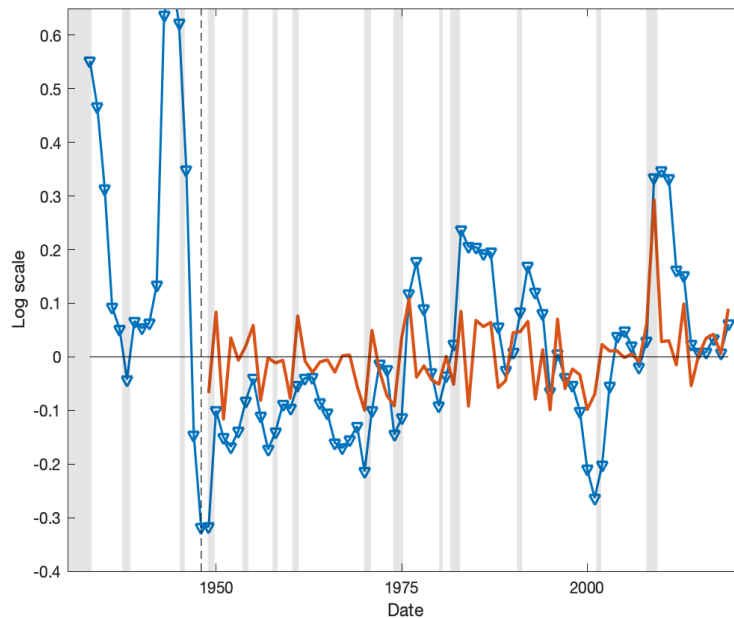
4 Filtered debt

Figure 5: Filtered value of debt $v_t^f = v_t - \frac{1}{3}(v_{t-1} + v_{t-2} + v_{t-3})$ (symbols) and the VAR innovation in the value of debt, $v_t - E_{t-1}(v_t)$. The vertical dashed line denotes 1947, the beginning of the sample used for VAR and statistical analysis. Shaded areas are NBER recessions.

Figure 5 presents filtered debt and the debt innovation. Both measures pick up familiar cyclical movements, and de-emphasize the large long-term variation visible in the previous figure. For example, you can see the big increase in debt to GDP ratio following the 2008 financial crisis, the decrease of the 1990s, the buildup in the early 1980s, and variation in debt through the recessions of the 1970s. Despite the very large debt to GDP ratio at the end of the sample, filtered debt is small because the debt to GDP ratio did not increase at an unusually large pace.

Table 3 presents the OLS estimates of the filtered VAR. The coefficients are similar to the unfiltered VAR though reduced in magnitude. The lower diagonals reflect quicker dynamics as we expect of filtered variables.

Figure 6 presents the components of the variance decomposition of filtered data at each point in time, providing intuition for their variance decompositions. Befitting filtered data, the

	r^n_{t+1}	g_{t+1}	π_{t+1}	s_{t+1}	v_{t+1}	i_{t+1}	y_{t+1}
r^n_t	-0.01	-0.05	-0.10**	-0.44**	0.58**	-0.18**	-0.02
g_t	0.06	0.07	-0.00	1.20**	-1.21**	0.12*	-0.09*
π_t	0.09	-0.13	0.10	0.04	0.07	-0.06	-0.11*
s_t	0.13**	0.01	0.02	0.11	-0.01	-0.03	-0.05**
v_t	0.01	0.02*	-0.03*	0.14**	0.88**	0.00	-0.01
i_t	-0.32	-0.51**	0.39**	0.88*	-1.08*	0.76**	0.41**
y_t	2.99**	0.48*	-0.34*	-0.75	3.60**	-0.50**	-0.04
$100 \times std(\varepsilon_{t+1})$	2.18	1.53	1.12	4.75	6.55	1.27	0.82
Corr $\varepsilon, \varepsilon_\pi$	-0.29	-0.24	1.00	-0.14	-0.11	0.21	0.31
R^2	0.62*	0.27*	0.46*	0.40*	0.77*	0.47*	0.37*
$100 \times std(x)$	3.52	1.80	1.53	6.13	13.78	1.75	1.04

Table 3: Filtered VAR estimate. Variables are the difference to a three year lagged moving average. Sample 1947-2018. One (two) stars means the estimate is one (two) Monte Carlo standard errors away from zero.
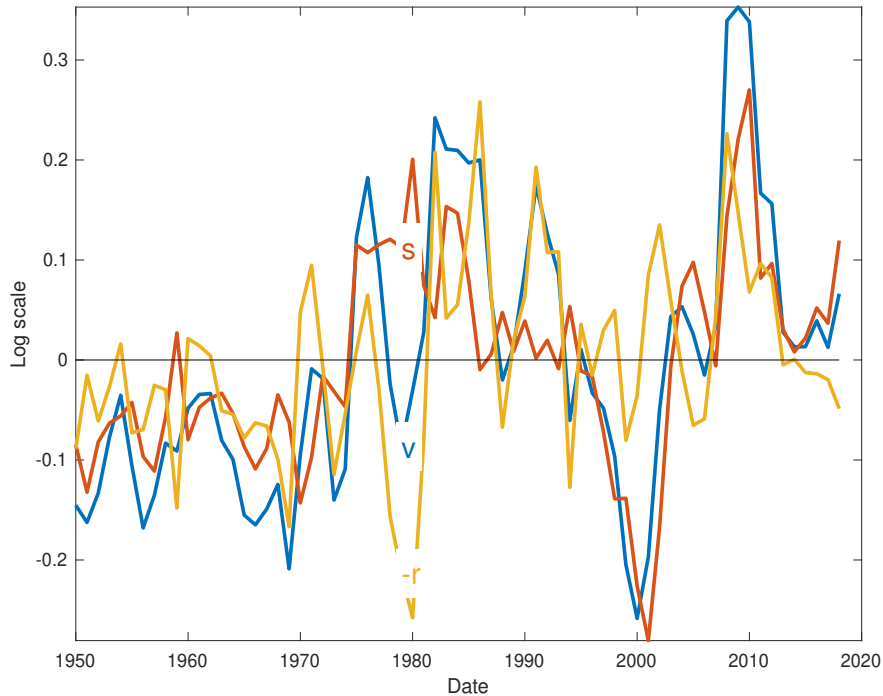

Figure 6: Decomposition of the filtered value of debt to GDP ratio. Each line presents the VAR estimate of the components of the filtered value of debt at each date, i.e. $E_t \sum_{j=1}^{\infty} \tilde{s}_{t+j}$.

variation in value is much higher frequency. The end of the sample seems like a normal level of the value of filtered debt, because debt was rising or declining much more quickly than usual in the previous three years. In 2000 and 2010 we see the strong correlation between the value of the debt and expected surpluses. 1980 is a different story, in which the sharp decline in value of debt lines up naturally with the higher expected returns of that era.

By filtering, we emphasize the medium and higher frequency fluctuation in surpluses, visible in Figure 4, and their relationship to the same fluctuations in filtered debt seen in Figure 5. The lower frequency variation in cumulative discount rates from Figure 4 is filtered out.

5 Full sample

Here I present computations using the full sample 1930-2018. Despite the enormous deficits of WWII, and their different cyclical pattern, the broad results are similar.

	r^n_{t+1}	g_{t+1}	π_{t+1}	s_{t+1}	v_{t+1}	i_{t+1}	y_{t+1}
r^n_t	-0.23**	0.06	-0.02	-0.12	-0.14	-0.06*	0.05*
g_t	0.02	0.42**	0.25**	0.52*	-1.17**	0.07*	-0.01
π_t	-0.11*	0.05	0.53**	-0.75**	0.05	0.02	0.02
s_t	0.01	-0.02	-0.02	0.65**	-0.61**	0.00	-0.01*
v_t	0.01	0.01*	0.01	0.08**	0.91**	-0.00	-0.00*
i_t	-0.32*	-0.35*	0.26	0.63	-0.87	0.79**	0.31**
y_t	1.85**	0.40*	-0.05	0.59	0.90	0.14	0.52**
$100 \times std(\varepsilon_{t+1})$	2.22	2.15	2.28	7.34	9.04	1.24	0.77
Corr $\varepsilon, \varepsilon_\pi$	-0.14	0.21	1.00	-0.07	-0.28	0.15	0.17
R^2	0.68*	0.32*	0.56*	0.54*	0.96*	0.84*	0.91*
$100 \times std(x)$	3.92	2.61	3.44	10.80	42.76	3.05	2.60

Table 4: VAR estimate. Sample 1930-2018. One (two) stars means the estimate is one (two) Monte Carlo standard errors away from zero.

Table 4 presents OLS estimates of the VAR in the full sample. They are broadly similar to Table 2, with a few exceptions. Growth does not predict surpluses as strongly, and disinflation predicts surpluses more strongly. WWII deficits came with growth, the opposite of the postwar pattern, and followed the deflation of the 1930s. Table 5 presents OLS estimates of the filtered VAR in the full sample.

Table 6 presents variance decompositions in the full sample. They paint a broadly similar picture. Even more, 59% rather than 52%, of the variation in the value of debt corresponds to discount rate variation, and less, 43% rather than 55%, to expected future surpluses. The main

	r^n_{t+1}	g_{t+1}	π_{t+1}	s_{t+1}	v_{t+1}	i_{t+1}	y_{t+1}
r^n_t	-0.04	0.00	-0.07	-0.13	0.16	-0.17**	-0.02
g_t	-0.02	0.29**	0.25*	0.29	-0.85*	0.06*	-0.02
π_t	-0.01	-0.02	0.37**	-0.90**	0.53*	-0.04	-0.00
s_t	0.04*	-0.03*	0.05*	0.59**	-0.57**	-0.01	-0.01*
v_t	0.01	0.04**	0.02*	0.23**	0.72**	-0.00	-0.00
i_t	-0.08	-0.49*	0.24	0.34	-0.16	0.70**	0.30**
y_t	2.83**	0.57*	-0.00	0.45	1.82*	-0.50**	0.05
$100 \times std(\varepsilon_{t+1})$	2.22	2.15	2.28	7.34	9.04	1.24	0.77
Corr $\varepsilon, \varepsilon_\pi$	-0.14	0.21	1.00	-0.07	-0.28	0.15	0.17
R^2	0.58*	0.30*	0.34*	0.46*	0.78*	0.45*	0.33*
$100 \times std(x)$	3.44	2.57	2.81	9.96	19.13	1.67	0.94

Table 5: Filtered VAR estimate. Variables are the difference to a three year lagged moving average. Sample 1930-2018. One (two) stars means the estimate is one (two) Monte Carlo standard errors away from zero.

	$(\cdot)/var(\tilde{v}_t)$		
Component	Plain	Innovation	Filtered
$var(\tilde{v}_t)$	1.00	1.00	1.00
$cov(\tilde{v}_t, \Sigma\tilde{s}_{t+j})$	0.43	0.04	0.69
$-cov[\tilde{v}_t, \Sigma(\tilde{r}^n_{t+j} - \tilde{\pi}_{t+j})]$	0.59	0.93	0.22
$cov[\tilde{v}_t, \Sigma\tilde{g}_{t+j}]$	-0.02	0.03	0.09
$-cov(\tilde{v}_t, \Sigma\tilde{r}^n_{t+j})$	0.68	1.20	0.06
$cov(\tilde{v}_t, \Sigma\tilde{\pi}_{t+j})$	-0.09	-0.27	0.16
$\sigma(\tilde{v}_t)$	0.43	0.09	0.19

	$(\cdot)/var(\tilde{v}_t)$					
Component	Plain		Innovation		Filtered	
	25%	75%	25%	75%	25%	75%
$cov(\tilde{v}_t, \Sigma\tilde{s}_{t+j})$	0.39	0.77	-0.03	0.56	0.68	0.81
$cov[\tilde{v}_t, \Sigma(\tilde{r}^n - \tilde{\pi})_{t+j}]$	-0.56	-0.23	-0.89	-0.40	-0.23	-0.13
$cov(\tilde{v}_t, \Sigma\tilde{g}_{t+j})$	-0.04	0.09	-0.02	0.19	0.05	0.11
$-cov(\tilde{v}_t, \Sigma\tilde{r}^n_{t+j})$	0.58	0.18	1.03	0.48	0.13	0.07
$cov(\tilde{v}_t, \Sigma\tilde{\pi}_{t+j})$	-0.11	0.15	-0.32	0.10	0.04	0.12
$\sigma(\tilde{v}_t)$	0.34	0.44	0.08	0.09	0.18	0.22

Table 6: Decomposition of the variance of the value of debt, full sample 1930-2018. Bottom panel: Monte Carlo quantiles.

change is that nominal rates correspond less to inflation, and growth disappears from 7% to 2%. Given the story, verified above, that a good deal of the WWII debt to GDP ratio was resolved by greater growth, the disappearance of growth from the decomposition is interesting. The *rise* in debt to GDP ratio in WWII had nothing to do with growth, which was positive rather than negative. The decomposition's insistence on averaging all episodes rather than focusing on one story bears fruit.

In the full sample, the variance of filtered value likewise corresponds more, 22% rather than 9%, to discount rate variation, and less, 69% rather than 85%, to surpluses. The variance of value innovations now corresponds entirely, 93%, to discount rates and not at all, 4%, to surpluses.