The Value of Government Debt

John H. Cochrane[*]

August 12, 2019

Abstract

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.

---

[*]Hoover Institution, Stanford University and NBER. Data, code, and updates are at http://faculty.chicagobooth.edu/john.cochrane.

1.   Introduction

This paper measures the sources of the value of government debt. I use data provided by Hall, Payne, and Sargent (2018) on the market value of U.S. federal debt and the rate of return on the portfolio of that debt.

I start by looking backward: What combination of initial debt, surpluses and deficits, inflation, bond returns, and growth produce each date's debt to GDP ratio? How were government debts incurred, and how were they paid off? I find, unsurprisingly, that the deficits of the 1930s and 1940s account for a lot of the 100% debt to GDP ratio at the end of WWII. However, about half of the rise in debt to GDP ratio in the Great Depression reflects real returns due to deflation and the fall in GDP. The rise debt to GDP ratio in WWII is only about two-thirds as large as the cumulative deficits, with low returns and the rise in GDP now tempering the rise in debt to GDP. About a third of the postwar fall in debt to GDP did come from sustained primary surpluses. The rest came from the larger fall in the cumulative growth-adjusted return $r - g$. Nominal returns and inflation largely cancel, with rising GDP contributing the rest. Starting in the mid 1970s, sustained primary surpluses turn to repeated large primary deficits, and $r - g$ turns slightly positive. Subsequent variation in the value of debt largely corresponds to variation in cumulated primary surpluses and deficits. Variation in rates of return, though large, is not sustained enough to produce much variation in the value of debt. The inflation of the 1970s did little to devalue debt, as debt rolled over faster than unexpected inflation could devalue it. Using the correct return on the entire government debt portfolio, there is little sign of a substantial post-2000 decline in $r - g$ that might make debt more sustainable.

The core of the analysis looks forward. The value of debt at each date is equal to the present discounted value of future surpluses. What combination of expected future surpluses, growth, and discount rates account for the value of debt to GDP ratio at each date?

I start with an ex-post rational calculation, which also establishes a useful set of facts for the later variance decomposition. Suppose that people know what future surpluses, growth, and real government bond returns will be. How do those terms combine to account for the value of debt at each date? This calculation tells a similar story backwards. After 1975-1980, variation in value of the debt largely follows the value of future surpluses, with little contribution from growth and discount rates. The high value of WWII debt comes about 2/3 from low growth-adjusted discount rates, with 1/3 future primary surpluses.

Of course people are not clairvoyant. For example, people may not have expected that the postwar era would have such unprecedented growth and low real returns. People may have ex-

pected that much higher taxes would be needed to pay off WWII debt. Without trying to measure what people expected in a given circumstance, we can measure the path of surpluses, discount rates, and growth *on average* following periods of high debt. This calculation can tell us why debt values vary on average, though it does not tell us people's expectations at each date.

In this paper's central calculation, then, I follow the asset pricing literature that decomposes the volatility of the price-dividend ratio to contributions of expected dividend growth and discount rates. I start with the identity, that 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 less its covariance with future real returns less growth.

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 expected growth-adjusted discount rates. The discount rate contribution comes mostly from nominal returns. Essentially none of the variation in debt to GDP ratios corresponds to changing growth forecasts. In the full sample, going back to 1930, even more of the variation in the debt to GDP ratio comes from discount rates, and less from variation in expected future surpluses.

Much of this variation in debt to GDP ratios and discount rates comes at very low frequency. To focus on business cycle variation, I develop a variant of the variance decomposition for filtered data, and for innovations. Since the discount rate movement is lower frequency than the surplus movement, the variance of filtered debt corresponds more to surplus variation. The variance of innovations in the value of debt, however, is dominated by discount rate movements. When interest rates decline, the value of debt rises, so this fluctuation in the value of (marked to market) debt corresponds to lower subsequent expected returns.

1.1.   Framing and literature

Why are these calculations interesting, and novel? We live in a time of unusually large peacetime government debt to GDP ratios. How will these debts be resolved – by primary surpluses, by low returns, or by large GDP growth rates? What set of expectations sustains these large debts – and what will happen if those expectations of surpluses, growth, and low returns change? Can we count on growth to exceed returns for the foreseeable future, suggesting a rather painless further expansion of government debt, as Blanchard (2019) has recently argued? Or are we primed for a debt crisis? We cannot know, of course, but the lessons of history captured by these calculations are illuminating: We can learn how past variation in debt to GDP ratios was resolved, and we can learn how debts are resolved on average. The latter illuminates the underlying set

of expectations, just as learning that high price-dividend ratios are resolved on average by low returns illuminates the expectations that produce high prices (see Cochrane (2011) for a short summary). Beyond analysis of current events, the expectations underlying the value of debt are a central part of many macroeconomic models.

Most analysis of the history and likely resolution of debts focuses on the likelihood of greater surpluses, via greater tax revenues or less spending, or unexpected low returns due to inflation or default. Less attention is focused on the possibility of higher growth rates, or the consequences to long-run growth of higher taxes, or on the discount rate (expected return) of government debt. A big message of this paper is that discount rates matter a lot to understanding debt dynamics, especially at the very low and very high frequencies.

I adapt asset pricing techniques to the study of government debt and deficits. The general approach to linearizing the valuation identity follows Campbell and Shiller (1988). I linearize in the level of the surplus to GDP ratio, or surplus to value ratio, rather than 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 $var(x) = var(y) + var(z) + 2cov(y, z)$. I use a VAR rather than direct long-horizon regressions to estimate long-run covariances.

The analysis of government finances, how debt is paid off, grown out of, or inflated away, is a long literature. Hall and Sargent (2011), Hall and Sargent (1997) are the most important immediate precursors. (Their references review well the prior literature.) I use data provided by Hall, Payne, and Sargent (2018). 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, not coupon payments or the short-term interest rate. I measure the surplus and deficit by how much the government actually borrows, not the NIPA deficit. My backward decomposition gives similar results to Hall and Sargent's. 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. The central novelty of this paper is the forward-looking decompositions, and in particular the asset-pricing-based decomposition of the variance of debt to GDP ratios stemming from the present value relation.

This paper uses the valuation identity, value of debt equals the present value of surpluses, and its time-t innovation, to study the value of debt. Cochrane (2019) takes a time-$t + 1$ innovation of the present value identity to study the fiscal roots of inflation and monetary policy. Though they share data, a VAR, and a linearization of the present value identity, the issues are orthogonal and independent. Inflation turns out not to matter here, and the value of the debt

drops out of the analysis there.

2. Identities

I start with a linearized version of the government debt flow identity,

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

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^n_{t+1} 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 value to GDP ratio. Both definitions lead to the same linearized identity. I infer s_{t+1} from the identity (1), and 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 evaluates the accuracy of the identity.

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

$$v_t = v_0 - \sum_{j=1}^{t} s_j + \sum_{j=1}^{t} \left(r^n_j - \pi_j - g_j \right). \tag{2}$$

This identity tells us where the debt to GDP ratio came from at each date – what combination of past surpluses and deficits, yes, but also bond returns, inflation, and growth, drove the debt to GDP ratio at each date t.

Iterating forward, we have a present value identity,

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

The value of government debt, divided by GDP, is the present value of future surpluses, discounted at the ex-post real return adjusted by GDP growth. (I derive the linearization around the point $r = g$. If one linearizes around $r > g$, discounting terms β^j appear in the above sums, with $\beta < 1$. However, such discounting is not important for this application, so I use $\beta = 1$ for simplicity.)

The identity (3) holds ex-post, because it discounts at the ex-post return. It is, essentially, the definition of ex-post return, rearranged. Taking conditional expectation, with any informa-

tion set that contains v_t, or any set of probabilities, it therefore also holds ex-ante. We can write

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

In particular, (4) holds with the information set of a VAR. We do not need to assume that agents have no more information than the VAR. However, the resulting estimates only reflect agent's expectations on average, conditioned down to the VAR variables.

So long as the debt to GDP ratio v_t is stationary, the last term converges, and looking at variation about the mean it converges to zero. The last term also converges to zero if we linearize around $r > g$, which adds β^j terms to the sums.

Let tildes stand for three possible transformations of variables,

$$\tilde{v}_t = v_t \text{ or } \theta(L)v_t \text{ or } (E_t - E_{t-1})\, v_t$$

and similarly for the other variables. Applying any of these linear transformations to the the present value identity (3), the same identify holds for the transformed versions.

Multiplying by $\tilde{v}_t - E(\tilde{v}_t)$ and taking expectations, we obtain a variance decomposition for the value of debt,

$$var\left(\tilde{v}_t\right) = \sum_{j=1}^{\infty} cov(\tilde{v}_t, \tilde{s}_{t+j}) - \sum_{j=1}^{\infty} cov\left[\tilde{v}_t, \left(\tilde{r}_{t+j}^n - \tilde{\pi}_{t+j} - \tilde{g}_{t+j}\right)\right].$$

Dividing by $var\left(\tilde{v}_t\right)$, we can express the result as fractions of variance due to each component. Such components are also the coefficients of single regressions by which $\tilde{v}_t$ forecasts the other terms. They measure the answer to 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 high growth rates or low real returns? And if the latter, were the low real returns the result of low nominal returns or high inflation?

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.

Why transform? A quick look at the data, below, reveals an obvious issue: U. S. debt is dominated by the runup in WWII, its steady decline, and it so-far-unresolved second runup. Facts are facts, but two to three effective data points are not much on which to take variances

and covariances. It is also economically interesting to separate debt variation into components. War finance clearly has different roots than cyclical surpluses and deficits, and war debts are likely resolved in different ways from business cycle or other debts.

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

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

and

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

but we cannot write

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

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 $(1 - L)$ yields $E_t y_{t+1} - E_{t-1} y_t = 0$. Instead, $E_t(y_{t+1} - y_t) = -y_t \neq 0$. 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. (Filtering is less useful for asset pricing applications, as the hypothesis that filtered returns are unpredictable is not particularly interesting.)

I use a VAR to calculate conditional expectations. Denote the VAR

$$x_{t+1} = A x_t + \varepsilon_t,$$

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

$$v_t = a_v' x_t.$$

To estimate the variance decomposition for plain $\tilde{v}_t = v_t$ case, write the present value identity (4) as

$$v_t = a_s'(I - A)^{-1} A x_t - a_{rg}'(I - A)^{-1} A x_t, \tag{5}$$

where

$$a_{rg} \equiv a_{r^n} - a_{\pi} - a_g.$$

The variance decomposition of the value of debt, with $\tilde{v}_t = v_t$, is then

$$a_v'Va_v = a_s'(I - A)^{-1}AVa_v - a_{rg}'(I - A)^{-1}AVa_v \tag{6}$$

where

$$V \equiv cov(x_t, x_t')$$

is the covariance of the data.

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

$$\Delta E_t v_t = \Delta E_t \sum_{j=1}^{\infty} s_{t+j} - \Delta E_t \sum_{j=1}^{\infty} \left(r_{t+j}^n - \pi_{t+j} - g_{t+j} \right) \tag{7}$$

where $\Delta E_t \equiv E_t - E_{t-1}$. Therefore, in VAR notation

$$a_v'\varepsilon_t = a_s'(I - A)^{-1}A\varepsilon_t - a_{rg}'(I - A)^{-1}A\varepsilon_t$$

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

$$a_v'\Sigma a_v = a_s'(I - A)^{-1}A\Sigma a_v - a_{rg}'(I - A)^{-1}A\Sigma a_v$$

where

$$\Sigma \equiv cov(\varepsilon_t, \varepsilon_t')$$

is the covariance matrix of the VAR shocks. This formula is the same as (6), 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.

(In Cochrane (2019), I take the time $t+1$ innovation $E_{t+1} - E_t$ of the present value identity (4), which eliminates the value of debt v_t on the left hand side, in order to focus on the roots of inflation. Here I take the time t innovation $E_t - E_{t-1}$ of the present value identity, to focus on the innovation of the value of debt.)

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 formula (6).

These decompositions are all measurement, and there is no particularly interesting hypothesis to test. The underlying present value relation is an identity, so there is nothing interesting to test there either. Which component of an identity matters is an interesting measurement, however.

A tempting mistake is to calculate from the VAR at each date

$$E_t \sum_{j=1}^{\infty} s_{t+j} - E_t \sum_{j=1}^{\infty} \left(r_{t+j}^n - \pi_{t+j} - g_{t+j} \right)$$

and then compare the result to v_t. Again, the present value relation is an identity, so if one correctly includes v_t in the VAR, the answer comes out to v_t exactly. Leaving v_t out of the VAR violates the conditioning-down assumption that leaves v_t and not $E(v_t|$ a lesser information set) on the left hand side.

We can, however, compute the terms on each date to see how the VAR decomposes the value of debt v_t to its various underlying components. From (5),

$$v_t = \left[a_s' - \left(a_{r^n}' - a_\pi' - a_g' \right) \right] (I - A)^{-1} A x_t, \tag{8}$$

we can compute and plot at each date the terms on the right hand side. Such expectations use more information than just v_t to forecast at each date, but not the full information set used by bond holders.

3.   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 infer the primary surplus from the flow identity (1). This calculation measures how much money the government actually borrows. NIPA surplus data, though broadly similar, does not obey the identity. I infer the surplus from the linearized identity to produce data that exactly matches that identity, rather than carry around an approximation error in all the calculations. Cochrane (2019) shows that the approximation error is small for these purposes.

I measure the debt to GDP ratio by the ratio 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. We want only a detrending divisor, and an indicator of the economy's long-run level of tax revenue and spending. Consumption is a decent
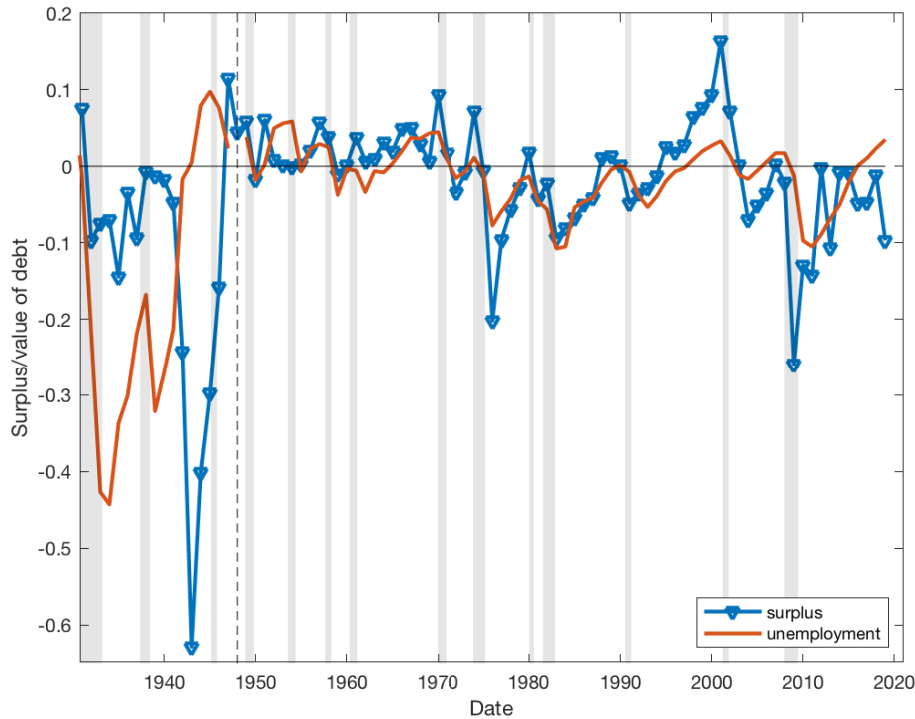
stochastic trend for GDP.

Figure 1: Surplus and negative of the unemployment rate. The unemployment rate is rescaled to fit on the same graph.

Figure 1 presents the surplus. The vertical dashed line indicates 1947 at which I consider a subsample below. There *are* primary surpluses. One's impression of endless deficits comes from the full deficit including interest payments on the debt. NIPA measures also show regular positive primary surpluses. After the great depression and WWII there is an era of steady primary surpluses, which, as we will see, helped to pay off 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. After 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.

I use CRSP data for the three-month Treasury rate. To measure a long-term bond yield, I use the 10-year constant maturity government bond yield from 1953 on, the yield on long-term United States bonds from 1941 to 1952, and the Fed H.15 discontinued composite yield on treasury bonds with maturity over ten years for 1929 to 1941.
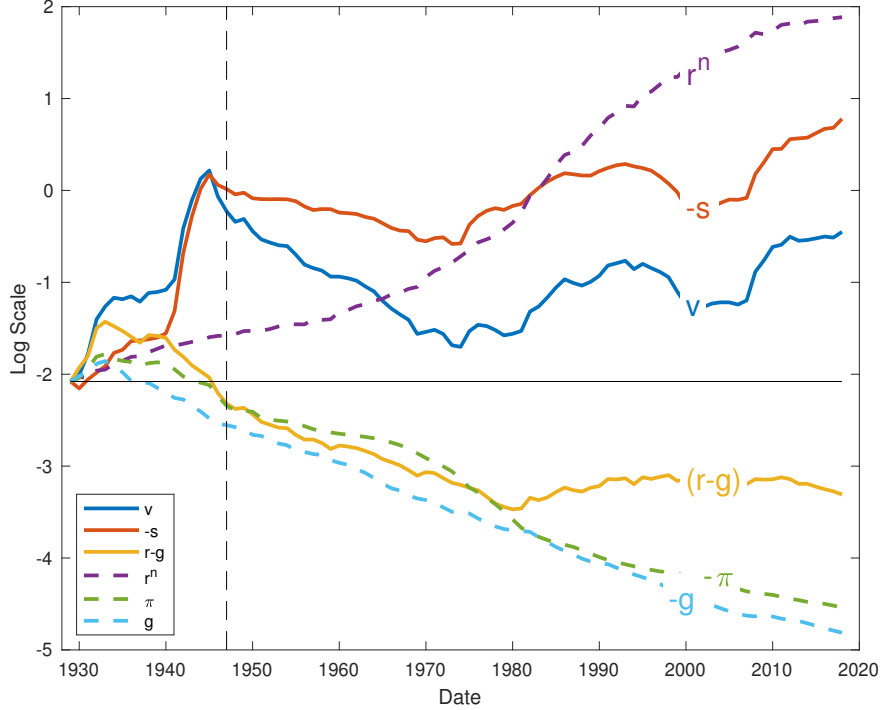
4.   The history of debt

Figure 2: Backwards decomposition of the value of debt.

Figure 2 presents the terms of the value decomposition (2). At each date t, it plots the terms of

$$v_t = v_0 - \sum_{j=1}^{t} s_j + \sum_{j=1}^{t} \left(r_j^n - \pi_j - g_j \right)$$

Each line presents what debt v_t would be at date t, starting from v_0, if that were the only term, setting the others 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 break $r - g$ down into its three components, nominal return r^n, inflation π, and growth g. Though the method is somewhat different, the results are similar to Hall and Sargent (2011) Figure 7.

In the v_t 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 rises in the great depression and in WWII, to a debt-to-GDP ratio greater than one i.e. $v_t = 0 = \log(1)$. The $-s$ line is the primary deficit, the negative of surplus, i.e. how surplus contributes to debt. The rise in debt correlates with the surplus line, i.e. by the deficits of the 1930s and WWII. Returns matter as well, however. About half of the rise in market value of debt to GDP ratio in the great depression comes from the $r - g$ term, reflecting the fall in GDP and deflation of the early 1930s. Likewise,

the rise in market value of 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 value of debt to GDP ratio fell steadily. About a third of this fall comes from primary surpluses, as shown by the fall in the $-s$ line. The rest came from the larger fall in the cumulative growth-adjusted return $r - g$. From 1940 to 1980, real returns are less than growth – the cumulative $r - g$ falls constantly. Breaking $r - g$ out to nominal returns, inflation, and growth, we see that nominal returns and inflation largely cancel, leading to the standard conclusion, that about 2/3 of the fall in debt to GDP ratio came from rising GDP. However, about 1/3 of the fall did come from the persistent primary surpluses of the first 30 postwar years, seen in Figure 1. It wasn't all growth.

The inflation of the 1970s is visible, in a speeding up of the fall in the π line. But returns in the r^n line rise as well, so the net effect $r = r^n - \pi$ is not as large as one would suppose. To devalue outstanding debt, inflation must come swiftly, relative to the maturity structure of outstanding debt. Expected inflation cannot devalue short-term debt, because its interest rate rises. 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. Variation in the value of debt came largely from surpluses, growth, and real returns – in the US, in this episode.

(That even the inflation of the 1970s does not contribute in a big way to the value of the debt is one reason why Cochrane (2019), which studies the fiscal roots of inflation, focuses on an accounting in which the value of the debt drops out, rather than this value of the debt accounting, and why the two calculations turn out to have separate lessons for separate issues.)

A sharp break 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 $r - g$ is small. That cumulative effect includes a change in behavior. First, $r - g$ is no longer negative; Cumulative $r - g$ rises slightly. Second, underlying that behavior, the nominal r^n and real $r^n - \pi$ returns rise – bonds did well in the 1980s and 1990s. Growth does not change that much in this graph, and inflation declines. Since 2000 bond returns have declined – the r^n line does not rise as fast – but inflation and growth also declined. Cumulatively, $r - g$ flattened around 2000 and gently declines after 2010.

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, 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.

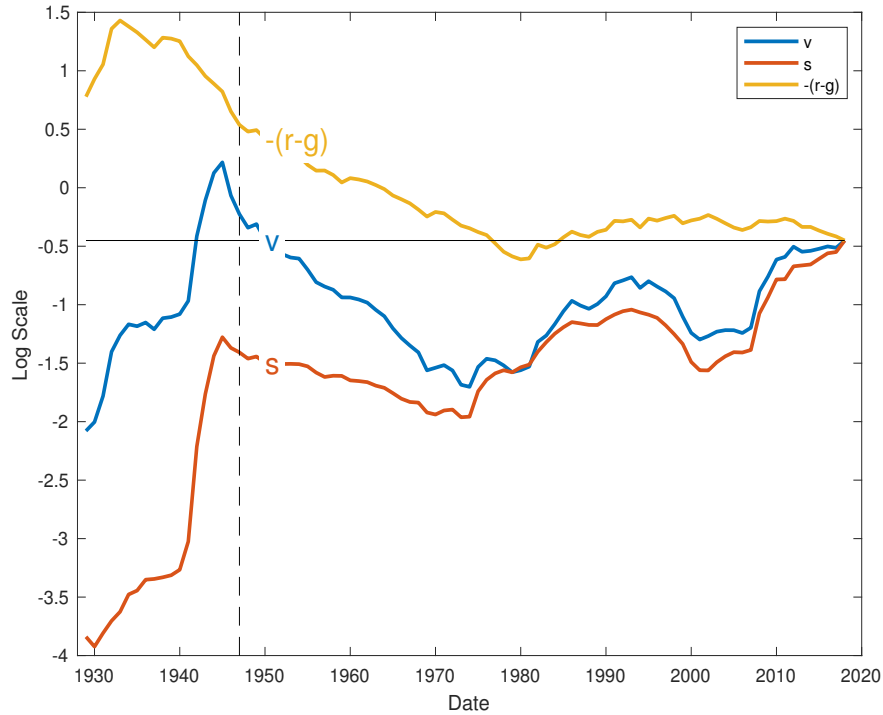## 4.1. A forward-looking value decomposition

Figure 3: Forward decomposition of the market value of debt to GDP ratio.

Figure 3 presents a forward-looking value decomposition,

$$v_t = \sum_{j=1}^{T} s_{t+j} - \sum_{j=1}^{T} \left(r^n_{t+j} - \pi_{t+j} - g_{t+j} \right) + v_{t+T}.$$

Regarding the value of the debt as the present value of future surpluses, 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 largely follows the variation in the sum of future surpluses, with little contribution from discount rates $r - g$. Why, for example, is the value of debt low in 1980? A big part of that value is foreknowledge that the value of the debt will be high in 2018, v_{t+T} – there will not be a default, an inflation, a growth disaster, or another big hit to $r - g$. Beyond that, the value of debt is low because on net, there will be a string of primary

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

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 backwards. Why was the value of debt so high at the end of WWII? About 2/3 of it was low growth-adjusted discount rates, with about 1/3 knowledge of primary surpluses to come.

Again, though, it is likely 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 move to the central calculation using *expected* present values.

4.2. Decompositions

I base the variance decomposition calculations on a first-order VAR consisting of the nominal return on the government bond portfolio r^n, the consumption growth rate g, inflation π, surplus s, value v, the three-month interest rate i and the 10 year bond yield y. The latter are important forecasting variables for growth, inflation, and long-term bond returns. (This is the same VAR as in Cochrane (2019).)

Table 1 presents OLS estimates of the VAR coefficients. Each column is a separate regression. (This is the same VAR as in Cochrane (2019).) 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 1

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.

Looking down the columns, the long and short interest rates help to forecast the government bond portfolio return and growth; the short rate helps to forecast inflation. Growth and debt predict higher surpluses. The value of debt is very persistent, with an 0.98 coefficient on its lag, and thereby constitutes the dominant root for long-term forecasts. Higher surpluses and GDP growth also lower the value of the debt. Interest rates forecast the value of the debt, largely mirroring their forecast of returns on the debt. Interest rates are persistent, and respond to growth though not to inflation. It is comforting that the VAR offers sensible coefficients. Long-run forecasts are driven by eigenvalues of this coefficient matrix, however, which are not easy to relate to individual coefficients.

I focus on the sample since 1947, since WWII is obviously a different phenomenon than the subsequent cyclical variation, and its deficits are so huge they dominate any estimates. Appendix Table 6 presents OLS estimates of the VAR in the full sample. They are broadly similar Table 1 , 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.

In addition to analyzing the log value of debt to GDP $\tilde{v}_t = v_t$, I filter by using the difference between the log value of debt and three lags,

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

and I use the VAR innovation in the value of debt,

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

Figure 4 presents the filtered value of debt and innovation to the value of debt. Both measures pick up familiar cyclical movements, and de-emphasize the large long-term variation of the debt v_t visible in the previous two figures. For example, you can see the big increase in debt 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 at the end of the sample, filtered debt is small because debt did not increase at an unusually large pace.

Appendix Tables 7 and 8 present the OLS estimates of the filtered VAR. The coefficients are similar though reduced in magnitude. The lower diagonals reflect quicker dynamics as we
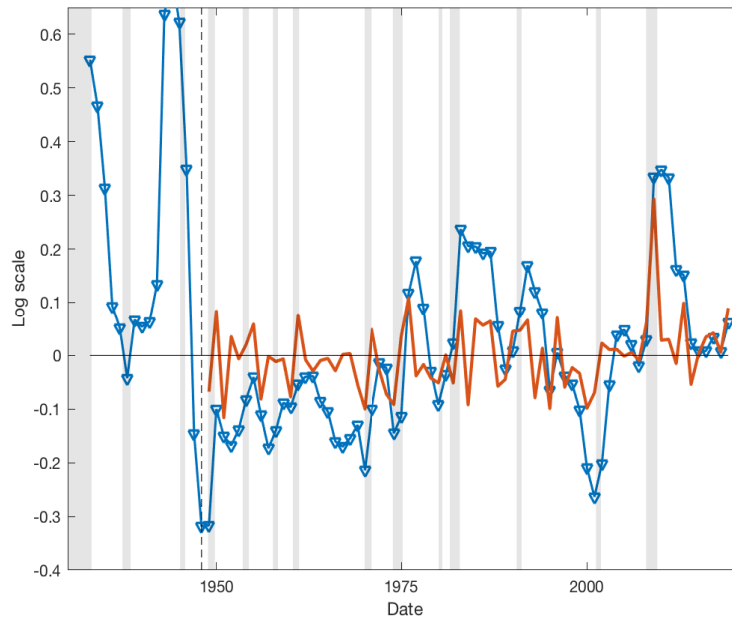
Figure 4: Filtered value of debt v_t^f (symbols) and the VAR innovation in the value of debt, $v_t - E_{t-1}(v_t)$. The filter is $v_t - \frac{1}{3}(v_{t-1} + v_{t-2} + v_{t-3})$. The vertical dashed line denotes the beginning of the sample used for VAR and statistical analysis. Shaded areas are NBER recessions.

expect of filtered variables.

4.3. Variance decomposition

As with asset pricing volatility tests, we cannot tell what expectations justify a value on a given date, but we can tell what expectations justify a given value *on average*, by seeing what combination of surpluses and returns follow that value on average. Under rational expectations these conditional expectations are conditioned-down measures of agents' expectations, and right on average.

Table 2 presents the decompositions of the variance of each of these measures of the value of debt $\tilde{v}_t$. The table entries are fractions of the variance of each measure of the value of debt. Table 3 presents the same decompositions for the full sample. Tables 4 and 5 present 25% and 75% quantiles of the Monte Carlo distribution of each of these statistics.

The last row of Table 2 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 variance of the debt innovation and filtered debt are lower, as they eliminate the large low-frequency variation, leaving standard deviations of 7

Component	Plain	$(\cdot)/var(\tilde{v}_t)$ Innovation	Filtered
$var(\tilde{v}_t)$	1.00	1.00	1.00
$cov(\tilde{v}_t, \Sigma \tilde{s}_{t+j})$	0.55	0.28	0.85
$-cov[\tilde{v}_t, \Sigma(\tilde{r}^n - \tilde{\pi} - \tilde{g})_{t+j}]$	0.45	0.72	0.15
$cov(\tilde{v}_t, \Sigma \tilde{r}^n_{t+j})$	-0.90	-1.63	-0.14
$cov(\tilde{v}_t, \Sigma \tilde{\pi}_{t+j})$	-0.38	-0.78	-0.05
$cov[\tilde{v}_t, \Sigma(\tilde{r}^n - \tilde{\pi})_{t+j}]$	-0.52	-0.85	-0.09
$cov(\tilde{v}_t, \Sigma \tilde{g}_{t+j})$	-0.07	-0.13	0.07
$\sigma(\tilde{v}_t)$	0.37	0.07	0.14

Table 2: Decomposition of the variance of the value of debt. Sample 1947-2018.

Component	Plain	$(\cdot)/var(\tilde{v}_t)$ 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
$-cov[\tilde{v}_t, \Sigma(\tilde{r}^n - \tilde{\pi} - \tilde{g})_{t+j}]$	0.57	0.96	0.31
$cov(\tilde{v}_t, \Sigma \tilde{r}^n_{t+j})$	-0.68	-1.20	-0.06
$cov(\tilde{v}_t, \Sigma \tilde{\pi}_{t+j})$	-0.09	-0.27	0.16
$cov[\tilde{v}_t, \Sigma(\tilde{r}^n - \tilde{\pi})_{t+j}]$	-0.59	-0.93	-0.22
$cov(\tilde{v}_t, \Sigma \tilde{g}_{t+j})$	-0.02	0.03	0.09
$\sigma(\tilde{v}_t)$	0.43	0.09	0.19

Table 3: Decomposition of the variance of the value of debt. Sample 1930-2018.

| Component | $(\cdot)/var(\tilde{v}_t)$ | | | | | |
| | Plain | | Innovation | | Filtered | |
	25%	75%	25%	75%	25%	75%
$var(\tilde{v}_t)$	1.00	1.00	1.00	1.00	1.00	1.00
$cov(\tilde{v}_t, \Sigma\tilde{s}_{t+j})$	0.51	0.75	0.23	0.54	0.71	0.81
$-cov[\tilde{v}_t, \Sigma(\tilde{r^n} - \tilde{\pi} - \tilde{g})_{t+j})]$	-0.91	-0.48	-1.54	-0.91	-0.27	-0.16
$cov(\tilde{v}_t, \Sigma\tilde{r}^n_{t+j})$	0.25	0.49	0.46	0.77	0.19	0.29
$cov(\tilde{v}_t, \Sigma\tilde{\pi}_{t+j})$	-0.39	-0.19	-0.75	-0.38	-0.10	-0.03
$cov[\tilde{v}_t, \Sigma(\tilde{r}^n - \tilde{\pi})_{t+j})]$	-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
$\sigma(\tilde{v}_t)$	0.29	0.40	0.06	0.07	0.12	0.15

Table 4: Monte Carlo quantiles for the decomposition of the variance of the value of debt. Sample 1947-2018.

| Component | $(\cdot)/var(\tilde{v}_t)$ | | | | | |
| | Plain | | Innovation | | Filtered | |
	25%	75%	25%	75%	25%	75%
$var(\tilde{v}_t)$	1.00	1.00	1.00	1.00	1.00	1.00
$cov(\tilde{v}_t, \Sigma\tilde{s}_{t+j})$	0.39	0.77	-0.03	0.56	0.68	0.81
$-cov[\tilde{v}_t, \Sigma(\tilde{r^n} - \tilde{\pi} - \tilde{g})_{t+j})]$	-0.58	-0.18	-1.03	-0.48	-0.13	-0.07
$cov(\tilde{v}_t, \Sigma\tilde{r}^n_{t+j})$	0.23	0.61	0.44	1.03	0.19	0.32
$cov(\tilde{v}_t, \Sigma\tilde{\pi}_{t+j})$	-0.11	0.15	-0.32	0.10	0.04	0.12
$cov[\tilde{v}_t, \Sigma(\tilde{r}^n - \tilde{\pi})_{t+j})]$	-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
$\sigma(\tilde{v}_t)$	0.34	0.44	0.08	0.09	0.18	0.22

Table 5: Monte Carlo quantiles for the decomposition of the variance of the value of debt. Sample 1930-2018.

(innovation) and 14 (filtered) percentage points respectively.

The second and third rows of Table 2 give the central variance decompositions. Looking down the first column,

- *Half (55%) of the variance of the debt to GDP ratio $var(v_t)$ corresponds to variation in forecasts of future surpluses. Half (45%) comes from discount rate variation. Practically all of this discount rate variation comes from real returns (-0.52), and essentially none from variation in forecasted growth (-0.07). Most of the variation in expected real returns comes from variation in expected nominal returns (-0.90) rather than inflation (-0.38).*

That growth g is not important, on average here, illustrates the important difference between this variance accounting, the average experience following a large debt, and the previous ex-post history, focusing on one data point. Much of the WWII debt was repaid, ex-post, by growth. However, the subsequent debt episodes did not resolve by a lot of ex-post growth, and the VAR does not find the value of the debt to forecast growth. Growth is not large *on average* when debt is higher, and not small on average when debt is lower (e.g. the low debt to GDP 1970s followed by good, not bad, growth in the 1980s) so expected growth is not an important driver of the value of debt from an ex-ante point of view. 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.

Figure 3 captures the intuition of this result. 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.

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. (About half the variation in returns comes from each source, since returns add contemporaneous unexpected dividend growth.)

One big difference is important however: Here I account for variation in the *total market value* of debt, including (and dominated by) issuances and redemptions. The standard asset

pricing calculation accounts for the variation of the value of a particular security. To see the difference, note that the 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. Yet the variation in total value of debt will add rises and falls in debt as deficits and surpluses vary. 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. Discount rates operate on top of that factor, affecting the value of each individual bond.

The amount of low-frequency variation in the debt, as well as in cumulative surpluses and discount rates, is unfortunate for this sort of empirical work, which is why I investigate the variance of the filtered debt. Turning to the "Filtered" column of Table 2, we see important differences:

- *Variation in the filtered value of debt corresponds more, 85%, to future surpluses, and less, 15%, to discount rates.*

By filtering, we emphasize the medium and higher frequency fluctuation in surpluses, visible in Figure 3, and their relationship to the same fluctuations in flittered debt seen in Figure 4. The lower frequency variation in cumulative discount rates from Figure 3 is filtered out. Recession-related, and somewhat longer (Vietnam, 1980s deficits, 1990s surpluses) variation in the value of debt is driven by expected surpluses, even more than the secular movement in debts.

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

- *Variation in the unexpected value of the debt to GDP ratio corresponds less, 28%, to revisions in expected future surpluses, and more, 72%, to revisions in expected future discount rates. The discount rate variation comes almost entirely from a revision in expected real returns, (-0.85), not a revision in growth (-0.13). In turn, the innovation in expected real returns comes from a large revision in expected nominal returns (-1.63) offset by a smaller change in expected inflation (-0.78).*

Reviewing the underlying identity (7), the decomposition correlates the innovation in the value of debt $(E_t - E_{t-1}) v_t$ with *subsequent* surpluses and discount rates only, starting at e.g. $(E_t - E_{t-1}) r_{t+1}^n$, not contemporaneous returns which would contributed directly to the market value of debt. This result then reflects the classic identity that when bond prices rise, yields go down. An unexpected return r_t^n raises the value of debt v_t, and lowers the yield and expected

nominal return on long-term bonds $E_t r_{t+j}^n$. The 1.63 percentage point decline in nominal returns is matched by an 0.78 percentage point decline in inflation, yielding an 0.85 percentage point decline in real expected returns.

In sum, the value of debt inherits some of 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. As it is for banks and stocks, so it is for government debt.

Though they integrate over a clearly influential data point, WWII, and a potentially different regime, the full sample estimates in Table 5 paint a broadly similar picture. Even more, 57% rather than 45%, 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 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, 31% rather than 15%, to discount rate variation, and less, 69% rather than 85%, to surpluses. The variance of value innovations now corresponds entirely, 96%, to discount rates and not at all 4%, to surpluses.

The Monte Carlo quantiles in Tables 4 and 5 give a sense of the accuracy of these measurements. This is just measurement, nothing is tested here, so the point of the table is only to assess the accuracy of that measurement. The signs and broad magnitudes are reasonably measured, in the sense that signs and rough magnitudes are in the quartile range. Measuring long-run regressions is always difficult. 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, as can be seen by cases in which the point estimate is not at the midpoint of the quartiles. Direct measurement of long-run regressions, e.g. of the sum of future surpluses on the current value of debt, are likely even more uncertain.

4.4.   Time-series decomposition

Figure 5 presents VAR estimates of the decomposition of the value of debt at each date calculated from (8), e.g. $E_t \sum_{j=1}^{\infty} s_{t+j}$, labeled "$s$." The conditional expectations use all variables in the VAR,

not just the value of the debt. The value v line is equal to the sum of the surplus s line and the discount rate $-(r-g)$ lines. The dashed lines break out the latter to the real expected return $-r$ and growth rate g components.

The figure shows visually that variation in value v is composed about equally of variation in expected surpluses s and adjusted discount rates $-(r-g)$. There are two great waves of expected surplus, while discount rates trend down and then back up again. The discount rate operates at lower frequency. Discount rate variation comes almost entirely from variation in expected returns, with growth playing a tiny role. The current large value of the debt is, by this analysis, 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 long-term fiscal policy analysis. It does include bond yields as a strong signal of expected returns.
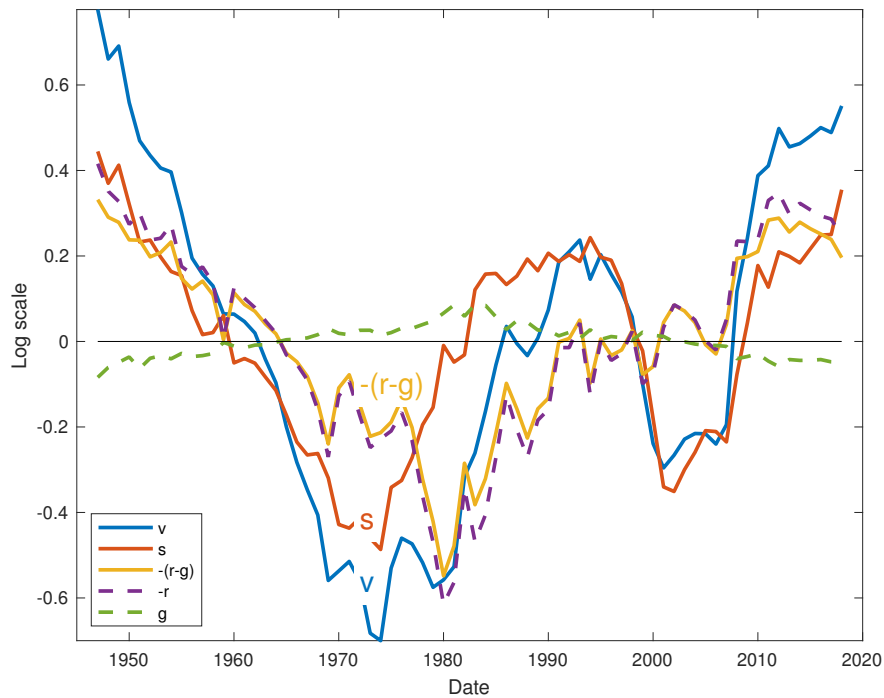

Figure 5: Decomposition of the value of debt to GDP ratio v_t. Each line presents the VAR estimate of the components of the value of debt at each date, i.e. $E_t \sum_{j=1}^{\infty} s_{t+j}$.

Figure 6 presents the same analysis for the filtered data. Befitting filtered data, the 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. It is clear to the eyeball here how variation in the cyclical value of debt corresponds more closely to variation in expected future surpluses, though high frequency discount rate variation still plays a role.
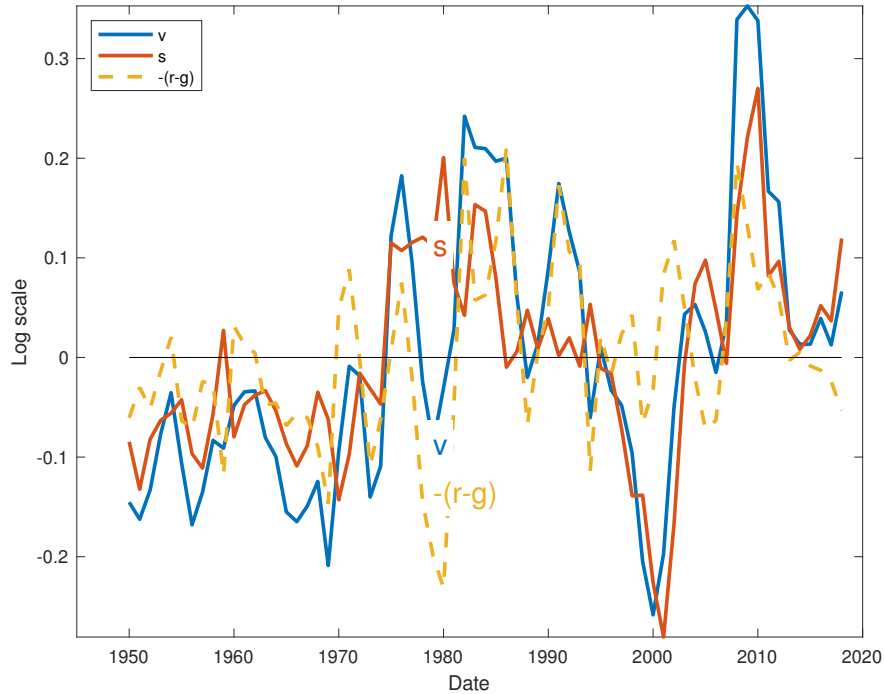
Figure 6: Decomposition of the filtered value of debt to GDP ratio. Each line presents the VAR estimate of the components of the filtered value of debt at each date, i.e. $E_t \sum_{j=1}^{\infty} \tilde{s}_{t+j}$.

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, 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 Federal 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/.

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.

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.

6. Additional Tables

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

	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
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 7: 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.

	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
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 8: 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.