

Debt and Deficits: Fiscal Analysis with Stationary Ratios

John Y. Campbell Can Gao Ian W. R. Martin*

April 2023

Abstract

We study cointegrating relationships among fiscal variables and output and use them to introduce a new measure of the government's fiscal position. In the US since World War II, we find that the primary surplus-GDP ratio and the government debt-GDP ratio are nonstationary, which invalidates standard analytical approaches that assume them to be stationary. The tax revenue-debt ratio and the government expenditure-debt ratio are also nonstationary but their difference, the primary surplus-debt ratio, is stationary, as is the tax revenue-GDP ratio. We develop a new framework for fiscal analysis that takes account of these facts. Empirically, we find that a deterioration in the fiscal position forecasts a decline in government spending over the long run. It does not forecast increases in tax revenue; nor does it forecast low returns for bondholders. Fiscal adjustment to tax and expenditure shocks occurs primarily through mean-reversion in tax and expenditure growth, with a negligible contribution from expected and unexpected debt returns. We find similar results for postwar UK data.

*John Campbell: Harvard University and NBER, <https://scholar.harvard.edu/campbell/home>.
Can Gao: University of St. Gallen, <https://sites.google.com/view/can-gao/home>. Ian Martin:
London School of Economics, <https://personal.lse.ac.uk/martiniw>.

When a government is in a weak fiscal position, holders of government debt must earn low returns over the long run, or taxes must rise over the long run, or spending must fall over the long run; or some combination of all three possibilities must occur. As we will show, this follows essentially as a matter of accounting. But which of the three channels is most important empirically?

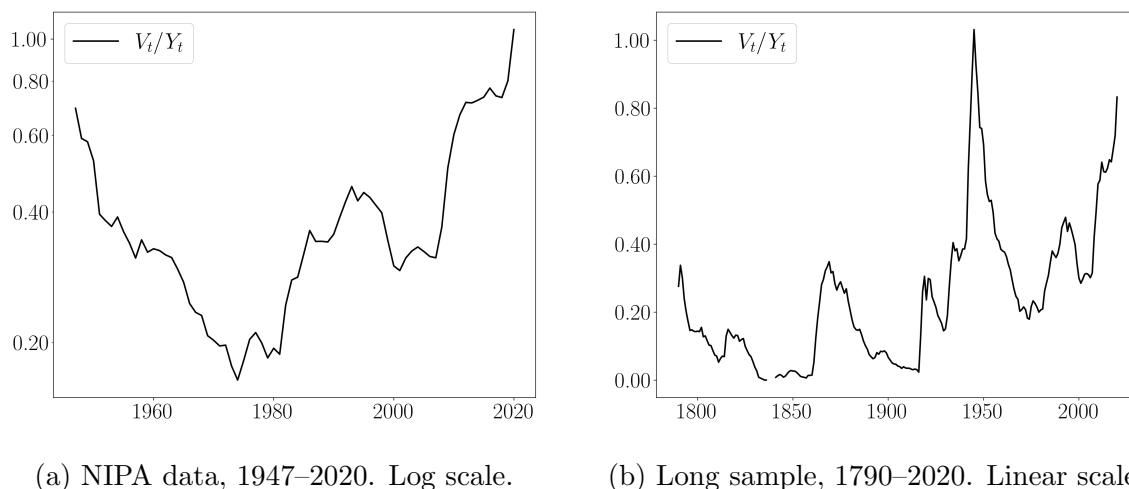
Any answer to this question requires a suitable definition of the “fiscal position”. We will argue that some seemingly natural definitions are problematic. Certainly the primary surplus of a government is an essential ingredient. The primary surplus—the excess of tax revenue over government expenditure—is the flow of resources that the government devotes to servicing its debt. When it is positive, the growth rate of the value of the debt is less than the return on the debt. When it is negative—that is, when the government runs a primary deficit—the debt grows at a faster rate than the return on debt. Under the standard assumption that the expected return on the debt exceeds its growth rate, the value of the debt is the expected discounted value of the primary surpluses that will service it in the future.

To be useful in fiscal analysis, the primary surplus must be scaled in some way so that the resulting ratio is stationary. A common approach, followed for example by [Bohn \(1998\)](#), [Cochrane \(2001, 2022, 2023\)](#) and [Jiang, Lustig, Van Nieuwerburgh, and Xiaolan \(2021b\)](#), is to divide the primary surplus by GDP to create the primary surplus-GDP ratio, and to divide the value of debt by GDP to create the debt-GDP ratio. These ratios are typically treated as stationary and analyzed in relation to one another.

Contrary to this approach, we find that neither the primary surplus-GDP ratio nor the debt-GDP ratio behave like stationary time series in US data since World War II. As [Figure 1](#), Panel a, shows, the debt-GDP ratio in particular has drifted persistently up and down for long periods of time, showing no strong tendency to return to a constant mean. Standard unit root tests fail to reject the null hypothesis that these ratios have a unit root, and cointegration tests fail to find statistically significant evidence that the primary surplus and government debt are cointegrated with GDP.

From a theoretical perspective, the nonstationarity of debt-GDP is not particularly surprising: for example, [Barro \(1979\)](#) writes, “There is no force that causes the ratio of debt to income to approach some target value”. Even if one believes that economic forces act to make the primary surplus-GDP ratio and the debt-GDP ratio truly stationary in the very long run—and the longer series shown in [Figure 1](#), Panel b does not support this view—the persistence of these time series implies that it is inadvisable to model

Figure 1: The debt-GDP ratio is nonstationary in US data.



them using the standard techniques of stationary time-series analysis (Campbell and Perron, 1991).¹

An alternative way to scale the primary surplus is to divide it by the value of debt to create the primary surplus-debt ratio. In an economy in which the return on the debt and the growth rate of the debt are stationary, the primary surplus-debt ratio should also be stationary: and indeed, standard unit root tests reject the null hypothesis that this ratio has a unit root in favor of the alternative that it is stationary.

The primary surplus-debt ratio is analogous in the fiscal context to the dividend-price ratio on a stock. Just as a corporation pays dividends to the owners of its stock, so the government pays primary surpluses to the owners of its debt. This suggests the possibility of analyzing the primary surplus-debt ratio using a Campbell and Shiller (1988) loglinearization to relate it to future log returns on debt and log growth rates of primary surpluses.

Two problems arise in doing so, and both result from the fact that the primary surplus can be negative. First, the log growth rate of the primary surplus is ill-defined when the surplus is negative. Second, an exogenous increase in the debt, which worsens the fiscal position of the government, can either raise or lower the primary surplus-debt ratio depending on whether the primary surplus is positive or negative. Thus, the effect of a given shock to the primary surplus-debt ratio depends on the sign of the ratio. Both

¹Appendix A.1.1 describes our data sources.

these problems also afflict the standard analysis of the primary surplus-GDP ratio.

In this paper we develop an alternative loglinear analysis, related to the work of [Giannitsarou, Scott, and Leeper \(2006\)](#) and [Berndt, Lustig, and Yeltekin \(2012\)](#), that solves these problems. Our approach is to approximate the primary surplus-debt ratio in a way that can be loglinearly related to the growth rates of tax revenue and of government expenditure. Both revenue and expenditure are always positive, so their log growth rates are well defined; and our loglinear approximation to the primary surplus-debt ratio has the appealing property that an increase in debt always reduces it, whether the primary surplus is currently positive or negative.

The approximations developed by [Giannitsarou et al. \(2006\)](#) and [Berndt et al. \(2012\)](#) are similar in spirit but rely on the assumption that the tax revenue-debt and government expenditure-debt ratios are stationary, so that one can approximate around their means. In the US data we find to the contrary that neither of these ratios are stationary. Instead, their logs are cointegrated with a cointegrating vector that is close to but not equal to a unit vector. We use this finding of cointegration to develop an approximation, related to the work of [Gao and Martin \(2021\)](#), that does not rely on inappropriate stationarity assumptions.

After presenting our approximation, we use it to explore the dynamics of debt, tax revenue, and government expenditure in US data since World War II. Despite the non-stationarity of the ratios of debt and the primary surplus to GDP, we find that the tax revenue-GDP ratio appears to be stationary. The government expenditure-GDP ratio, on the other hand, appears nonstationary, which accounts for the nonstationarity of the primary surplus-GDP ratio. Given these results, we estimate a vector autoregression (VAR) model including two growth rates—the return on debt and the growth rate of tax revenue—and the two ratios we have found to be stationary: the primary surplus-debt ratio (or rather our loglinearized approximation to it) and the tax revenue-GDP ratio.

Our main empirical findings are as follows. First, even though the primary surplus-debt ratio is stationary it is extremely persistent. A shock to this ratio has important effects even after several decades. Second, expected returns on government debt, while time-varying, are not variable or persistent enough to contribute importantly to the dynamics of the primary surplus-debt ratio. Third, the slow mean-reversion of the primary surplus-debt ratio occurs in the short run through changes in tax revenue, but in the longer run more than all the adjustment occurs through changes in government

expenditure. This finding relies critically on the inclusion of the tax revenue-GDP ratio in the VAR model, and it reflects the fact that faster growth of tax revenue raises the tax revenue-GDP ratio, predicting slower growth of GDP and eventually slower growth in tax revenue. Fourth, we examine the fiscal adjustment to shocks in tax revenue and government expenditures, and find that this occurs almost entirely through mean-reversion in the growth rates of taxes and expenditures. Expected returns on government debt again have little importance, and the same is true for unexpected returns on debt contemporaneous with tax and expenditure shocks.

We repeat the analysis for the UK and find similar results. While the evidence is not completely decisive, the surplus-debt ratio appears stationary and the debt-GDP ratio nonstationary, as in the US. We also find that the tax-GDP ratio is stationary whereas the spending-GDP ratio is nonstationary; again, these mirror our findings for the US, though the tax-GDP ratio is somewhat more persistent in the UK. We therefore estimate the same VAR system as we do for the US, and find the same sign pattern on statistically significant coefficients; as in the US, the variance decomposition reveals that shocks to the fiscal position are resolved, in the long run, by movements in spending rather than in taxes or returns.

Two caveats should be kept in mind when interpreting our results. First, because we conduct a reduced-form time-series analysis, we cannot make causal statements about fiscal dynamics. For example, our finding that an increase in the tax revenue-GDP ratio predicts slower GDP growth does not prove that high taxes cause lower growth as argued by [Reinhart and Rogoff \(2010\)](#) and [Alesina, Favero, and Giavazzi \(2020\)](#).

For the same reason we cannot resolve the debate about the fiscal theory of the price level ([Sargent and Wallace, 1981](#); [Leeper, 1991](#); [Sims, 1994](#); [Woodford, 1995](#); [Cochrane, 2001, 2023](#)). According to traditional analysis, the ability of the primary surplus-debt ratio to predict future fiscal adjustment is causal, in that a given value of the debt forces the government to run future primary surpluses that will pay it off. According to the fiscal theory of the price level, the predictive relationship reflects reverse causality: the debt has the value that is consistent with an exogenous path of future surpluses, as in a forward-looking asset pricing model of the sort analyzed by [Campbell and Shiller \(1987, 1988\)](#). If the debt promises to make fixed nominal payments, the required adjustment in value can occur largely through changes in the price level, although also in part through changes in long-term nominal interest rates ([Cochrane, 2001](#)).

Second, we take the returns on government debt as given, measuring them in the

data without requiring them to satisfy the restrictions of any asset pricing model. We do not address the question, studied by [Jiang, Lustig, Van Nieuwerburgh, and Xiaolan \(2021a\)](#), of whether the measured return is too low to be consistent with the risk of the government debt, or the related question, discussed by [Greenwood, Hanson, and Stein \(2015\)](#), [Krishnamurthy and Vissing-Jorgensen \(2012\)](#), [Reis \(2022\)](#), and [Mian, Straub, and Sufi \(2022\)](#), of whether government debt offers a convenience yield that investors value separately from its return.

The organization of the paper is as follows. In [Section 1](#) we present unit root and cointegration tests. These establish the basic empirical facts that motivate the framework for fiscal analysis introduced in [Section 2](#). We apply the framework to empirical analysis of US data in [Section 3](#), and repeat the analysis using UK data in [Section 4](#). [Section 5](#) concludes. An online appendix ([Campbell, Gao, and Martin, 2023](#)) presents supplementary details.

1 Some background facts

By definition, the gross return on government debt is

$$R_{t+1} = \frac{V_{t+1} + T_{t+1} - X_{t+1}}{V_t}. \quad (1)$$

Here R_{t+1} is the return on debt (including money) from time t to $t + 1$, V_t is the total value of debt (including money) in period t , T_{t+1} is tax income and X_{t+1} is expenditure. Everything is in real terms. We define the surplus as $S_t = T_t - X_t$ and assume throughout that the gross return R_{t+1} is strictly positive. Note that the debt return R_{t+1} should only be interpreted as a riskless interest rate in the special case in which all government debt is short-term real debt. We allow debt to be risky: the realized return on debt is low if, for example, real yields rise, or if there is a sudden unexpected inflation or explicit default.

As a first step toward a simple benchmark, let us imagine that conditional expectations of growth in tax, spending, and the debt are all equal to some constant, G , and that the conditionally expected return on debt equals R . Equation (1) then implies

$$R = \mathbb{E}_t \frac{V_{t+1}}{V_t} + \mathbb{E}_t \frac{T_{t+1}}{T_t} \frac{T_t}{V_t} - \mathbb{E}_t \frac{X_{t+1}}{X_t} \frac{X_t}{V_t} = G \left(1 + \frac{S_t}{V_t} \right). \quad (2)$$

It follows that the surplus-debt ratio is a constant:

$$\log \left(1 + \frac{S_t}{V_t} \right) = \log R - \log G. \quad (3)$$

We write surplus-debt in this form for comparability with the more general analysis below. When $R > G$, the government must run primary surpluses to pay off its debt. By contrast, if $R \leq G$ the government need not run surpluses: even an unexpected increase in debt—for example, to fight a war—never needs to be paid off. In this case, the value of the debt reflects the presence of a rational bubble. In our more general analysis of Section 2, we will rule out this possibility a priori.

Equation (3) exhibits the surplus-debt ratio as a natural quantity of interest, analogous to the dividend-price ratio in the Gordon growth model. Figure 2 shows the evolution of the surplus-debt ratio, S_t/V_t , in the US from 1947 to 2020. As the surplus can take negative values, we plot the series on a linear scale. (We provide a detailed description of our data sources in Appendix A.1.1, and summary statistics are provided in Table A.1.) Although surplus-debt is not constant as it would be in a Gordon-growth-type model, it does appear to be stationary, and this impression is confirmed by an augmented Dickey–Fuller (ADF) test, reported in Table A.2, which rejects the presence of a unit root at the 99% confidence level. By contrast, an ADF test does not reject the presence of a unit root in the surplus-GDP ratio, S_t/Y_t .

Under the weaker assumption that *unconditional* expected tax growth, spending growth and debt growth are all equal to G , and that the *unconditional* expected return on debt equals R , the argument that led to (3) implies that

$$\mathbb{E} \log \left(1 + \frac{S_t}{V_t} \right) = \log R - \log G. \quad (4)$$

An uncomfortable feature of the post-war data is that the time-series average of the surplus-debt ratio is negative over the sample period, as illustrated in Figure 2. If we believe that this sample average is an accurate measure of the true population average, we are forced to conclude that the value of the debt reflects the presence of a rational bubble. We rule out this possibility by imposing a positive population mean $\mathbb{E} \log(1 + S_t/V_t) > 0$ in our empirical work.

The left panel of Figure 3 breaks the primary surplus $S_t = T_t - X_t$ into its constituent parts, plotting the tax-debt and spending-debt ratios separately. Again, the impression

Figure 2: The surplus-debt ratio is stationary in postwar data. Linear scale. NIPA data, 1947–2020.





Figure 3: The spending-debt and tax-debt ratios are nonstationary. The spending-GDP ratio is also nonstationary, but the tax-GDP ratio is stationary.



which emerges from these figures is confirmed by ADF tests: neither $\tau v_t = \log T_t/V_t$ nor $xv_t = \log X_t/V_t$ is stationary, despite the fact that the surplus-debt ratio is stationary. These facts place important constraints on how we set up our analysis.

The right panel of Figure 3 plots tax-GDP and spending-GDP over time. By now it may come as no surprise that spending-GDP is not stationary. But tax-GDP *is* stationary. (We report ADF tests in Table A.2.) This important empirical fact supplies us with another stationary variable to take into account when we analyze fiscal dynamics.

2 A framework for fiscal analysis

The simple benchmark (3) is unrealistic in various important ways: for one thing, it implies that surplus cannot switch sign. To set up an empirically useful framework, we will have to account for the fact that expected tax growth, spending growth, debt returns, and so on, vary over time. We now present a general framework that does so.

To make a start, rewrite equation (1) as

$$R_{t+1} = \frac{V_{t+1}}{V_t} \left(1 + \frac{S_{t+1}}{V_{t+1}} \right). \quad (5)$$

Taking logs of (5), we have

$$r_{t+1} = \Delta v_{t+1} + \log \left(1 + \frac{S_{t+1}}{V_{t+1}} \right). \quad (6)$$

The measure of the surplus-debt ratio that appears on the right-hand side of (6) echoes the dividend-price ratio measure, $\log(1 + D_{t+1}/P_{t+1})$, used by [Gao and Martin \(2021\)](#). It allows surplus to go negative; moreover, the measure is in natural units, in the sense that $\log(1 + S_{t+1}/V_{t+1})$ is approximately equal to S_{t+1}/V_{t+1} if surplus-debt is small. It can be written in terms of the log tax-debt ratio, $\tau v_t = \log(T_t/V_t)$, and the log spending-debt ratio, $xv_t = \log(X_t/V_t)$, as

$$\log \left(1 + \frac{S_{t+1}}{V_{t+1}} \right) = \log (1 + e^{\tau v_{t+1}} - e^{xv_{t+1}}). \quad (7)$$

2.1 A loglinear measure of the fiscal position

To construct a tractable measure of the fiscal position, we linearize equation (7) in τv_t and xv_t . In doing so, we exploit the fact that while neither tax-debt, τv_t , nor spending-debt, xv_t , is stationary over the postwar sample (as discussed in the previous section and shown in [Figure 3](#)) they do appear to be cointegrated. [Table A.3](#) reports results of Johansen tests that indicate a cointegrating relationship: that is, $\tau v_t - \beta xv_t$ is stationary for some constant β . The estimates of β are close to but slightly less than one. Likewise, $\log(1 + S_t/V_t)$ is stationary, as discussed in the previous section. We use these facts to guide our linearization.

Specifically, linearizing $\log(1 + e^{\tau v_{t+1}} - e^{xv_{t+1}})$ around $(\tau v_{t+1}, xv_{t+1}) = (\log a, \log b)$,

where a and b are both positive, we have

$$\log(1 + e^{\tau v_{t+1}} - e^{xv_{t+1}}) = k + \frac{1}{1+a-b} (a\tau v_{t+1} - bxv_{t+1}) \quad (8)$$

up to higher order terms in τv_{t+1} and xv_{t+1} , where

$$k = \log(1+a-b) + \frac{b \log b - a \log a}{1+a-b}. \quad (9)$$

We choose a and b to satisfy two conditions.

First, we want to linearize around the unconditional mean of $\log(1+S_{t+1}/V_{t+1})$: that is, we require

$$\log(1+a-b) = \mathbb{E} \log\left(1 + \frac{S_t}{V_t}\right). \quad (10)$$

We write $\mathbb{E} \log(1+S_t/V_t) = -\log \rho$. Assuming $\mathbb{E} \log(1+S_t/V_t)$ is well defined, this implies that $\rho > 0$. If the surplus-debt ratio were constant, we would have $1/\rho = 1 + S_t/V_t$, so in the context of the discussion following equation (3) $\rho < 1$ is equivalent to $R > G$. Henceforth we will impose this constraint, which is equivalent to requiring that the government must ultimately pay off its debt. Thus, to summarize, ρ must lie between zero and one. With this new notation, equation (10) becomes

$$1+a-b = \frac{1}{\rho}. \quad (11)$$

Second, we want the right-hand side of (8) to be stationary, as the left-hand side is. Given the cointegrating relationship between τv_t and xv_t , this requires that

$$\frac{b}{a} = \beta. \quad (12)$$

Equations (11) and (12) jointly determine a and b in terms of β and ρ . We have

$$a = \frac{1}{1-\beta} \frac{1-\rho}{\rho} \quad \text{and} \quad b = \frac{\beta}{1-\beta} \frac{1-\rho}{\rho}. \quad (13)$$

As a and b are positive, and $0 < \rho < 1$ by assumption, we must have $0 < \beta < 1$.

Plugging these choices of a and b back into (8), we have our linearization

$$\log\left(1 + \frac{S_{t+1}}{V_{t+1}}\right) = \log(1 + e^{\tau v_{t+1}} - e^{xv_{t+1}}) = \underbrace{k + \frac{1-\rho}{1-\beta}(\tau v_{t+1} - \beta x v_{t+1})}_{sv_{t+1}}, \quad (14)$$

where the first equality follows from the definition of surplus. Here k is as in equation (9) with a and b given by (13).

We will refer to the quantity on the far right-hand side of equation (14) as sv_{t+1} and will use it as our measure of the government's fiscal position. That is, we define

$$sv_t = k + \frac{1-\rho}{1-\beta}(\tau v_t - \beta x v_t) \quad (15)$$

where

$$k = \rho \log \rho + (1-\rho) \log \frac{1-\rho}{1-\beta} - \frac{1-\rho}{1-\beta} \beta \log \beta, \quad (16)$$

so that sv_t is a linearization of $\log(1 + S_t/V_t)$ that, like $\log(1 + S_t/V_t)$, is stationary.

The two quantities differ in one important way, however. As the level of debt rises with surplus held fixed, sv_t declines whether the surplus is positive or negative. This follows from the definition (15), given that ρ and β lie between zero and one. Similarly, sv_t declines when tax falls or when spending rises with other quantities held fixed. Thus we can think of sv_t as a measure of the fiscal position: it is high when the government is in a strong fiscal position, and low when the government is in a weak fiscal position. By contrast, the more conventional measures S_t/V_t and $\log(1 + S_t/V_t)$ are harder to interpret: as the debt grows, they go *down* if surplus is positive, but *up* if the surplus is negative.

2.2 A present value model for the fiscal position

The linearity of sv_t allows us to relate it to fundamentals in a linear present value framework. Inserting the linearization (14) into the exact identity (6), we have

$$r_{t+1} = \Delta v_{t+1} + sv_{t+1}. \quad (17)$$

Taking differences of (15) and rearranging, we have

$$(1 - \rho)\Delta v_{t+1} = \frac{1 - \rho}{1 - \beta}\Delta\tau_{t+1} - \beta\frac{1 - \rho}{1 - \beta}\Delta x_{t+1} - \Delta sv_{t+1}. \quad (18)$$

We use (18) to eliminate Δv_{t+1} from identity (17), giving (after some rearrangement)

$$sv_t = (1 - \rho) \left[r_{t+1} - \frac{1}{1 - \beta}\Delta\tau_{t+1} + \frac{\beta}{1 - \beta}\Delta x_{t+1} \right] + \rho sv_{t+1}. \quad (19)$$

We now solve forward in the usual way, to find that

$$sv_t = (1 - \rho) \sum_{j=0}^{T-1} \rho^j \left[r_{t+1+j} - \frac{1}{1 - \beta}\Delta\tau_{t+1+j} + \frac{\beta}{1 - \beta}\Delta x_{t+1+j} \right] + \rho^T sv_{t+T}. \quad (20)$$

In the limit as $T \rightarrow \infty$, since stationarity implies that sv_t is not explosive so $\lim_{T \rightarrow \infty} \rho^T sv_{t+T} = 0$, we have the dynamic generalization of the static present value formula (3) that we were seeking:

$$sv_t = (1 - \rho) \sum_{j=0}^{\infty} \rho^j \left[r_{t+1+j} - \frac{1}{1 - \beta}\Delta\tau_{t+1+j} + \frac{\beta}{1 - \beta}\Delta x_{t+1+j} \right]. \quad (21)$$

In other words, if the government is in a strong fiscal position (sv_t is high), then either the holders of government debt will earn high returns, or taxes will grow slowly, or government expenditure will grow rapidly, or some combination of the above will occur, at some point in the future. This relationship is a loglinear approximation to an accounting identity, so it holds ex post. It also holds ex ante for rational expectations, and indeed for any subjective expectations that respect identities.

Two further points about equation (21) are worth noting. First, the discounting with discount factor $\rho < 1$ implies that the longer the various sources of fiscal adjustment are delayed, the larger they must ultimately be. Second, the multiplication of tax growth by $1/(1 - \beta)$ and of spending growth by $\beta/(1 - \beta)$ —which are large numbers given that β is close to one—reflects the fact that when the average primary surplus is small relative to the average levels of tax revenue and government expenditure, small percentage changes in either taxes or spending have large proportional effects on the primary surplus and hence on our measure of the fiscal position.

3 Empirical results in US data

3.1 Parameter calibration

As $\mathbb{E} \log(1 + S_t/V_t) = -\log \rho$, we could in principle use the sample mean of $\log(1 + S_t/V_t)$ to pin down ρ , given a sufficiently long sample. In postwar data, however, the average surplus-debt ratio is negative, so this procedure would set ρ greater than one, and would bake in an “ $R < G$ ” assumption. In order to impose a restriction that the government must pay off its debt, we therefore set ρ less than one as an a priori choice.

In our baseline analysis, we set $\rho = 0.999$ so that the implied unconditional expectation of $\log(1 + S_t/V_t)$ is not too far from its sample mean in postwar data. For consistency, we therefore de-mean returns, tax growth and the fiscal position in our VAR estimation using “theory means” $\mathbb{E} r_t = 0.030$, $\mathbb{E} \Delta \tau_t = 0.029$, and $\mathbb{E} sv_t = 0.001$, and estimate with zero intercepts. (That is, we set mean of tax or spending growth equal to the sample mean of tax growth and the mean log return to $0.03 = 0.029 + 0.001$ for consistency with our assumption that $\rho = 0.999$.) In the case of the log tax-GDP ratio, we de-mean using the sample mean over the period 1947–2020, $\mathbb{E} \tau y_{t+1} = -1.787$.


We then choose β so that sv_t optimally approximates $\log(1 + S_t/V_t)$ in a least-squares sense. That is, β is chosen to solve the problem

$$\min_{\beta} \sum_t \left(\log(1 + S_t/V_t) - \underbrace{\left[k + \frac{1-\rho}{1-\beta} (\tau v_t - \beta x v_t) \right]}_{sv_t} \right)^2, \quad (22)$$

where k is given in equation (16). With $\rho = 0.999$, this procedure sets $\beta = 0.997$. The resulting time series of sv_t is shown in Figure 4, together with $\log(1 + S_t/V_t)$ which it approximates.

In Appendix A.2.1, we conduct a sensitivity analysis by setting $\rho = 0.99$, and show that this choice has little effect on our conclusions. Finally, let us emphasize that our approach allows for the possibility that there are extended periods in which the *conditional* expectation of $\log(1 + S_t/V_t)$ is negative; what we want to rule out is the possibility that the mean is negative unconditionally, for all time.

Figure 4: sv_t and $\log(1 + S_t/V_t)$.



3.2 A VAR system

The approximate identity (21) relates our measure of the fiscal position, sv_t , to future debt returns, tax growth, and spending growth. It formalizes the fact that when the government is in a weak fiscal position (i.e., sv_t is low) we must subsequently have some combination of low debt returns, high tax growth, and low spending growth.

To determine which of these channels is most important empirically, we estimate a VAR for the variables r_t , $\Delta\tau_t$, sv_t , and the log tax-GDP ratio, $\tau y_t = \log T_t/Y_t$. A Johansen test, reported in Table A.4, confirms that the VAR is well specified.

By including τy_t , we ensure that the VAR takes into account the stationary relationship between tax and output. We do not include Δx_t as it is mechanically related to the first three included variables via the approximate identity (19). Indeed, we treat the identity as holding exactly, so that we can infer Δx_{t+j} from (19),

$$\frac{\beta}{1-\beta}\Delta x_{t+j} = \frac{sv_{t+j-1} - \rho sv_{t+j}}{1-\rho} - r_{t+j} + \frac{1}{1-\beta}\Delta\tau_{t+j}. \quad (23)$$

Similarly, we do not include GDP growth Δy_t because it is mechanically related to included variables by the identity

$$\Delta y_{t+j} = \Delta\tau_{t+j} - \Delta\tau y_{t+j}. \quad (24)$$

Table 1: VAR coefficient estimates for $(r_t, \Delta\tau_t, sv_t, \tau y_t)$, US data 1947–2020.

OLS standard errors are reported in square brackets.

	r_{t+1}	$\Delta\tau_{t+1}$	sv_{t+1}	τy_{t+1}	Δy_{t+1}
r_t	0.060 [0.110]	-0.259 [0.115]	-0.220 [0.080]	-0.324 [0.103]	0.065 [0.053]
$\Delta\tau_t$	-0.072 [0.092]	0.355 [0.096]	-0.037 [0.066]	0.367 [0.085]	-0.011 [0.044]
sv_t	-0.102 [0.127]	-0.136 [0.133]	0.763 [0.092]	-0.221 [0.119]	0.085 [0.062]
τy_t	0.287 [0.091]	-0.419 [0.095]	0.003 [0.066]	0.676 [0.084]	-0.095 [0.044]
R^2	17.11%	40.58%	60.87%	63.58%	8.13%

The estimated VAR is shown in the first four columns of Table 1. The more persistent variables—tax growth, the fiscal position, and tax-GDP ratio—are relatively predictable, with R^2 between 40% and 64%, and are each strongly predicted by their lags. Restricting to coefficients with t -statistics above three for the purposes of discussion, we see that a high tax-GDP ratio τy_t predicts high returns on debt and low tax growth; high tax growth predicts a high future tax-GDP ratio; and high debt returns predict a lower future tax-GDP ratio.

The rightmost column of Table 1 shows the implied VAR forecast of GDP growth calculated using equation (24). The explanatory power of the model for GDP growth is modest at 8%, but the tax-output ratio does enter significantly and predicts slow GDP growth.

3.3 Decomposing the variance of the fiscal position

We can use the VAR to understand what fluctuations in the fiscal position, sv_t , imply about the subsequent evolution of debt returns, tax growth, and spending growth. Stacking the variables into a vector $\mathbf{z}_{t+1} = (r_{t+1}, \Delta\tau_{t+1}, sv_{t+1}, \tau y_{t+1})'$ and arranging the entries of Table 1 into a coefficient matrix \mathbf{A} , we have $\mathbb{E}_t \mathbf{z}_{t+j} = \mathbf{A}^j \mathbf{z}_t$. If we write \mathbf{e}_n for a vector with one in the n th entry and zeroes elsewhere, we therefore have $\mathbb{E}_t r_{t+j} = \mathbf{e}'_1 \mathbf{A}^j \mathbf{z}_t$,

$\mathbb{E}_t \Delta\tau_{t+j} = \mathbf{e}'_2 \mathbf{A}^j \mathbf{z}_t$, and so on.

We can use identity (20) to derive finite-horizon variance decompositions in the form

$$1 = \frac{\text{cov}(sv_t, (1 - \rho) \sum_{j=0}^{T-1} \rho^j r_{t+1+j})}{\text{var } sv_t} + \frac{\text{cov}(sv_t, -(1 - \rho) \sum_{j=0}^{T-1} \rho^j \frac{1}{1-\beta} \Delta\tau_{t+1+j})}{\text{var } sv_t} + \frac{\text{cov}(sv_t, (1 - \rho) \sum_{j=0}^{T-1} \rho^j \frac{\beta}{1-\beta} \Delta x_{t+1+j})}{\text{var } sv_t} + \frac{\text{cov}(sv_t, \rho^T {}_t sv_{t+T})}{\text{var } sv_t}. \quad (25)$$

This decomposition can be derived by taking the time- t conditional expectations of both sides of (20), computing covariances with sv_t and, finally, scaling by $\text{var } sv_t$ so that the four terms on the right-hand side of (25) add up to 100%. The decomposition tells us the relative contribution of future debt returns, tax growth, spending growth, and persistent variation in the fiscal position to explaining the variability of the fiscal position at any given horizon. As we let the horizon increase, the contribution of the long-horizon future fiscal position declines to zero and we are left with a three-variable infinite-horizon variance decomposition for the fiscal position.

Table 2, Panel A reports the results of this exercise over various different horizons T . In each row of the table, the four entries correspond to the four terms on the right-hand side of (25); notice that we include a minus sign inside the second covariance term on the right-hand side of (25), so that positive entries in the column labelled “tax” indicate a negative covariance between the fiscal position and subsequent tax growth.

Panel B reports bootstrapped 95% confidence intervals for these estimates. Each bootstrap sample is computed by first drawing a new VAR coefficient matrix using the point estimates and the covariance matrix of the estimated coefficients. Using this VAR coefficient matrix, we generate the news series and do the variance decomposition. We repeat this procedure 10,000 times and report the 2.5% and 97.5% quantiles.

At short horizons, variation in sv_t is largely reflected in short-run future sv_t : if the fiscal position is weak this year, it probably will be next year too. At all horizons, there is essentially no relationship between the fiscal position and expected real returns; this contrasts with the evidence that dividend yields do forecast returns on the stock market. As a result, the fiscal position sv_t must in the long run forecast tax growth or spending growth, or both. We find that a poor fiscal position (low sv_t) is associated with high expected tax growth and low expected spending growth over the medium run. Over the long run, though, essentially all of the burden of adjustment falls on spending: a weak fiscal position forecasts spending cuts, not tax growth.

Table 2: A variance decomposition for the fiscal position, sv_t .

Panel A: Variance decomposition for sv_t				
Horizon	return	tax	spending	future sv
1	-0.0%	4.2%	14.5%	82.7%
3	-0.0%	19.4%	32.4%	49.5%
10	-0.1%	3.4%	85.3%	12.7%
30	-0.1%	0.4%	100.9%	0.2%
∞	-0.1%	0.3%	101.2%	0.0%

Panel B: Bootstrap intervals				
Horizon	return	tax	spending	future sv
1	[-0.0%, 0.0%]	[-1.2%, 27.8%]	[3.6%, 43.3%]	[37.7%, 90.8%]
3	[-0.1%, 0.1%]	[0.5%, 36.6%]	[9.4%, 65.9%]	[9.3%, 80.9%]
10	[-0.3%, 0.1%]	[-26.3%, 19.2%]	[52.8%, 101.8%]	[-0.1%, 61.1%]
30	[-0.5%, 0.1%]	[-69.6%, 18.7%]	[82.0%, 146.5%]	[-0.0%, 29.7%]
∞	[-0.7%, 0.1%]	[-108.6%, 18.7%]	[82.8%, 210.5%]	[-0.0%, 0.0%]

3.4 Decomposing the fiscal adjustment to tax and expenditure shocks

As our framework allows us to analyze the behavior of tax and spending separately, we can also ask whether deficits driven by shocks to taxes look different from deficits driven by shocks to spending.

We address this question by using the identity (20) to explore the implications of unexpected shocks to taxes or spending. Applying the “news operator”, $\Delta \mathbb{E}_{t+1} =$

$\mathbb{E}_{t+1} - \mathbb{E}_t$, to both sides of (20) and rearranging, we have

$$\begin{aligned}
\underbrace{\Delta \mathbb{E}_{t+1} \tau_{t+1}}_{\text{short-run tax news}} &= (1 - \beta) \underbrace{\Delta \mathbb{E}_{t+1} \sum_{j=0}^{T-1} \rho^j r_{t+1+j}}_{\text{return news}} - \underbrace{\Delta \mathbb{E}_{t+1} \sum_{j=1}^{T-1} \rho^j \Delta \tau_{t+1+j}}_{\text{long-run tax news}} + \\
&+ \beta \underbrace{\Delta \mathbb{E}_{t+1} \sum_{j=0}^{T-1} \rho^j \Delta x_{t+1+j}}_{\text{spending news}} + \frac{1 - \beta}{1 - \rho} \underbrace{\Delta \mathbb{E}_{t+1} \rho^T s v_{t+T}}_{\text{future fiscal position news}}. \quad (26)
\end{aligned}$$

This identity allows us to trace out the consequences of an unexpected shock to taxes. We refer to such a shock as short-run tax news, $N_{\text{SR tax},t+1} = \Delta \mathbb{E}_{t+1} \tau_{t+1}$. A positive short-run tax shock must be reflected in some combination of (i) news about returns, $N_{\text{return},t+1} = \Delta \mathbb{E}_{t+1} \sum_{j=0}^{T-1} \rho^j r_{t+1+j}$; (ii) news about declines in long-run tax growth, $N_{\text{LR tax},t+1} = \Delta \mathbb{E}_{t+1} \sum_{j=1}^{T-1} \rho^j \Delta \tau_{t+1+j}$; (iii) news about spending growth, $N_{\text{spending},t+1} = \Delta \mathbb{E}_{t+1} \sum_{j=0}^{T-1} \rho^j \Delta x_{t+1+j}$; and/or (iv) news about the future fiscal position, $N_{\text{future sv},t+1} = \Delta \mathbb{E}_{t+1} \rho^T s v_{t+T}$. This last term becomes negligible once the horizon, T , is sufficiently long.

Taking covariances of both sides of (26) with short-run tax news, $N_{\text{SR tax},t+1} = \Delta \mathbb{E}_{t+1} \tau_{t+1}$, and rearranging, we have

$$\begin{aligned}
1 &= \frac{\text{cov}((1 - \beta)N_{\text{return},t+1}, N_{\text{SR tax},t+1})}{\text{var } N_{\text{SR tax},t+1}} + \frac{\text{cov}(-N_{\text{LR tax},t+1}, N_{\text{SR tax},t+1})}{\text{var } N_{\text{SR tax},t+1}} \\
&+ \frac{\text{cov}(\beta N_{\text{spending},t+1}, N_{\text{SR tax},t+1})}{\text{var } N_{\text{SR tax},t+1}} + \frac{\text{cov}\left(\frac{1-\beta}{1-\rho} N_{\text{future sv},t+1}, N_{\text{SR tax},t+1}\right)}{\text{var } N_{\text{SR tax},t+1}}. \quad (27)
\end{aligned}$$

Panel A of Table 3 reports the four terms on the right-hand side of the identity (27) for a range of horizons, T ; the four terms in each row would add up to precisely 100% if our loglinear approximation were exact. Panel B shows bootstrapped 95% confidence intervals calculated as in the variance decomposition reported in Table 2.

In the very short run, at horizon $T = 1$, unexpected declines in tax are associated with unexpected contemporaneous *increases* in spending. This movement is in the “wrong” direction, which exacerbates the shock to the fiscal position (hence the entry greater than 100% in the rightmost column of Panel A). At longer horizons, an unexpected short-run tax cut forecasts rises in long-run tax growth, but as these do not fully offset the effect of the original tax cut it also forecasts declines in long-run spending. That

Table 3: A variance decomposition for short-run tax news.

Panel A: Variance decomposition for short-run tax news				
T	return	LR tax	spending	future sv
1	-0.0%	—	-37.6%	139.2%
3	0.1%	43.8%	-16.3%	74.0%
10	0.0%	75.7%	9.3%	16.5%
30	0.0%	77.2%	24.1%	0.3%
∞	0.0%	77.1%	24.4%	—

Panel B: Bootstrap intervals				
T	return	LR tax	spending	future sv
1	[-0.1%, -0.0%]	[-0.0%, 0.0%]	[-45.9%, -29.5%]	[131.0%, 147.5%]
3	[-0.0%, 0.1%]	[11.2%, 69.5%]	[-61.2%, 23.9%]	[25.2%, 135.2%]
10	[-0.2%, 0.2%]	[43.9%, 97.7%]	[-33.0%, 33.3%]	[-0.4%, 70.0%]
30	[-0.4%, 0.2%]	[16.9%, 98.7%]	[1.8%, 68.5%]	[-0.1%, 29.9%]
∞	[-0.5%, 0.1%]	[-18.3%, 98.6%]	[2.7%, 120.5%]	[-0.0%, 0.0%]

said, we should note that as the confidence intervals are wide, our results are not decisive about the relative importance of tax and spending adjustment.

We can carry out a similar exercise for spending rather than taxes, rewriting the identity (26) as

$$\underbrace{\Delta \mathbb{E}_{t+1} x_{t+1}}_{\text{short-run spending news}} = -\frac{1-\beta}{\beta} \underbrace{\Delta \mathbb{E}_{t+1} \sum_{j=0}^{T-1} \rho^j r_{t+1+j}}_{\text{return news}} + \frac{1}{\beta} \underbrace{\Delta \mathbb{E}_{t+1} \sum_{j=0}^{T-1} \rho^j \Delta \tau_{t+1+j}}_{\text{tax news}} + \underbrace{-\Delta \mathbb{E}_{t+1} \sum_{j=1}^{T-1} \rho^j \Delta x_{t+1+j}}_{\text{long-run spending news}} - \frac{1-\beta}{\beta(1-\rho)} \underbrace{\Delta \mathbb{E}_{t+1} \rho^T s v_{t+T}}_{\text{future fiscal position news}}. \quad (28)$$

We write $N_{\text{tax},t+1}$ for the tax news term that appears on the right-hand side of identity (28). This is the sum of short-run tax news and long-run tax news, as defined in (26): $N_{\text{tax},t+1} = N_{\text{SR tax},t+1} + N_{\text{LR tax},t+1}$. Similarly, we write $N_{\text{SR spending},t+1}$ for short-

Table 4: A variance decomposition for short-run spending news.

Panel A: Variance decomposition for short-run spending news				
T	return	tax	LR spending	future sv
1	-0.0%	-15.9%	—	117.4%
3	-0.1%	-14.3%	37.1%	78.6%
10	-0.1%	-25.3%	107.7%	19.1%
30	-0.1%	-29.9%	131.2%	0.3%
∞	-0.1%	-30.0%	131.6%	—

Panel B: Bootstrap intervals				
T	return	tax	LR spending	future sv
1	[-0.1%, -0.0%]	[-0.0%, 0.0%]	[-0.0%, 0.0%]	[113.7%, 120.8%]
3	[-0.1%, -0.0%]	[-15.6%, 18.7%]	[13.0%, 59.0%]	[50.5%, 110.6%]
10	[-0.4%, -0.0%]	[-38.1%, 9.2%]	[61.0%, 127.3%]	[0.4%, 82.1%]
30	[-0.7%, 0.0%]	[-91.2%, 8.8%]	[107.8%, 182.1%]	[-0.0%, 39.5%]
∞	[-0.9%, 0.0%]	[-140.5%, 8.8%]	[108.7%, 258.6%]	[-0.0%, 0.0%]

run spending news and $N_{\text{LR spending},t+1}$ for long-run spending news, so that $N_{\text{spending},t+1}$ as defined after identity (26) is equal to the sum $N_{\text{SR spending},t+1} + N_{\text{LR spending},t+1}$.

We can now decompose the variance of short-run spending news as the sum of its covariances with news about returns, about tax growth, about long-run spending growth, and about the long-run fiscal position:

$$\begin{aligned}
 1 = & \frac{\text{cov}\left(-\frac{1-\beta}{\beta}N_{\text{return},t+1}, N_{\text{SR spending},t+1}\right)}{\text{var } N_{\text{SR spending},t+1}} + \frac{\text{cov}\left(\frac{1}{\beta}N_{\text{tax},t+1}, N_{\text{SR spending},t+1}\right)}{\text{var } N_{\text{SR spending},t+1}} \\
 & + \frac{\text{cov}\left(-N_{\text{LR spending},t+1}, N_{\text{SR spending},t+1}\right)}{\text{var } N_{\text{SR spending},t+1}} + \frac{\text{cov}\left(-\frac{1-\beta}{\beta(1-\rho)}N_{\text{future sv},t+1}, N_{\text{SR spending},t+1}\right)}{\text{var } N_{\text{SR spending},t+1}}.
 \end{aligned} \tag{29}$$

Panel A of Table 4 reports the four terms on the right-hand side of the identity (29) for a range of horizons, T . Panel B reports the corresponding bootstrapped 95% confidence intervals.

In the very short run, at horizon $T = 1$, unexpected increases in spending are associated with unexpected contemporaneous *decreases* in tax. Again, this movement is in the “wrong” direction, which exacerbates the shock to the fiscal position.

At longer horizons, a deterioration in the fiscal position due to an unexpected rise in short-run spending does not forecast an increase in tax over the long run (as was the case for an unexpected decline in short-run tax) but a *decline*. As a result, a positive spending news shock forecasts a large decline in long-run spending growth that more than offsets the original increase, as indicated by the entry greater than 100% in the column labelled “LR spending.”

We note, finally, that the evidence in Tables 3 and 4 shows that whether the fiscal position worsens due to unexpected declines in tax or unexpected increases in spending, there is almost no association with news about returns, either contemporaneously or in the long run.

3.5 The importance of the tax-GDP ratio

These results depend critically on our inclusion of the stationary variable τy_t in our VAR model. As we will now see, a three-variable VAR in $(r_t, \Delta\tau_t, sv_t)$ —which does not “know” that the log tax-output ratio is stationary—would suggest that variations in sv_t are largely resolved by future tax growth.

Table A.6 reports the results of a VAR that includes returns, r_{t+1} , tax growth, $\Delta\tau_{t+1}$, and the fiscal position, sv_{t+1} . (See Table A.5 for the associated Johansen test.) The coefficient estimates are consistent with those reported in Table 1, but returns and tax growth are substantially less predictable in an R^2 sense when the tax-GDP ratio is not included.

Table A.7 reports the result of a variance decomposition of the fiscal position that applies the identity (25) using this new VAR system. At the shortest horizon, $T = 1$, the results echo our baseline findings shown in Table 2: variation in the fiscal position is largely reflected in the short-run future fiscal position. But at longer horizons the picture is very different: the VAR suggests that variation in the fiscal position is resolved more through tax than through spending adjustment. This conclusion, which differs sharply from our earlier finding that the fiscal position is entirely resolved by adjustments in spending, reflects the fact that the VAR does not take into account the stationarity of the tax-GDP ratio.

The variance decompositions of tax and spending news also look quite different if we neglect the importance of the tax-GDP ratio. Using the identities (27) and (29) to decompose the variance of unexpected shocks to taxes or to spending, as before, we find results shown in Tables A.8 and A.9. These should be compared to Tables 3 and 4.

When the tax-GDP ratio is not included in the VAR, roughly 30% of variation in short-run tax news is accounted for by adjustments in long-run tax news, and roughly 70% by long-run spending adjustments. This contrasts with our earlier findings, in Table 3, which attributed a larger fraction of the variation to adjustments in long-run taxes. The reason is that following positive tax news (an unexpected rise in taxes), the full system understands that this drives up the tax-GDP ratio; as this ratio is stationary, the full system predicts a greater role for the offsetting decrease in taxes in the long run.

The reduced system attributes the variance of short-run spending news shocks roughly equally to adjustments in tax and spending, though with wide confidence intervals that include zero both for tax and spending. Here there is a sharp contrast with our baseline results, which attribute the variance of spending news (more than) entirely to long-run spending—to the extent that the 95% confidence intervals in Tables 4 and A.9 for long-run spending at horizon $T = \infty$ do not overlap.

3.6 Impulse response functions

We now use the estimated VAR coefficient matrix shown in Table 1 (i.e., the full system including the tax-GDP ratio τy_t) to plot impulse response functions.

Figure 5 shows how the system evolves following a debt-financed increase in spending (black lines), or a debt-financed decline in taxes (red lines). These shocks have identical effects on surplus, so analyses that focus directly on surplus without separating into its constituent parts, tax and spending, would impose identical dynamic responses to the two shocks by construction. By contrast, our framework generates very different responses to the two shocks.

The black lines in the panels of Figure 5 indicate the response to a sudden increase in spending (x_t) at time zero that is financed by an increase in debt (v_t) in such a way that the debt return (r_t) is held constant at $t = 0$.

At impact, the increase in debt represents a sharp deterioration in the fiscal position (Panel a). The effect is reasonably persistent, with a half-life on the order of three years. In part this is because subsequent bond returns are expected to be positive. The dotted

Figure 5: Debt-financed spending or tax cut, 4D system, US data.

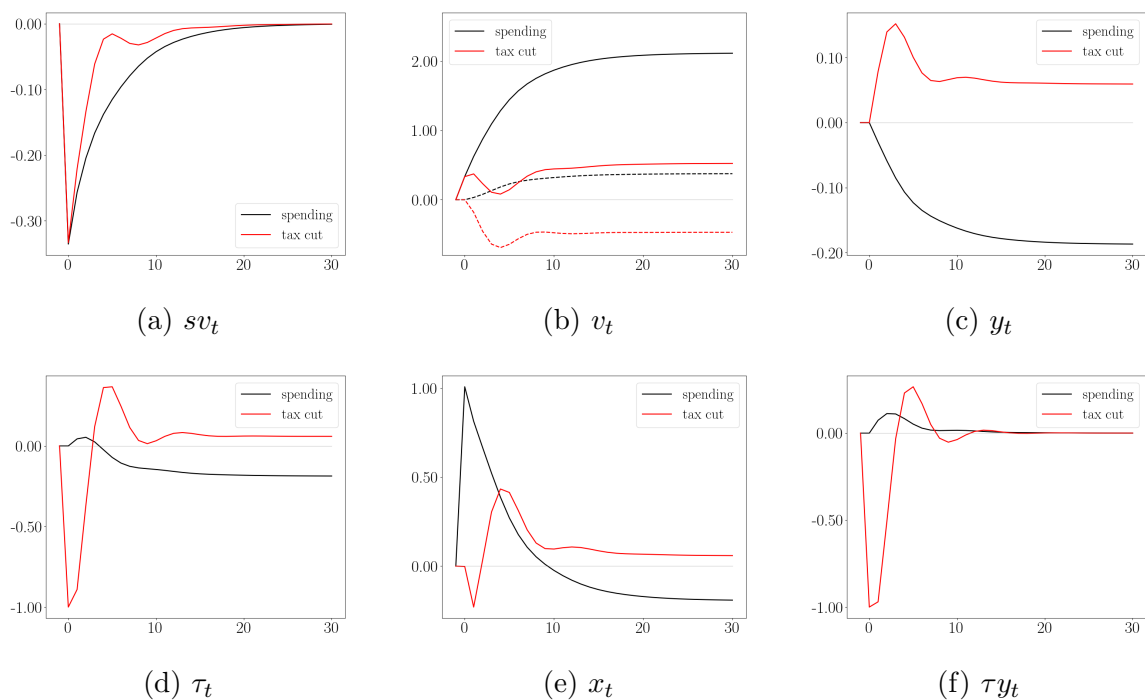



Figure 6: Debt-financed or tax-financed spending, 4D system, US data.



black line in panel b shows the effect of cumulated bond returns on the value of the debt, and the solid black line shows the path of the debt value itself. Both returns and, more importantly, cumulated primary deficits swell the debt over the next 10–15 years. There is a small short-time rise in tax revenue (Panel d) which reverses in the long-run; the reversion of tax-GDP, τy_t , to its mean (Panel f) is achieved by a decline in output (Panel c).

The red lines in Figure 5 trace out the corresponding responses to a decline in taxes (τ_t) that is financed by an increase in debt in such a way that the debt return is held constant, as before.

At impact, there is again a sharp deterioration in the fiscal position (Panel a), but the recovery is faster for a decline in taxes than it was for an increase in spending. There are two contributing factors. First, the dotted red line in panel b shows that bond returns are expected to be negative following the shock, so the value of outstanding debt does not increase as much. Second, a positive shock to GDP, shown in panel c, rapidly restores the level of tax revenue—and, indeed, boosts it above its level prior to the shock—so that the primary surplus recovers more rapidly.


To emphasize the differential impact of the way in which a spending shock is financed, Figure 6 shows the response to a spending shock that is either financed by debt (as in Figure 5) or by taxes in such a way that the primary surplus is unaffected. The former case is shown as a black line and is identical to the black line in Figure 5. The latter case is shown as a red line. In frameworks that analyze surplus directly, a tax-financed spending shock must have zero impact by construction. In our framework, the red lines in Figure 6 show that the fiscal position is unchanged at impact of a tax-financed spending shock, but the fiscal position deteriorates over time (panel a). There are two reasons for this deterioration. First, debt returns turn positive (panel b). Second, output declines (panel c), and as the tax-GDP ratio is stationary, taxes and spending are lower in the long run (panels d and e).

4 Empirical results in UK data

We now repeat the analysis for the UK.

To set the scene, Figure 7, Panel a shows that the debt-GDP ratio also appears nonstationary in postwar UK data. This visual impression is confirmed by an ADF test which fails to reject the presence of a unit root in the log debt-GDP ratio (with a

Figure 7: The debt-GDP and surplus-debt ratios in postwar UK data, 1947–2016.



(a) The debt-GDP ratio. Log scale.

(b) The surplus-debt ratio. Linear scale.


p -value of 0.562: see Table A.18 of the Appendix). Panel b of the same figure shows the surplus-debt ratio, for which the ADF test does reject the presence of a unit root (with a p -value of 0.014, also reported in Table A.18).

Figure 8, Panel a plots the tax-debt ratio and spending-debt ratio for the UK on a log scale. Both ratios appear nonstationary, as confirmed by ADF tests reported in Table A.18 with p -values of 0.579 and 0.335, respectively. Figure 8, Panel b plots the tax-GDP ratio and spending-GDP ratio, also on a log scale. The spending-GDP ratio appears nonstationary, as in the US, and the ADF test fails to reject the presence of a unit root (p -value 0.759). The evidence for the tax-GDP ratio is more mixed: the ADF test gives a p -value of 0.114. Given our earlier results for the US, we proceed under the assumption that the tax-GDP ratio is stationary.

As before, we approximate the surplus-debt ratio by sv_t , as in equation (15). The sample mean of the surplus-debt ratio is positive over our sample period, so we set $-\log \rho$ equal to the sample mean of $\log(1 + S_t/V_t)$, which is reported (together with other summary statistics for the UK data) in Table A.17. Having done so, we pick β to minimize (22), as we did for the US. This procedure sets $\rho = 0.958$ and $\beta = 0.944$. Figure 9 shows the resulting series sv_t , together with the surplus-debt ratio, $\log(1 + S_t/V_t)$, which it approximates.

Table 5 shows the estimated coefficients in VAR featuring r_t , $\Delta\tau_t$, sv_t , and τy_t ; we demean all variables by their sample means, $\mathbb{E} r_t = 0.066$, $\mathbb{E} \Delta\tau_t = 0.025$, $\mathbb{E} \tau y_t = -1.194$,

Figure 8: The spending-debt and tax-debt ratios are nonstationary in the UK. The spending-GDP ratio is also nonstationary; the evidence for the tax-GDP ratio is mixed.



(a) Tax- and spending-debt ratios. Log scale. (b) Tax- and spending-GDP ratios. Log scale.

Figure 9: sv_t and $\log(1 + S_t/V_t)$ in the UK.

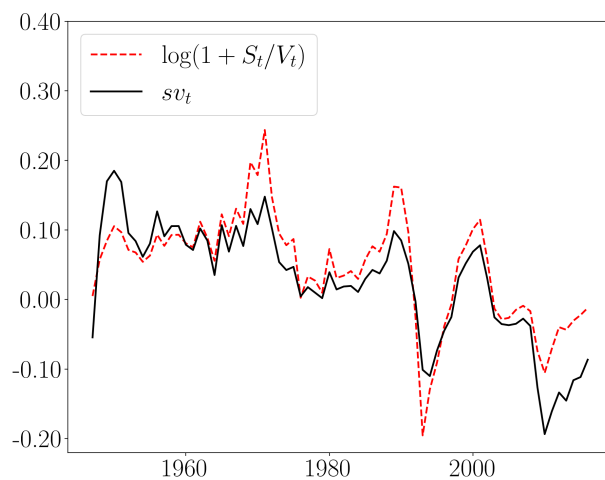


Table 5: VAR coefficient estimates. UK data, 1947–2016.

OLS standard errors are reported in brackets.

	r_{t+1}	$\Delta\tau_{t+1}$	sv_{t+1}	τy_{t+1}	Δy_{t+1}
r_t	−0.228 [0.119]	−0.064 [0.037]	−0.037 [0.047]	−0.131 [0.040]	0.067 [0.029]
$\Delta\tau_t$	0.600 [0.328]	0.446 [0.103]	0.102 [0.129]	0.342 [0.110]	0.104 [0.079]
sv_t	0.024 [0.147]	−0.084 [0.046]	0.873 [0.058]	−0.113 [0.049]	0.029 [0.035]
τy_t	0.166 [0.176]	−0.201 [0.055]	−0.072 [0.069]	0.827 [0.059]	−0.028 [0.042]
R^2	9.04%	30.83%	79.94%	77.84%	14.37%

and $\mathbb{E}sv_t = 0.043$. Johansen test results are reported in Tables [A.19–A.21](#).

In the US VAR, we singled out certain coefficients whose t -statistics were above three: a high tax-GDP ratio τy_t predicted high returns on debt and low tax growth; high tax growth predicted a high future tax-GDP ratio; and high debt returns predicted a lower future tax-GDP ratio. With just one exception, the corresponding coefficients for the UK are also significant, with t -statistics above three and with the same sign. The exception is that high τy_t forecasts high returns on debt, but the estimate is not significantly different from zero. Similarly—though less surprisingly—tax growth, the fiscal position and the tax-GDP ratio were all strongly predicted by their lags in the US data, and we find that the same is true for the UK.

4.1 Variance decompositions

As before, we use the estimated VAR, together with the identity [\(20\)](#), to derive variance decompositions for the UK surplus-debt ratio, sv_t .

The results, which are shown in [Table 6](#), are strikingly similar to the corresponding results reported for the US in [Table 2](#). Very little of the variation in surplus-debt ratio is explained by variation in returns on the debt at any horizon. At long horizons, variation in the surplus-debt ratio is resolved by adjustments in spending. Indeed, the evidence

Table 6: Variance decomposition for sv_t . UK data, 1947–2016.

Panel A: Variance decomposition for sv_t				
Horizon	return	tax	spending	future sv
1	0.2%	-0.1%	14.1%	87.2%
3	-0.1%	3.1%	34.3%	64.1%
10	-0.6%	-11.8%	86.1%	27.7%
30	-0.8%	-30.1%	129.3%	3.0%
∞	-0.8%	-32.3%	134.6%	0.0%

Panel B: Bootstrap intervals				
Horizon	return	tax	spending	future sv
1	[-1.9%, 2.0%]	[-11.5%, 15.6%]	[10.9%, 41.8%]	[57.1%, 89.0%]
3	[-3.7%, 3.2%]	[-20.4%, 25.9%]	[18.9%, 68.8%]	[32.4%, 78.1%]
10	[-7.4%, 6.0%]	[-51.6%, 21.4%]	[52.8%, 125.0%]	[5.3%, 52.6%]
30	[-11.4%, 9.1%]	[-96.3%, 17.9%]	[84.4%, 184.8%]	[-0.0%, 19.3%]
∞	[-12.7%, 10.0%]	[-114.6%, 17.3%]	[90.2%, 215.7%]	[0.0%, 0.0%]

suggests that movements in tax go in the “wrong” direction: a weak fiscal position is associated, in the long run, with *reduced* tax growth. As a result, spending growth must contract even more to resolve the weak fiscal position.

Table 7 uses the identity (27) to understand the correlates of tax news, as Table 3 did for the US, and with similar results. A positive tax news shock forecasts a decline in future tax growth, an increase in spending, and an increased return on the debt: quantitatively, our point estimates imply that about 59 per cent of the variance of tax news shocks is explained by adjustments in long-run tax growth, about 35 per cent by adjustments in long-run spending growth, and about 6 per cent by adjustments in returns.

Table 8 uses the identity (29) to do the analogous exercise for shocks to spending, as Table 4 did for the US. At short horizons, we find that there is some compensating adjustment in both tax and spending in response to an initial spending shock; at long horizons, we find, as in the US, that the burden of adjustment falls more than entirely

Table 7: Variance decompositions for short-run tax news. UK data, 1947–2016.

Panel A: Short-run tax news				
T	return	LR tax	spending	future sv
1	3.1%	—	25.7%	71.1%
3	7.2%	−4.7%	52.8%	44.6%
10	5.9%	61.8%	28.6%	3.6%
30	5.9%	59.5%	34.2%	0.4%
∞	5.9%	59.1%	34.9%	—

Panel B: Bootstrap intervals				
T	return	tax	LR spending	future sv
1	[2.2%, 4.0%]	[-0.0%, 0.0%]	[18.8%, 32.7%]	[63.6%, 78.5%]
3	[3.04%, 11.8%]	[-39.2%, 21.6%]	[15.9%, 95.0%]	[7.8%, 91.7%]
10	[-0.5%, 12.5%]	[23.4%, 94.5%]	[-14.4%, 68.1%]	[-20.9%, 37.7%]
30	[-1.8%, 12.9%]	[15.8%, 103.1%]	[-4.9%, 73.3%]	[-4.7%, 7.3%]
∞	[-2.0%, 13.3%]	[12.3%, 105.1%]	[-7.9%, 79.7%]	[-0.0%, 0.0%]

on long-run spending, so that, for example, a rise in short-run spending is offset by a larger decline in long-run spending.

4.2 The importance of the tax-GDP ratio

As before, we estimate a three-variable VAR in $(r_t, \Delta\tau_t, sv_t)$ and use it to assess the importance of the stationarity of τy_t . Table A.22 reports the results of a VAR that includes returns, r_{t+1} , tax growth, $\Delta\tau_{t+1}$, and the fiscal position, sv_{t+1} . The coefficient estimates are consistent with those reported in Table 5, but tax growth is substantially less predictable when the tax-GDP ratio is not included.

Table A.23 reports a variance decomposition of the fiscal position that applies the identity (25) using this VAR. At longer horizons, the VAR attributes more of the variation in the fiscal position to tax adjustment than was the case in Table 6. This conclusion, which contrasts with our earlier finding that variation in the fiscal position is entirely resolved by adjustments in spending, reflects the fact that the VAR does not take into

Table 8: Variance decompositions for short-run spending news. UK data, 1947–2016.

Panel A: Short-run spending news				
T	return	tax	LR spending	future sv
1	-1.5%	13.0%	—	89.3%
3	-2.9%	22.2%	15.8%	64.8%
10	-2.5%	6.9%	68.2%	28.2%
30	-2.7%	-11.8%	112.3%	3.1%
∞	-2.7%	-14.1%	117.6%	—

Panel B: Bootstrap intervals				
T	return	tax	LR spending	future sv
1	[-2.2%, -0.7%]	[-0.0%, 0.0%]	[-0.0%, 0.0%]	[85.6%, 93.0%]
3	[-3.9%, 0.0%]	[-3.9%, 23.3%]	[0.8%, 31.6%]	[47.6%, 80.7%]
10	[-8.2%, 3.0%]	[-36.5%, 25.7%]	[30.4%, 95.6%]	[7.9%, 53.4%]
30	[-12.4%, 6.4%]	[-83.9%, 21.9%]	[66.4%, 160.8%]	[-0.0%, 18.7%]
∞	[-13.9%, 7.4%]	[-101.9%, 21.4%]	[73.3%, 188.8%]	[-0.0%, 0.0%]

account the stationarity of the tax-GDP ratio. These results are consistent with our findings for the US.

Finally, we confirm that (as in the US) variance decompositions of tax and spending news also look quite different if we neglect the importance of the tax-GDP ratio. Using the identities (27) and (29) to decompose the variance of unexpected shocks to taxes or to spending, as before, we find results shown in Tables A.24 and A.25. These should be compared to the results for the full VAR system that are reported in Tables 7 and 8.

Once again, the results are consistent with the corresponding results for the US. The full system suggests that shocks to short-run tax news are resolved, to a considerable degree, by offsetting adjustments in long-run taxes. The reduced system, which does not appreciate the stationarity of the tax-GDP ratio, attributes more of the adjustment to spending.

Similarly, the reduced system attributes the variance of short-run spending news shocks roughly equally to adjustments in tax and spending. Again there is a sharp

contrast with our baseline results, which attribute the variance of spending news (more than) entirely to long-run spending.

4.3 Impulse response functions

Returning to the full system with VAR coefficient matrix shown in Table 5, we now plot impulse response functions.

Figure 10 shows the impulse responses following a debt-financed increase in spending (black lines), or a debt-financed decline in taxes (red lines). We find, as in the corresponding plots for the US (Figure 5), quite different responses to these two different ways in which the fiscal position can deteriorate. In some respects the results are similar to those we found for the US. Spending increases have a more persistent negative effect on the fiscal position, and a larger positive effect on the size of the debt, than do tax declines; and spending increases forecast long-run declines in tax, spending, and GDP. In the case of the UK, however, the point estimates suggest that debt-financed reductions in tax also forecast long-run declines in tax, spending, and GDP: we found that the opposite was true in the US.

Figure 11 shows the responses to a given spending shock that is financed either with debt (as in Figure 10) or with taxes. A tax-financed spending shock forecasts a deteriorating fiscal position (panel a) for the same two reasons as in the US: debt returns turn positive (panel b), and GDP declines over the medium to long run (panel c), so that both taxes and spending decline in the long run due to the stationarity of the tax-GDP ratio.

5 Conclusion

Conventional tests do not reject the presence of a unit root in the debt-GDP ratio in US postwar data. We have presented a framework for fiscal analysis that takes this uncomfortable fact into account by making the surplus-debt ratio—which does appear to be stationary in postwar data—the central object of interest.

Our framework allows us to analyze the contributions of taxes and spending to surplus separately, and so to draw a distinction between, say, declines in tax revenue and increases in government expenditure. There are good economic reasons to think that these two variables might not have symmetrical properties: spending might exhibit occa-

Figure 10: Debt-financed spending or tax cut, 4D system, UK data.

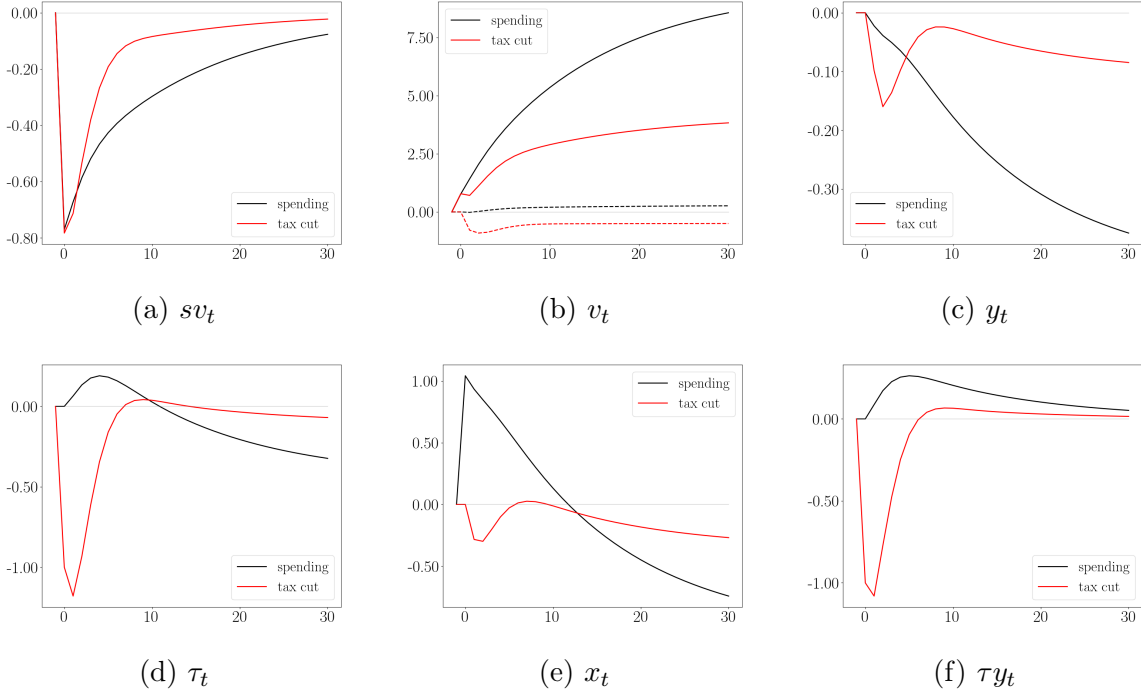
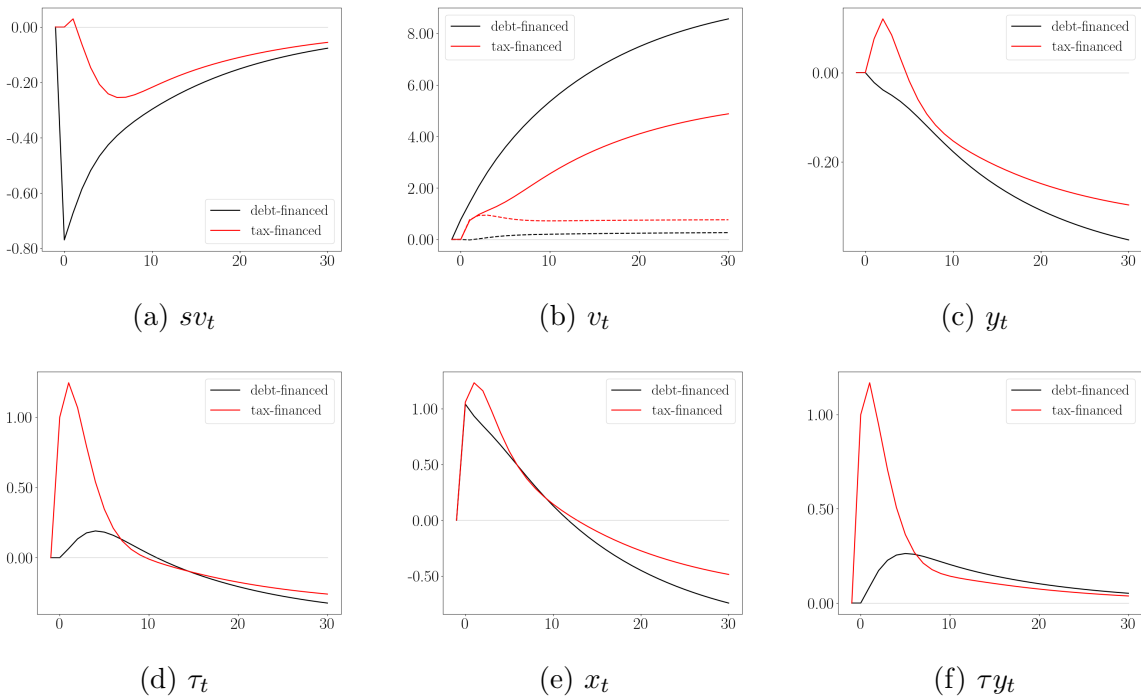


Figure 11: Debt-financed or tax-financed spending, 4D system, UK data.



sional spikes at times of war, for example, whereas we might expect tax revenue to evolve relatively smoothly over time. Concretely, we find that despite the nonstationarity of the surplus-GDP ratio and the expenditure-GDP ratio, the tax-GDP ratio does appear to be stationary, a fact that has important implications for our empirical findings.

We organize our empirical work by deriving a loglinear approximation to the surplus-debt ratio that summarizes the fiscal position of the government. Our key identity relates the fiscal position to future returns on government debt and to future tax and spending growth rates, just as the identities derived by [Campbell and Shiller \(1988\)](#) relate the dividend yield on a security to that security's future returns and dividend growth rates. A weak fiscal position must be followed by some combination of low long-run returns on government debt, high long-run tax growth, and low long-run spending growth.

We use this identity to interpret variation in the fiscal position over time in postwar data from the US and the UK. The fiscal position has almost no forecasting power for future returns; instead, it forecasts adjustment in the primary surplus. More specifically we find that in the long run the burden of adjustment falls essentially entirely on *spending*, with a weak fiscal position predicting long-run declines in spending. These findings differ sharply from those reported in the literature that carries out variance decompositions for stock market returns, following [Campbell \(1991\)](#), and where it is generally argued that valuation ratios have more forecasting power for returns than for cashflow growth.

The comparative importance of cashflow growth reflects, in part, the simple fact that while surplus—tax minus expenditure—is a relatively small number, tax and expenditure are *large* numbers. Thus, say, a 1% change in the level of spending can have a very large proportional impact on the surplus. Meanwhile, the limited role for tax by comparison with spending reflects the fact that taxes are linked to GDP via stationarity of the tax-GDP ratio; and fiscal variables do not strongly predict GDP growth.

We also use our identity to analyze the fiscal adjustment to tax and spending shocks. Again we find that debt returns, both unexpected returns at the time the shocks occur and subsequent predictable returns, play almost no role in fiscal adjustment. Instead, mean-reverting tax and spending growth satisfy the government's intertemporal budget constraint allowing debt value to remain stable. While our framework does not allow us to say which variables are exogenous and which are endogenous, this pattern does tell us that if, as the fiscal theory of the price level asserts, debt value is endogenous, postwar governments in the US and the UK have chosen fiscal policies that avoid large predictable

or unpredictable returns to debtholders. It is possible, perhaps even probable, that our framework would attribute a more significant role to debt returns in countries that have experienced turbulent macroeconomic crises. A priority for future research should be to apply our analysis to other countries where data are available on the market value (as opposed to the face value) of the public debt.

References

- Alberto Alesina, Carlo Favero, and Francesco Giavazzi. *Austerity: When it Works and when it Doesn't*. Princeton University Press, 2020.
- Robert J. Barro. On the determination of the public debt. *Journal of Political Economy*, 87(5):940–971, 1979.
- Antje Berndt, Hanno Lustig, and Şevin Yeltekin. How does the US government finance fiscal shocks? *American Economic Journal: Macroeconomics*, 4(1):69–104, 2012.
- Henning Bohn. The behavior of US public debt and deficits. *Quarterly Journal of Economics*, 113(3):949–963, 1998.
- John Y. Campbell. A variance decomposition for stock returns. *Economic Journal*, 101(405):157–179, 1991.
- John Y. Campbell and Pierre Perron. Pitfalls and opportunities: What macroeconomists should know about unit roots. *NBER Macroeconomics Annual*, 6:141–201, 1991.
- John Y. Campbell and Robert J. Shiller. Cointegration and tests of present value models. *Journal of Political Economy*, 95(5):1062–1088, 1987.
- John Y. Campbell and Robert J. Shiller. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1(3):195–228, 1988.
- John Y. Campbell, Can Gao, and Ian W. R. Martin. Online appendix to Debt and Deficits: Fiscal Analysis with Stationary Ratios. *Working paper*, 2023.
- John H. Cochrane. Long-term debt and optimal policy in the fiscal theory of the price level. *Econometrica*, 69(1):69–116, 2001.
- John H. Cochrane. The fiscal roots of inflation. *Review of Economic Dynamics*, 45:22–40, 2022.
- John H. Cochrane. *The Fiscal Theory of the Price Level*. Princeton University Press, 2023.

- Martin Ellison and Andrew Scott. Managing the UK national debt 1694–2018. *American Economic Journal: Macroeconomics*, 12(3):227–57, 2020.
- Can Gao and Ian W. R. Martin. Volatility, valuation ratios, and bubbles: An empirical measure of market sentiment. *Journal of Finance*, 76(6):3211–3254, 2021.
- Chryssi Giannitsarou, Andrew Scott, and Eric M. Leeper. Inflation implications of rising government debt. In *NBER International Seminar on Macroeconomics*, pages 393–439. 2006.
- Robin Greenwood, Samuel G. Hanson, and Jeremy C. Stein. A comparative-advantage approach to government debt maturity. *Journal of Finance*, 70(4):1683–1722, 2015.
- George J. Hall and Thomas J. Sargent. Debt and taxes in eight US wars and two insurrections. *The Handbook of Historical Economics*, pages 825–880, 2021.
- Zhengyang Jiang, Hanno Lustig, Stijn Van Nieuwerburgh, and Mindy Z. Xiaolan. The US public debt valuation puzzle. *NBER Working Paper No. 26583*, 2021a.
- Zhengyang Jiang, Hanno Lustig, Stijn Van Nieuwerburgh, and Mindy Z. Xiaolan. What Drives Variation in the US Debt/Output Ratio? The Dogs that Didn’t Bark. *NBER Working Paper No. 29351*, 2021b.
- Louis Johnston and Samuel H. Williamson. What was the US GDP then?, 2023. URL <https://www.measuringworth.com/datasets/usgdp/>.
- Arvind Krishnamurthy and Annette Vissing-Jorgensen. The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267, 2012.
- Eric M. Leeper. Equilibria under active and passive monetary and fiscal policies. *Journal of Monetary Economics*, 27(1):129–147, 1991.
- Atif R. Mian, Ludwig Straub, and Amir Sufi. A Goldilocks theory of fiscal deficits. *National Bureau of Economic Research*, 2022.
- Carmen M. Reinhart and Kenneth S. Rogoff. Growth in a time of debt. *American Economic Review*, 100(2):573–78, 2010.
- Ricardo Reis. Debt revenue and the sustainability of public debt. *Journal of Economic Perspectives*, 36(4):103–24, 2022.

Thomas Sargent and Neil Wallace. Some unpleasant monetarist arithmetic. *Quarterly Review*, 5(3):1–17, 1981.

Christopher A. Sims. A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy. *Economic Theory*, 4(3):381–399, 1994.

Michael Woodford. Price-level determinacy without control of a monetary aggregate. In *Carnegie-Rochester Conference Series on Public Policy*, volume 43, pages 1–46. Elsevier, 1995.