The evolution of deficits in the United States and implications for debt sustainability*

Jesse Bricker[†], William Peterman[‡], Rebecca Wasyk[§]

March 2026

Abstract

With the debt-to-GDP ratio nearly tripling in the past decades and projected to keep rising in the future, there has been a renewed discussion of debt sustainability in the United States. A key factor in the evolution of debt is the primary deficit. Leveraging the Congressional Budget Office's projections, this paper models the data generation process for the primary balance by incorporating a rich structure of serial correlation in the innovations to this process including allowing for differences by source (legislative, economic and technical reasons). Applying this model to debt accumulation, we demonstrate that including a richer structure for serial correlation leads to a wider distribution of debt outcomes, centered on significantly more debt accumulation. Moreover, we find that a consolidation with a one-time legislative change is less effective at stabilizing the debt-to-GDP ratio than a series of smaller but correlated contractionary legislative innovations consistent with the innovations in the pre-2000 period.

*We thank colleagues at the Board—including Laura Feiveson and Nick Turner—for valuable early feedback and comments from Doris Prammer and participants at the XXIII Banca d'Italia Public Finance Workshop and the 8^{th} Interdisciplinary Sovereign Debt Research and Management Conference (DebtCon) at Georgetown University. The analysis and conclusions set forth here are those of the authors and do not indicate concurrence by other members of the research staff, the Board of Governors, or the Federal Reserve System

[†]Federal Reserve Board, jesse.bricker@frb.gov

[‡]Federal Reserve Board, william.b.peterman@frb.gov

[§]Federal Reserve Board, rebecca.d.wasyk@frb.gov

1 Introduction

Over the last thirty years, the ratio of debt held by the public-to-GDP in the United States has tripled, reaching nearly 100%, a level not seen since World War II. Moreover, assuming current law holds, this debt ratio is projected to further increase to over 150% over the next thirty years (CBO (2025b)). Although there is not one definition of debt "sustainability" many consider the current trajectory "unsustainable" (see for example Edelberg et al. (2025) and Dynan and Elmendorf (2025)). One complication with defining sustainability is that the future path of the debt ratio is governed by three pieces—the primary deficit, the effective interest rate on borrowing, and the growth rate of GDP—each of which are uncertain over the future. Thus, recent exercises explore sustainability by treating each of these three determinants as serially independent stochastic objects and then determining the distribution of potential debt outcomes (IMF (2021), Blanchard (2023), European Commission (2023)).[1]

Importantly, though, changes in the path of these determinants need not be independent across time as is often assumed. For example, a legislative change that reduces tax revenues in the current period often changes the law for many years or permanently, leading to a higher path for the primary deficits in both the current and future periods. Alternatively, a weakening in the cyclical position may cause a reduction in tax revenue and an increase in the primary deficit that could be shorter- or longer-lived than the legislative shock. In both cases though, the initial shock is the first in a sequence of changes to the dynamics of the deficit path—or what we will call "innovations" to the future primary deficit path.

Thus, we revisit stochastic debt sustainability analysis, paying attention to how

[1] Often sustainability is defined as a debt ratio that stabilizes over a given time horizon. In other forms, it is the probability that interest payments on public debt do not exceed a given threshold (Furman and Summers (2020)), or a combination of debt stock and flows (Zenios et al. (2021))

and why the primary deficit evolves allowing for serial correlation. In particular, we leverage the semi-annual deficit projections from the Congressional Budget Office (CBO) that project the path of the deficit for the next five years at any given point in time. The revisions to these multi-year projections allow us to distinguish two types of innovations: (i) innovations that occur in the current period's deficit ("real-time") and (ii) innovations to the expected future deficit path as of today ("future path"), and then separately examine whether there is serial correlation both across real-time innovations and between real-time and future innovations. For example, a real-time innovation due to a permanent legislative change will likely not only alter the deficit today but also alter the expected path in the future as of today. Alternatively, correlation in real-time innovations may arise due to correlated economic shocks leading to correlation in real-time innovations. Because these CBO projection revisions are broken out by economic, legislative, and technical reasons, we are able to examine the serial correlation separately for each reason.

We begin by documenting that there are notable changes to the path of the primary deficit in the CBO projections and that these changes do not appear independent across time (Table 1, Figure 3, and Figure 4). The average CBO deficit revisions in the 1986-2019 sample period have pushed up the deficit in the real-time period, and these changes tend to persist across the next several horizons of the future projection (Table 1)—an initial sign of serial correlation in the primary deficit innovations. We find that each of the three main sources for innovations—economic, legislative, and technical—contribute with a varying degree of persistence indicating that each should be modeled separately.

Given this potential serial correlation, next we formally model the data generating process for the primary deficit, allowing for the correlation described above: across time in the real-time innovations, and between the real-time innovations and the innovations to the future path in a given projection separately by the type of innovation. We find that the data generating processes for innovations from CBO's projections due to legislative, economic, and technical reasons all exhibit both types of serial correlation to a varying degree.

Next, we appeal to a stochastic debt sustainability (SDSA) framework, simulating dynamic debt accumulation under our estimated primary balance data generating process. The serially-correlated data generating process described in our model tends to lead to considerable debt accumulation. Beginning with CBO's current projection for the deficit, simulating the stochastic process for innovations going forward we find that there is essentially no mass on outcomes that would be consistent with a sustainable path for debt. The debt ratio increases by at least one-hundred percentage points in more than half of the simulations, and by twenty five percentage points in almost 90% of our projections. Importantly, CBO projects that debt will only increase by about twenty five percentage points indicating that the stochastic process for innovations is likely to lead to a less sustainable path than in CBO's projection.

To better understand the role of serial correlation, we compare the outcomes above to a dynamic debt accumulation process with a simpler model used in other SDSAs: a stochastic process for the primary deficits using an AR(1) process. Allowing the different types of serial correlation in our model leads, on average, to a debt-to-GDP ratio that is fifty percentage points larger than the simpler model after twenty years and a much more dispersed distribution of outcomes. Importantly, we find that allowing for real-time innovations to be correlated with future path innovations is critical for the larger mean and variance of the distribution of outcomes. Moreover, we find that the debt accumulation in the full stochastic exercise is more so due to legislative innovations than due to technical or economic innovations (Figure 8).

As our stochastic simulations point to almost all of the mass of the distribution being inconsistent with a sustainable debt path, we next investigate what type of changes would be necessary to lead to sustainability. Overall, we find that a one-time large innovation consistent with a fiscal contraction can temporarily stabilize the debt-to-GDP ratio for a few years. However, debt will tend to resume its upward trajectory if legislative innovations continue to be, on average, expansionary as they have been, on average, over the last forty years. In contrast, we find that if the underlying distribution of legislative innovations shifts to being consistent with the pre-2000s era, which, on average, was a period of debt consolidation, then the debt path tends to eventually stabilize. Thus, our simulations point to a change in the underlying distribution of innovations as opposed to consolidation shock being necessary to stabilize debt.

Our work is related to the general stochastic debt sustainability analysis. However, most of these studies tend to assume that innovations to the primary deficit are independently distributed ((e.g. Darvas et al. (2025) Blanchard (2023), and European Commission (2023))) or allow for limited serial correlation (e.g. IMF (2021) and Darvas et al. (2025)), as opposed to our specification which allows for different types of persistence across time that varies by the source of the innovations. We find that allowing for this richer specification leads to a large difference in the distribution of outcomes indicating a higher variance but a lower probability of a sustainable outcome.

Our work is also related to more recent work that incorporates another specific types of persistent by including the feedback loop between current legislative policy changes and future deficit projections (Auerbach and Yagan (2025)). Prior to 2001, the expectation of higher future deficits tended to lead to deficit-reducing legislative changes; since 2001, though, this feedback loop has disappeared. Similar to this paper, Auerbach and Yagan (2025) model future deficits to project future debt ratios. But in this paper, we model the six separate components of deficit revisions: through economic, technical, and legislative changes to both outlays and revenues. This allows us to examine the differing persistence in these innovations from each source and also have more specific counterfactual exercises.

Finally, somewhat related to our work is the prior work on describing the CBO revisions, including a general evaluation of the CBO forecasts in CBO (2024). Past work has described serial correlation in the CBO revenue forecasts—that is, whether the forecast revision for a given fiscal year is correlated with the subsequent forecast revision for that same fiscal year (Auerbach (1999)). This would answer the question of whether each CBO forecast fully uses all available information, or whether there is some hedging in the forecast that eventually goes the full way.

This paper proceeds as follows. Section 2 describes the historical trends in debt accumulation and deficits in the United States. Section 3 includes a description of the CBO projections and revisions data that we draw from, and a model that describes the time series of those data. Section 4 includes the debt accumulation projections, and the final section concludes.

2 Deficits, Debt, and Sustainability Models

This section formally describes the components that determine how the debt-to-GDP ratio evolves over time. Additionally, it describes how to think about debt sustainability given that the factors explaining the dynamics of debt are not deterministic.

2.1 Debt Accumulation

Fiscal policy makers set the level of government revenues and outlays with a consideration for economic growth, financial stability, and social goals (Furman and Summers (2020)). There are important intertemporal linkages between current fiscal policy and the future: any shortfall between revenues collected and government outlays is financed through deficits and adds to the stock of government debt that future governments will inherit.[2] In order to understand the dynamics of debt, it is useful to start with the basic debt accumulation equation:

$$\Delta b_t = p_t + \frac{r_t - g_t}{1 + g_t} b_{t-1} \tag{1}$$

where the debt ratio (b_t) is the ratio of debt held by the public-to-GDP, and whose dynamics (Δb_t) depend on the current average real interest rate on government debt (r_t), the real growth rate of the economy (g_t), last period's debt ratio (b_{t-1}), and the current primary deficit (p_t)—the deficit, excluding net interest payments on the debt.

The first term of (1) implies that a current deficit serves to increase the level of the debt ratio one-for-one in the current period. Assuming $r_t = g_t$, without any other actions, this means an increase in the primary deficit at time t will permanently raise debt by the equivalent amount. However, there are other ways in which this might transmit into a different path for debt after t. For example, an increase in the primary deficit at time t may signal a change in the primary deficit in future periods, implying a change in the slope of debt over time.

Additionally, r may not equal g, which has implications for the path of the debt. For example, when $r_t < g_t$, the second term in (1) is negative which allows the debt ratio (b_t) to fall. This "growth dividend" leads the debt ratio to increase less than one-to-one with any deficit and this downward pressure on debt grows with the size of debt.[3] But when $r_t > g_t$, the opposite occurs and the debt ratio becomes explosive with deficits.

[2]However, setting outlays equal to revenues in every period is often not optimal policy (Furman and Summers (2019); Blanchard (2023)). For example, a deficit financed increase in spending may provide a countercyclical stabilizing force. Even though deficit financed spending can be an important component of countercyclical fiscal policy that ameliorates the business cycle costs, there is an understanding that there might be a limit on the amount of debt a country can issue.

[3]Over most of the past 30 years r_t has been less than g_t which has acted as a brake on the growth of the debt ratio, even sometimes allowing the debt ratio to fall during periods of primary deficit.

Figure 1a plots the ratio of public debt-to-GDP (the "debt ratio") in the United States since 1985, with red shaded areas denoting fiscal years where the primary balance is in surplus and grey shaded areas denoting recessions. Overall, the debt ratio was fairly stable through 2008 and then began to rise and has continued to rise. Figure 1b plots the primary deficit over time while Figure 1c plots $r_t - g_t$. Importantly, for much of the 1985 through 2005 period, the primary deficit was zero or less than zero (e.g. primary surplus) during years where interest costs outpaced GDP growth ($r > g$), with the surpluses acting as a brake on debt ratio growth. For much of the past two decades the primary balance was in deficit, but the growth dividend ($r < g$) acted as a debt ratio brake on the primary deficits. In particular, we find that over the last 45 years, the debt-to-GDP ratio has increased 72 percentage points, with the primary deficit leading to a 75 percentage point increase and the growth dividend pushing the ratio down 3 percentage points.

Given the oversized role of the primary deficit to the dynamics of debt over the last forty years, we will begin with a simplistic representation that assumes $r = g$, which suppresses feedback between r, g, and p. Thus, we assume that policy makers' decisions with regards to government spending and tax policy affect the primary deficit but abstract from the potential for these decisions to indirectly affect the interest rate on government debt (r_t) and the growth rate of the economy (g_t). But we address these linkages in subsection 4.5, allowing debt to evolve when $r \neq g$.

2.2 Stochastic Debt Sustainability Analysis (SDSA)

Given the stochastic nature of the path of deficits, r, and g, analyses of sustainability of public debt have focused on understanding whether there exists a plausible fiscal policy that stabilizes the debt-to-GDP ratio after a transitional fiscal consol-

Figure 1: Historical debt, deficit, and $r - g$ data

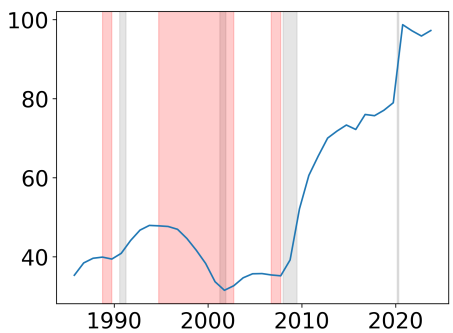

(a) Debt held by the public as % of GDP

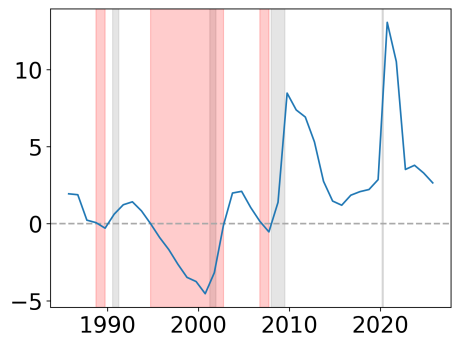

(b) Primary deficit as % of GDP

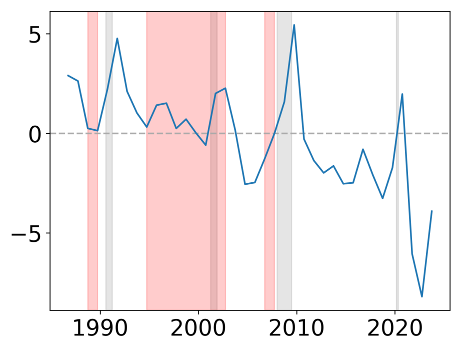

(c) r-g

Note: Grey shading indicates recessions. Red shading indicates years with primary balance in surplus. Where they overlap, both conditions apply. In the bottom panel, r is the nominal average interest paid on debt held by the public, and g is the growth rate of nominal GDP. Source: CBO and FRED.

idation period (Darvas et al. (2025)).[4] These debt sustainability assessments begin by finding the probability that the debt ratio will increase under current fiscal conditions under typical debt dynamics. The debt dynamics are projected many years in the future using current forecasts of interest rates (r), economic growth (g), and the primary balance (p), as in (1). The probability of an increasing debt path is then estimated by adding stochastic uncertainty to (1) in the form of shocks to r, g, and p—typically defined using past realizations of annual changes to these variables (Blanchard (2023)).[5] If an increasing debt path is probable, the next step is to introduce a plausible fiscal consolidation and re-assess whether debt will eventually stabilize.

The typical carried assumption in these models is that shocks to the primary deficit (p), r, and g are jointly *iid* normal with no persistence (European Commission (2023)). But deviations from this *iid* assumption can include building in limited persistence (as in IMF (2021) and Darvas et al. (2025)), or by potentially modelling the full persistence and dynamics in primary balance shocks—whereby an initial shock propagates to future periods via serial correlation—which is the direction taken in this paper. Accounting for this shock persistence will also affect the level of debt accumulation in the models, and in the variance of those levels.

Our analysis begins by exploring the historical data on primary deficit shocks, examining how and why primary deficits evolve. Then we examine the implications for debt dynamics projections when introducing serial correlation to the primary deficit shocks.

[4] Another factor leading to studying sustainability with the SDSA approach is that the dynamics of public debt can be non-linear with respect to the main determinants (r, g, and primary balance, as in (1) above), and the sensitivity to shocks increases with the level of the debt ratio (Reinhart and Rogoff (2010)). While both serve to amplify the risks associated with a rising debt ratio, the level where public debt becomes unsustainable is unknown.

[5] The next step is finding the primary balance that stabilizes the future debt ratio. When $r > g$ then the stabilizing primary balance is a surplus, but when $r < g$ then primary balance may still be in deficit.

3 Modeling the Path of the Primary Deficit

Previous approaches to modeling the primary deficit have assumed *iid* or AR(1) shocks. In this section, though, we examine the evolution of the primary deficit over time and demonstrate serial correlation. Next, we fit the data to a more sophisticated model that allows multiple forms of serial correlation along with differences depending on the source and type of the innovation to the path for the primary deficit.

3.1 The Expected Primary Deficit Path over Time

Figure 2a plots a series of primary deficit projections from the Congressional Budget Office (CBO) in blue and the actual primary deficit over time in red. These CBO projections are typically released twice per calendar year (though we only plot one projection per year in the figure for convenience).[6]

Figure 2a gives an indication that changes in the path of the primary deficit are not independent. In each period, new information is incorporated that leads to changes in the projection in that period as well as future periods, relative to the prior projection path. Figure 2b narrows in on the period from 1992-2002—as the U.S. went from a primary surplus to a primary deficit—to demonstrate this serial correlation. Throughout this period, information led to the subsequent projection path to be revised up (toward deficit), mostly in a level shift as this new information was carried forward in each year of the new projection. Second, there has been a slow rotation in these paths over time, again reflecting new information that gets included not just in the current period but incorporated in future points in the projection path.

When the new information described above is incorporated into deficit projections, it leads to a the sequence of changes to the dynamics of the debt path. Throughout the remainder of the paper, we will define these sequence of changes to the dynamics

[6]We limit our analysis to the projection for five years forward although the length of the projections changed in the later years. In particular, from 1985 to 1995, CBO provided 5-year forward projections and then began producing 10 year projections in 1996.

Figure 2: Projections of deficits and actual deficit

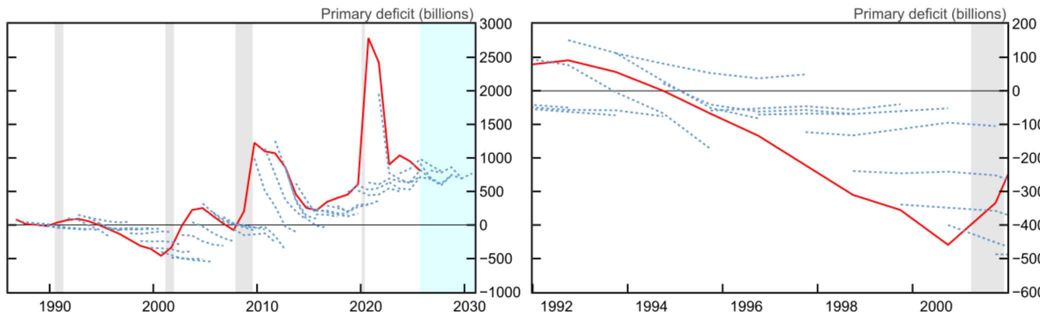

(a) 1986-2030 **(b)** 1992-2002

Note: Dotted lines represent the 5-year out projections of the deficit path for select years. The red line plots the actual deficit.

of the debt path as "innovations." Those innovations that occur in the current period will be called "real-time" innovations. These real-time innovations propagate through to future periods in that debt path (or "future" innovations).

Table 1 describes these innovations to the primary deficit (as a share of GDP) using the first five years of the CBO projection horizon. Over the full pre-Covid (1986-2019) sample period, the new information has tended to revise up the deficit path towards larger deficits. Prior to 2000, though, innovations tended to be deficit-reducing—real-time innovations tended to reduce the deficit projection by 0.37 percentage points of GDP, for example—with the years since 2000 tending to be revised up (second and third rows of Table 1, respectively).

3.2 Decomposition of Innovations to the Deficit Path

How and why do projected deficit paths evolve? The CBO projections data allow a decomposition of the innovations into changes attributable to three types of information: (1) new legislation passed since the most recent update, (2) new information about the state of the economy, and (3) a reassessment of the budget path unrelated to economic or legislative information (technical reasons), and do so separately for

Table 1: Mean CBO primary balance revisions, 1986-2019

	Projection horizon				
	0	1	2	3	4
	Full sample				
1986-2019	0.43	0.67	0.43	0.28	0.16
	Sub-periods				
1986-1999	-0.37	-0.06	-0.17	-0.20	-0.32
2000-2019	0.98	1.18	0.85	0.62	0.50
	Memo: 2020 and beyond				
2020+	2.53	0.67	0.20	-0.03	-0.01

Note: Expressed as a share of projected GDP (× 100) in each CBO budget outlook release (or percentage point of GDP). A revision that decreases (increases) primary balance deficit is expressed as a negative (positive) number. The "real-time" innovation is horizon 0.

revenues and outlays.

Figure 3 shows CBO's innovations decomposed by category (outlays excluding net interest or revenues), type (economic, legislative, or technical), and projection horizon. The x-axis is the fiscal year in which CBO released the information. There are notable differences in the magnitude of the innovations according to each dimension of the decomposition–category, type, and horizon.

The differences in magnitudes are evidenced in Figure 4, which shows the mean of the innovations over the sample time period (1986-2019), decomposed by category type, and horizon. The net height in each projection horizon column is the total reported in Table 1. Legislative innovations (orange bars) have tended to be the largest source of new information, and have tended to drive deficit increases on

Figure 3: Innovations by category, type and horizon

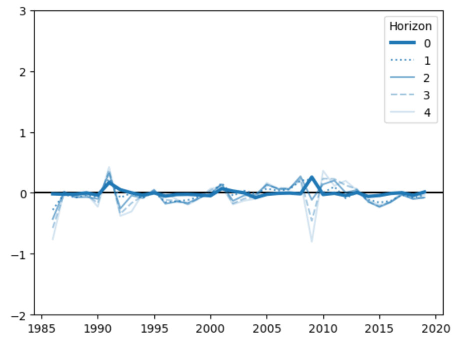

(a) economic innovations to outlays

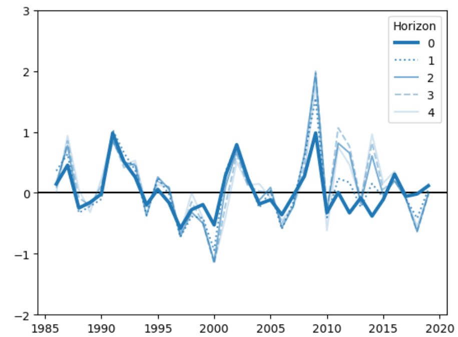

(b) economic innovations to revenues

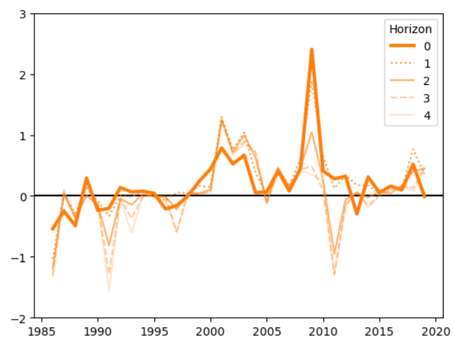

(c) legislative innovations to outlays

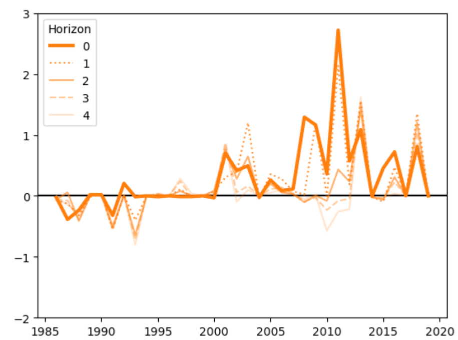

(d) legislative innovations to revenues

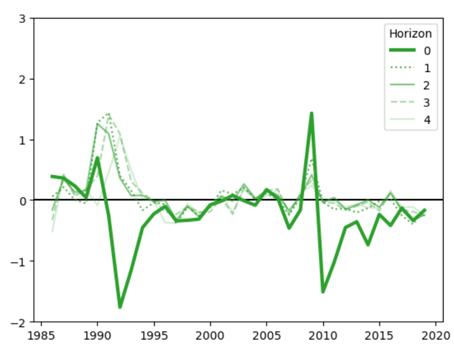

(e) technical innovations to outlays

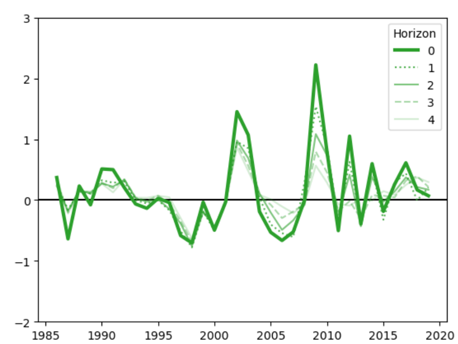

(f) technical innovations to revenues

Note: X-axis is the fiscal year of CBO's publication.

net over the full pre-Covid sample period.[7] Real-time technical innovations have tended to push down deficits on net between revenues (upward push) and outlays (downward push). Economic innovations tend to be small in magnitude, though also persist across the horizon in Figure 4.

Figure 4: Innovation means by category and type

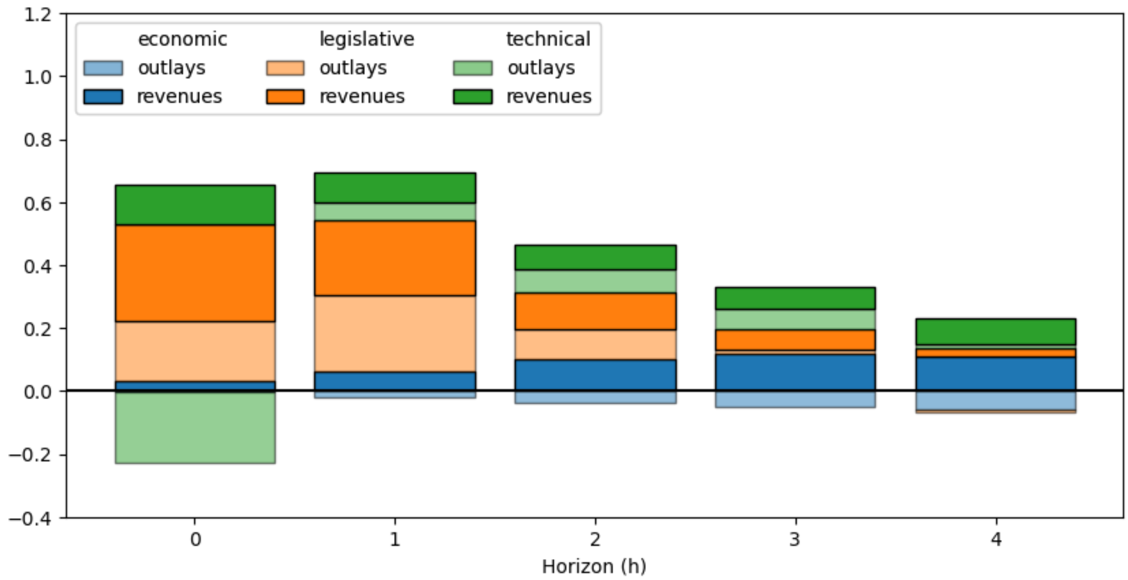

Note: Positive (negative) values indicate increase (decrease) in primary balance deficit.

The previous figures hint at several types of serial correlation in the deficit innovations. The first is serial correlation in the "real-time" innovations across adjoining data releases (t), which is seen in Figure 2a, as well as in the different innovation means during different periods in Table 1. The second type of correlation is between the real-time and the future innovations in a given projection path, which is seen

[7]Since 2000, new revenue and outlays legislation has tended to push deficit projections up with technical and economic revisions exerting smaller upward pressure on deficits (appendix Table 2, bottom panel). Prior to 2000, though, legislation exerted downward pressure on deficits, with revenue increases and outlays decreases roughly similar in magnitude (appendix Table 2, middle panel).

when looking across horizons in Table 1, Figure 3, and Figure 4. This correlation directly links the pass-through of the real-time innovation to future period in the projection path. Finally, there is correlation in the overall path, such that for further-out horizons the $t-1$ period can generally predict the value in t. This is most clearly shown in Figure 2a. This next subsection introduces a model to account for all three types of serial correlation, which are used to create a data generating process of the primary deficit and debt accumulation.

3.3 Modelling the Data Generating Process

Consistent with the data available from CBO, at each period t, our model consists of a projected deficit path for the current period and the next five years. Using the notation in (1), we denote the primary deficit at time t as p_t, and can express the projected path for the primary deficit as of time t for the next five periods in the vector form as

$$\begin{bmatrix} p_t \\ p_{t+1,1} \\ \vdots \\ p_{t+5,5} \end{bmatrix} \tag{2}$$

For the projected path, the first subscript of each element is the fiscal year for which we are estimating the deficit, and the second subscript is the horizon of the projection, i.e. the number of years prior to the fiscal year. Consistent with this notation, we can write $p_t = p_{t,0}$.

Moving from period t to period $t+1$, we update the existing projection from period t for the whole path with innovations, simulating new information learned. Our projection window changes, dropping time t and now including time $t+6$. We update the deficit path in (2) with innovations and insert an initial projection for time $t + 6$ as shown below:

$$\begin{bmatrix} p_{t+1,1} \\ p_{t+2,2} \\ p_{t+3,3} \\ p_{t+4,4} \\ p_{t+5,5} \\ \cdot \end{bmatrix} + \begin{bmatrix} \delta_{t+1,0} \\ \delta_{t+2,1} \\ \delta_{t+3,2} \\ \delta_{t+4,3} \\ \delta_{t+5,4} \\ \cdot \end{bmatrix} \rightarrow \begin{bmatrix} p_{t+1,0} \\ p_{t+2,1} \\ p_{t+3,2} \\ p_{t+4,3} \\ p_{t+5,4} \\ p_{t+6,5} \end{bmatrix} \tag{3}$$

In (3), the vector being added to the existing path is the vector of innovations to the projected deficit path. This represents all of the information learned at time $t + 1$ that modifies the path, which we denote by δ. Note that the final element of the projected deficit path in the right hand side of (3) is the initial deficit projection for the new period entering the projection window. Because we have only a 5-year projection of the deficit in the projection window, we have no existing projection of the deficit from the previous period to update (shown with a dot in the vectors to the left). Thus, at each period, we must introduce a projection for a new year. We model these initial projections as a function of the end of the projected path such that $p_{t+6,5} = \alpha p_{t+5,4}$. As we formalize below, this modelling approach allows for innovations to propagate forward to the path of the deficit outside the current projection period.

Returning to the evolution of the path within the projection period, as in CBO's publications, at each time t, we have a 5-year forward looking path for the deficit as in (2), and looking across multiple periods using (3), we have a sequence of projections of the deficit for each year that evolves with new information over a 5-year span:

$$
\begin{aligned}
&p_{t,5} \\
&\ \downarrow \\
&p_{t,4} = p_{t,5} + \delta_{t,4} \\
&\ \downarrow \\
&\ \vdots \\
&\ \downarrow \\
&p_{t,0} = p_{t,5} + \sum_{4 \leq h \leq 0} \delta_{t,h} = p_t
\end{aligned}
\tag{4}
$$

Thus, our model for the actual deficit in any given year has two components which we must describe–(1) the innovations to the projected deficit path and (2) the initial projection of the primary deficit when a new period enters the projection window.

We start by describing our model for the different versions of innovations to the projected deficit path at a given period. First, there is the real-time innovation which updates the current fiscal year's projection, $\delta_{t,0}$. Building off the observations in previous sections about correlation in the overall path, we allow for correlation between the current real-time innovation, $\delta_{t,0}$, and the previous period's real-time innovation, $\delta_{t-1,0}$. Secondly, we separately model how information learned today affects the future path. Because we update the entire deficit path within the projection window, our model allows for shifts in the entire future path due to new information. We capture this by modeling correlation between the current real-time innovation and the innovations to the future path, $\delta_{t+h,h}$ for $1 \leq h \leq 4$, and call these "direct" future path innovations. Our model for real-time innovations ($h = 0$) and direct future path innovations ($1 \leq h \leq 4$) is:

$$\delta_{t,h} = \begin{cases} \rho_h \delta_{t-1,0} + \epsilon_h, & h = 0 \\ \rho_h \delta_{t,0} + \epsilon_h, & 1 \leq h \leq 4 \end{cases} \tag{5}$$

where $\epsilon_h \sim \mathcal{N}(\mu_h, \sigma_h)$. Due to the differences in innovations according to category and type seen in Figure 3 and Figure 4, we use data from the CBO reports to estimate the parameters in (5) separately for each category and type. We suppress that notation here to focus on the relationship between real-time and future innovations.

In addition to updating the existing deficit path, at each period, our model must also produce an initial projection of the primary deficit five years in the future. We can think of having a deficit path that extends indefinitely, and updating it with information learned at the current time. However, we have no information about innovations outside the observable CBO window, which we refer to as "path persistent" future innovations ($h > 4$). We instead produce an initial projection of the primary deficit by allowing for correlation with the current period's updated projected deficit path. At time $t - 5$, we create an initial projection of the primary deficit at time t using the model:

$$p_{t,5} = \alpha \left(p_{t-1,5} + \sum_{\substack{c \in C \\ y \in Y}} \delta_{t-1,4} \right) \tag{6}$$

Here, we explicitly show the decomposition of the innovations according to the collection of categories, C, and types, Y.[8] However, since we do not have data from which to estimate parameters, we do not attempt to explicitly model the correlation between real-time innovations and these persistent future innovations separately by

[8] $C = \{$outlays excluding net interest, revenues$\}$, $Y = \{$economic, legislative, technical$\}$

source and type.

This model for "path persistent" future innovations provides a reasonable method to extend our projections over longer periods and allows for serial correlation to exist outside of the budget window. Figure 5 shows that the model fits CBO's historical initial projections. Performing a linear regression on CBO's historical projections yields $\alpha \approx 0.995$, which creates a strong persistence in the path of the primary deficit. However, since this is only our initial estimate of the deficit at time t, it will evolve according to (4), eventually incorporating 5 years of innovations.

Figure 5: CBO's 5 years out deficit projection versus model

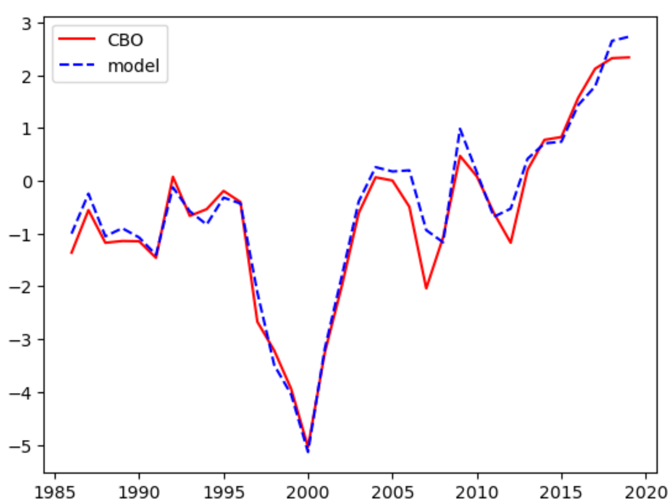

Note: The red line shows the 5-year-out projection of the primary balance as a share of GDP from CBO. The blue line is the estimate as described in (6).

Up until this point, we have focused primarily on how the model updates the deficit path at each period. However, as mentioned previously, we can also look across periods to see how the debt projection for a particular time has evolved, ultimately leading to our model for the primary deficit at time t:

$$p_t = p_{t,5} + \sum_{4 \geq h \geq 0} \sum_{\substack{c \in C \\ y \in Y}} \delta_{t,h} \tag{7}$$

where $\delta_{t,h}$ is given by (5) and $p_{t,5}$ is given by (6).

Guided by the data generating process described above for the primary deficit, we next create a debt path associated with these innovations. For our primary results, we simplify by initially assuming $r_t = g_t$ in (1), and simulate debt accumulation using the model:

$$b_T = b_0 + \sum_{t=1}^{T} \left(p_{t,5} + \sum_{4 \geq h \geq 0} \sum_{\substack{c \in C \\ y \in Y}} \delta_{t,h}^{c,y} \right) \tag{8}$$

and focus on the distribution of outcomes twenty years in the future.

4 Simulation Results

Next, we embed the data generating process in the previous section, which includes the three types of serial correlation, into a stochastic debt sustainability analysis (SDSA) following the debt dynamics from (8). We demonstrate that our full model that allows for the different types of serial correlation produces a very different distribution than a parsimonious model that follows an AR(1) process. Next, we decompose the layers of serial correlation in our model to better understand the role of each type of serial correlation in debt projections, and also decompose debt accumulation by source—legislative, economic, and technical.

4.1 SDSA with Simple or Full Serial Correlation

We begin by simulating future deficits and debt-to-GDP ratios using the estimated moments of the data generating process from the model above, incorporating the mean and variance of $\epsilon_h^{c,y}$, estimated $\rho_h^{c,y}$, and estimated α. We draw 10,000 realizations for the full path of the primary deficit projected over twenty years in order to develop a distribution of the effect of these innovations to the level of debt after twenty years. In particular, we determine the simulated debt ratio paths using the debt accumulation (8) after modeling data-driven deficit paths as in (6).

Figure 6a plots the distribution of the evolution of debt over time starting from CBO's projection for the current path for the primary deficit. The black line plots the path for the debt ratio that CBO projects. Alternatively, the blue dashed line is the modal outcome for debt assuming that the path for the primary deficit continues to experience innovations consistent with our model (8). Notably, these innovations lead to a much faster accumulation of debt.

Figure 6: Evolution of debt path, by initial primary deficit

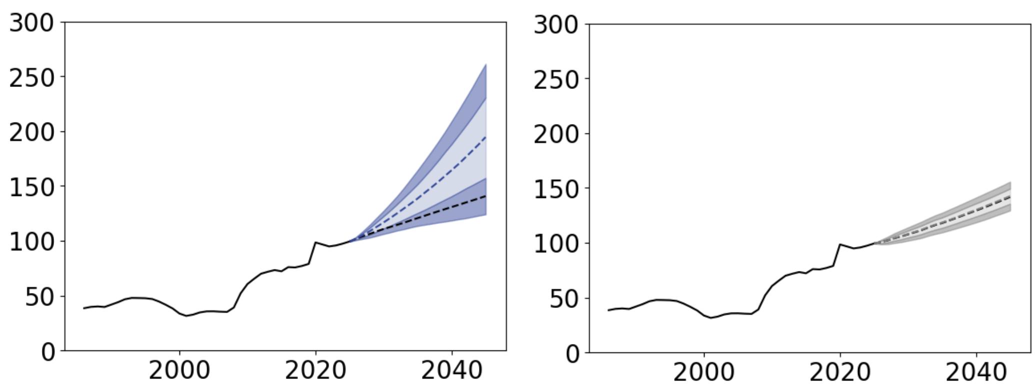

(a) Full Model

(b) AR(1)

Note: Black line is the historical time series of the US debt ratio: debt held by the public to GDP. The dashed black line is CBO's projected path. The inner, lightly shaded cone is the 75% confidence interval surrounding the median path of each model distribution and the dark shaded area is the 90% confidence interval. On the left, we show the distribution resulting from (8) and on the right we show an AR(1) model for the primary deficit.

The two shaded regions plot the 75 percent and 90 percent confidence intervals around the distribution of outcome. Though they are quite wide, both point to CBO's path being fairly unlikely—falling outside the 75 percent confidence interval of our model. Most importantly, none of the outcomes within the 90 percent confidence intervals would be consistent with a sustainable path, loosely defined as a debt-to-GDP ratio eventually stabilizing.

Most of the previous SDSA incorporate no or limited serial correlation and, as a comparison, Figure 6b plots the distribution of debt innovations assuming this simpler model.[9] Compared to the simpler model, our model leads to the expectation of significantly more debt and to more uncertainty around the potential outcomes. Thus, the additional serial correlation we introduce has important consequences for SDSA.

4.2 Roles of Types of Serial Correlation

Next we examine the role for each type of serial correlation we include in our model to determine why this model produces a different distribution than the more parsimonious approach of an AR(1) process. In particular, Figure 7a plots the distribution of primary deficit innovations after 20 years sequentially adding the different types of serial correlation in our model, as introduced in (5). Starting with no serial correlation ($\rho_h = 0$ for $h \geq 0$ in (5), labeled "None"), introducing just correlation across adjoining real-time innovations ($\rho_h = 0$ for $h > 0$, labeled "RT") has little effect on the distribution. Further allowing correlation between real-time and future innovations (as in (5), labeled "RT + F") slightly increases the variance in deficit outcomes. However, the complete model—which also includes path persistence (as in (6), labeled "RT + F + PP")—widens the distribution notably and shifts it to the right (to larger deficits).

In Figure 7b, we show the distribution of the innovations to the debt after 20 years in the scenarios outlined above. The innovation to the debt for a particular path is the sum of the primary deficit innovations on that path as in (8). Allowing

[9] The model shown here assumes $r = g$, meaning the confidence interval is formed purely using the primary deficit. In 4.5 we will remove this constraint.

Figure 7: Innovations to deficit and debt 20 years out, comparing serial correlation

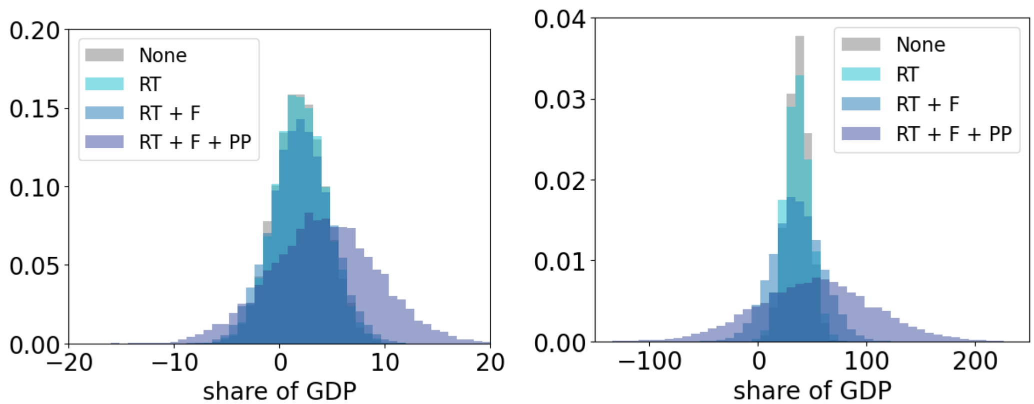

(a) deficit innovations

(b) debt innovations

Note: The left panel plots the distribution of primary deficit innovations after 20 years, adding in three layers of serial correlation sequentially (see (5)). The right panel shows the distribution of the innovations to the debt after 20 years—the sum of the primary deficit innovations on that path (as in (8)).

these deficit outcomes to accumulate over 20 years amplifies any serial correlation since debt is the stock of the changes in primary deficits over time. Accordingly, the complete model ("RT + F + PP"), which incorporates all types of serial correlation, has considerable variation and increases the debt, on average, by nearly 15 percentage points over the other distributions.

As in Figure 7a, the serial correlation across adjoining real-time innovations (or what Auerbach (1999) calls a "forecast error"), labeled "RT", is only slightly different than a model that includes no serial correlation (labeled "None"), and further allowing the serial correlation of innovations to the future path ("RT + F") leads to a somewhat more disbursed distribution. But, importantly, incorporating the path persistence serial correlation, the "RT + F + PP", leads to a notable increase in the variation, and shifts the distribution to the right. Thus, although all types of serial correlation effect the distribution of outcomes, the persistence of the path of primary deficit outside the projection window is the most important contributor to both the increase in the debt and also the uncertainty (and largely ignored by previous studies).

4.3 Innovations in Primary Deficit: Decomposition by Source

By developing a data generating process for the innovations for the whole path of the primary deficit at any given point in time, we can incorporate serial correlation and leveraging CBO's projections can also separate innovations by their different sources— economic changes, technical changes, and legislative changes. Figure 8 decomposes the distribution of debt innovations after twenty years into the effects from these three sources and the additional changes due to the path persistence correlation.[10] Overall, the distribution of the effects from economic and technical changes are tightly centered around zero while the effects from legislative innovations play an important role in the distribution of the overall change in debt being positive. Moreover, the transmission of these innovations through the path persistence further shifts the overall distribution to the right and also adds significantly to the variance.[11]

4.4 What Leads to a Sustainable Path?

Given that the trajectory of the debt path is unsustainable (Figure 6a), the obvious question is what is necessary to shift the distribution of outcomes to be more consistent with debt sustainability. Importantly, we demonstrate that the legislative innovations are one key component leading to the distribution of the debt path pointing to an upward slope (Figure 8). Thus, we consider two counterfactual scenarios, both involving large policy changes. The first is a counterfactual where policy makers

[10] We are unable to determine separate data generating processes for the path persistence correlation for the three sources of revisions and thus we separate the effect of this type of serial correlation and do not assign it separately to each source of revision.

[11] A decomposition of changes to the primary deficit—into those driven by revisions to revenues and outlays (excluding net interest) separately for each of the three sources of revision: economic, technical, and legislative—can be found in Figure 12a-Figure 12c. Figure 12a decomposes the "RT + future" model in Figure 7b into the three sources of revision, and panels Figure 12b and Figure 12c separate into the revenues and outlays categories.

Figure 8: Total (debt innovations) incl. path persistence

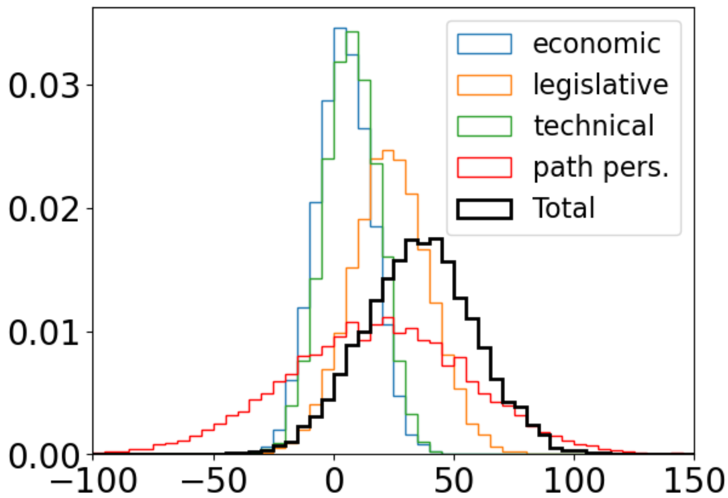

Note: The x-axis is debt as a share of GDP. Decomposes the "RT + F + PP" area in Figure 7b (black area here) into the three sources of revision—economic, legislative, and technical—along with a "path persistence" component of the model.

make a drastic one-time abrupt change in legislative policy such that the path of the expected primary deficit as of 2026 is set to zero. To further understand the role of the legislative innovations, we assume that there are no longer any innovations to the path of the primary deficit due to economic and technical changes but continue to draw the shocks from legislative innovations going forward.

Relative to the baseline projection, the modal outcomes from this counterfactual scenario are notably lower in levels (A in Figure 9b) and the distribution of the debt ratio 20 years out has fewer large increases (Figure 9a). But even this large policy change does not lead to a sustainable debt path, as the median debt path is still increasing at the end of the 20 year horizon and the debt ratio increases in more than 80 percent of all simulated paths.

In the second counterfactual scenario, we force a significant change to the legislative innovations by modeling the legislative innovations (and model coefficients) to match those consistent with the pre-2000 era. As noted in Figure 3, the legislative changes were more muted in this period, with periods of upward and downward revisions but on average the innovations were contractionary. In this counterfactual scenario, the modal path of debt largely stabilizes by the end of the 20 year period—a necessary condition for debt sustainability (B in Figure 9b).[12] However, because it takes time for the innovations to reduce the primary deficit, on average, a majority of estimated debt accumulation paths still lead to debt ratio increases before the debt ratio stabilizes at a larger level than the jumping-off point. Thus, we find that in order for the distribution of debt outcomes to be consistent with a path that eventually flattens out, a large one time change is less helpful than a change to the underlying distribution of innovations.

Figure 9: Debt evolution under legislative scenarios

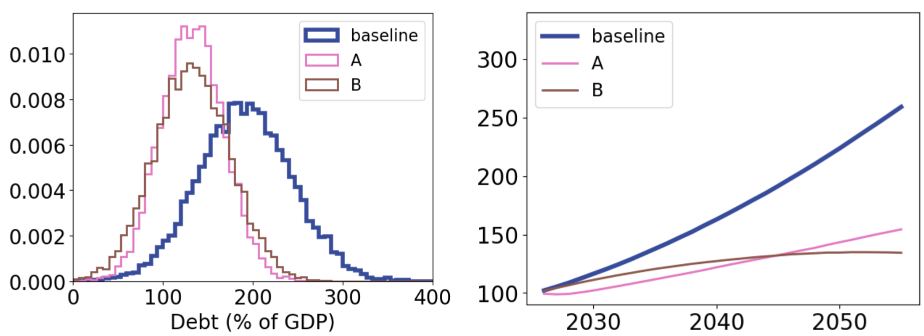

(a) Debt distribution 20 years out

(b) Median debt path over next 30 years

Note: Consider various legislative shocks: A = one-time consolidation shock (p_0 to zero and no future legislative changes). B= shift in legislative process (use parameters estimated from 1992-1999 period).

[12]As noted in Auerbach and Yagan (2025), there was legislative feedback in this period whereby increases in the debt ratio led to a dampening of deficits. And while scenario B in Figure 9b and Figure 9a includes only data from 1992-1999, a replication with the full pre-2000 data does not overturn the stabilization result.

4.5 Scenarios when $r \neq g$

Our simulations for the debt path thus far have assumed $r = g$, which allows debt dynamics to evolve purely as a function of the primary deficit (setting $r = g$ forces the second term in equation (1) to zero). This simplification abstracts from the fact that an increase in the debt ratio is often associated with slower GDP growth when the debt level is elevated (de Soyres et al. (2022)), and with an increase long-run interest rates (Gamber and Seliski (2019), Plante et al. (2025)).[13]

To more fully model debt accumulation taking into account some potential feedbacks, simulations in this final section initially set r and g at values according to CBO's latest near-term projections, and then allow them to dynamically update according to the simulated debt path. Specifically, for every one percentage point increase in the debt ratio, g will decrease by 1.4 basis points in the subsequent year (as in Heimberger (2021)), and the interest rate on newly issued debt—which gets incorporated into r—will increase by 2 basis points (as in Neveu and Schafer (2024)).[14]

The outcomes in this more-complete model—shown in Figure 10— broadly tell the same story as in Figure 6a. Overall, debt accumulates more quickly in our model than CBO. However the accumulation tends to be somewhat larger in the model with the feedbacks, particularly towards the end of the twenty year window, because as debt continues to accumulate, the pressure on the $r - g$ differential serves to heighten debt growth. Similarly, the distributions of debt accumulation shown in the counterfactual scenarios in Figure 11a paint a similar story, however, when

[13]The exact nature of these relationships remains uncertain. The feedback from debt to GDP growth may be closer to zero (Furman and Summers (2019)), and estimates of the feedback effect from debt to interest rates are sometimes smaller (Neveu and Schafer (2024)), or closer to zero (Peterman and Sager (2025)).

[14]The initial values of r and g are set as the average 5-year forward values in the CBO projections (CBO (2025a)). The average term on outstanding federal government debt is about six years, so we assume that one-sixth of debt is re-issued annually at this new interest rate.

feedbacks are included the outcomes are somewhat more variable than when $r = g$ in Figure 11b.

Figure 10: Evolution of debt path, by initial primary deficit, dynamic r and g

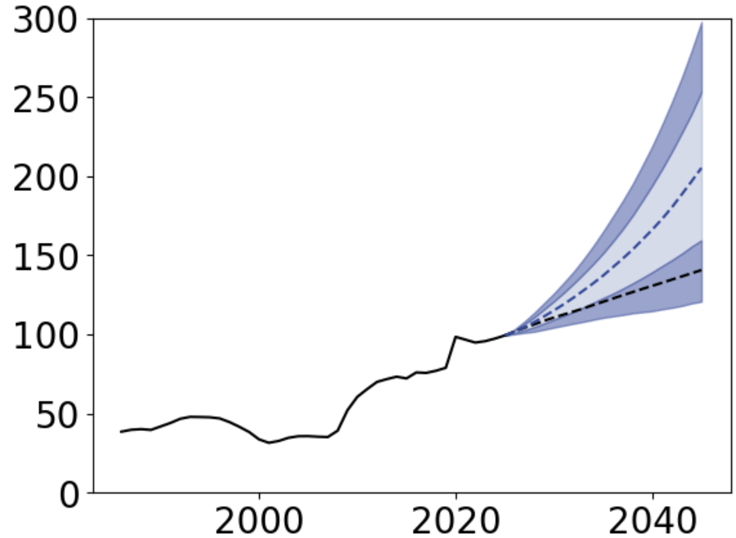

Note: As in Figure 6a, this figure shows stochastic projections of debt accumulation under equation (8) (blue dashed line and shaded areas), though the simulations shown here also include dynamic estimates of r and g. Black line is the historical time series of the US debt ratio (debt held by the public to GDP) and the dashed black line is CBO's projected path. The light blue cone is the 75% confidence interval surrounding the dashed blue line and the dark blue shaded area is the 90% confidence interval.

5 Conclusion

The current debt-to-GDP ratio is high by historical standards and is projected to continue to increase, leading to debt sustainability concerns. Typical projections of the debt ratio include stochastic processes for r, g, and the primary balance. But the assumptions that govern the dynamics of each process may also constrain the usefulness of the projections.

In this paper, we provide a detailed model of the data generating process of the primary deficit path, documenting notable serial correlation in the innovations to

Figure 11: Debt evolution under legislative scenarios, dynamic r and g

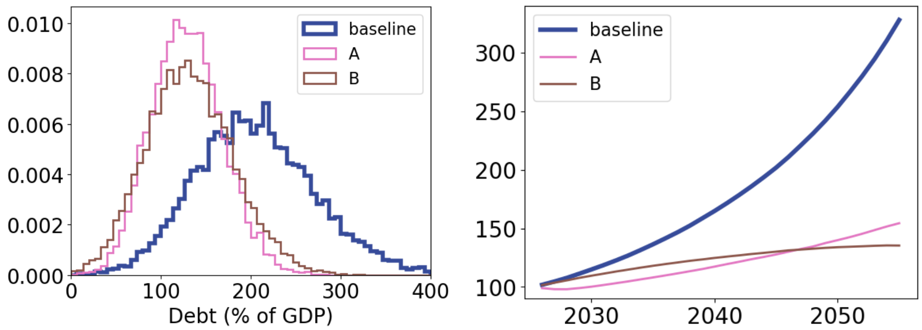

(a) Debt distribution 20 years out

(b) Median debt path over next 30 years

Note: As in Figure 9, this figure shows debt accumulation outcomes under alternate legislative scenarios, though the simulations shown here also include dynamic estimates of r and g. The legislative shocks: A = one-time consolidation shock (p_0 to zero and no future legislative changes). B= shift in legislative process (use parameters estimated from 1992-1999 period).

this path, and demonstrate that incorporating this serial correlation into a model of debt accumulation has a meaningful impact on debt accumulation projections. In most projections using our modeled primary balance process there was more debt accumulation than in projections made without, implying that current projections are closer to a lower bound of debt accumulation.

Legislation is the most important source of innovations, and we examine two types of potential legislative policy changes to help put the debt ratio on a sustainable path. A sustained legislative consolidation will help stabilize the debt ratio while a one-time consolidation policy will not.

References

Auerbach, A. (1999). On the performance and use of government revenue forecasts. *National Tax Journal*, 36(4).

Auerbach, A. J. and Yagan, D. (2025). Robust fiscal stabilization. Working Paper 33374, National Bureau of Economic Research.

Blanchard, O. (2023). *Fiscal policy under low interest rates*. MIT press.

CBO (2024). An evaluation of CBO's projections of deficits and debt from 1984 to 2023. Technical report.

CBO (2025a). The budget and economic outlook: 2025 to 2035.

CBO (2025b). The long-term budget outlook: 2025 to 2055.

Darvas, Z., Huertas, G., Welslau, L., and Zettelmeyer, J. (2025). What will it take to stabilize debt in advanced countries? Working Paper 28/2025, Breugel.

de Soyres, C., Kawai, R., and Wang, M. (2022). Public debt and real GDP: Revisiting the impact. Working Paper WP/22/76, IMF.

Dynan, K. and Elmendorf, D. (2025). Putting us fiscal policy on a sustainable path. Technical report, National Bureau of Economic Research.

Edelberg, W., Harris, B., and Sheiner, L. (2025). Assessing the Risks and Costs of Rising U.S. Federal Debt. Working papers, Economics Studies at Brookings.

European Commission (2023). Debt sustainability monitor 2022. Working Paper 199, Directorate General of Economic Affairs, European Commission.

Furman, J. and Summers, L. (2019). Who's afraid of budget deficits? *Foreign Affairs*.

Furman, J. and Summers, L. (2020). A reconsideration of fiscal policy in the era of low interest rates. Technical report, Harvard University and Peterson Institute for International Economics.

Gamber, E. and Seliski, J. (2019). The Effect of Government Debt on Interest Rates. Working Papers 2019-01, Congressional Budget Office.

Heimberger, P. (2021). Do Higher Public Debt Levels Reduce Economic Growth? Technical report, The Vienna Institute for International Economic Studies.

IMF (2021). Review of the debt sustainability framework for market access countries.

Neveu, A. R. and Schafer, J. (2024). Revisiting the Relationship Between Debt and Long-Term Interest Rates. *Congressional Budget Office WP*, 5.

Peterman, W. and Sager, E. (2025). Public debt and interest rates. Working paper, Federal Reserve Board of Governors.

Plante, M. D., Richter, A. W., and Zubairy, S. (2025). Revisiting the Interest Rate Effects of Federal Debt. *Federal Reserve Bank of Dallas Working Paper*, 2513.

Reinhart, C. M. and Rogoff, K. S. (2010). Growth in a time of debt. *American Economic Review*, 100(2):573–78.

Zenios, S. A., Consiglio, A., Athanasopoulou, M., Moshammer, E., Gavilan Gonzalez, A., and Erce, A. (2021). Risk Management for Sustainable Sovereign Debt Financing. *Operations Research*.

6 Appendix

The appendix describes several tables and figures in more detail, by breaking out the deficit into the part attributable to revenues and to outlays.

6.1 Means

The revisions described in Table 1 can be decomposed to understand the reasons for the changes in deficit projections, and the top panel of Table 2 shows that legislative changes can explain nearly all of the deficit revisions in the full sample period. For example, in the real-time innovations, the deficit tends to be revised up by about 0.43 percentage points of GDP (Table 1, top panel, first column). Of that, more than the entire amount can be attributed to legislative (0.50 percentage points of GDP), with non-legislative revisions pushing deficits down by 0.07 percentage points (Table 2, top panel).

Since 2000, new revenue and outlays legislation has tended to push deficit projections up with technical and economic revisions exerting smaller upward pressure on deficits (Table 2, bottom panel). Prior to 2000, though, legislation exerted downward pressure on deficits, with revenue increases and outlays decreases roughly similar in magnitude (Table 2, middle panel).

Table 2: Mean CBO revisions by type of revision

	Economic					Technical					Legislative				
	Projection horizon					Projection horizon					Projection horizon				
	0	1	2	3	4	0	1	2	3	4	0	1	2	3	4
	1986-2019 (All)					*1986-2019 (All)*					*1986-2019 (All)*				
Revenues	0.03	0.06	0.10	0.12	0.11	0.13	0.09	0.08	0.07	0.08	0.31	0.24	0.12	0.07	0.03
Outlays	-0.00	-0.02	-0.04	-0.05	-0.06	-0.23	0.06	0.07	0.07	0.02	0.19	0.24	0.10	0.02	-0.01
Deficit	0.03	0.04	0.06	0.07	0.05	-0.10	0.15	0.15	0.14	0.10	0.50	0.48	0.22	0.08	0.02
						By period									
	1986-1999					*1986-1999*					*1986-1999*				
Revenues	0.04	0.05	0.05	0.09	0.13	-0.03	-0.03	-0.02	0.00	0.02	-0.06	-0.10	-0.10	-0.11	-0.11
Outlays	-0.00	-0.06	-0.09	-0.11	-0.12	-0.23	0.17	0.19	0.21	0.07	-0.09	-0.10	-0.20	-0.28	-0.31
Deficit	0.04	-0.01	-0.04	-0.02	0.01	-0.26	0.14	0.16	0.21	0.10	-0.14	-0.20	-0.30	-0.38	-0.42
	2000-2019					*2000-2019*					*2000-2019*				
Revenues	0.02	0.07	0.13	0.14	0.09	0.24	0.18	0.15	0.12	0.12	0.56	0.47	0.27	0.18	0.12
Outlays	-0.00	0.01	0.00	-0.01	-0.02	-0.22	-0.03	-0.01	-0.03	-0.03	0.39	0.48	0.31	0.22	0.21
Deficit	0.02	0.07	0.13	0.13	0.07	0.02	0.15	0.14	0.08	0.10	0.95	0.96	0.58	0.40	0.33

Notes: See Table 1 notes. A positive value implies the revisions are increasing the primary deficit.

6.2 Decomposition of change in primary deficit: revenues, outlays

Figure 8 decomposes the distribution of debt innovations after twenty years into the effects from the three sources—economic, legislative, technical. Here we further decompose the changes to the primary deficit into those driven by revisions to revenues and outlays (excluding net interest).

Figure 12a replicates much of figure 8, and panels figure 12b and 12c separate these three sources of revision by the revenues and outlays categories. In simulations 20 years out, legislative innovations drive the dispersion in accumulated debt (figure 12a). Legislative innovations drive most of the dispersion in outlays (figure 12c) and are the biggest driver in revenues (figure 12b).

Figure 12: Accumulated innovations 20 years out, broadest serial correlation, by outlays and revenues

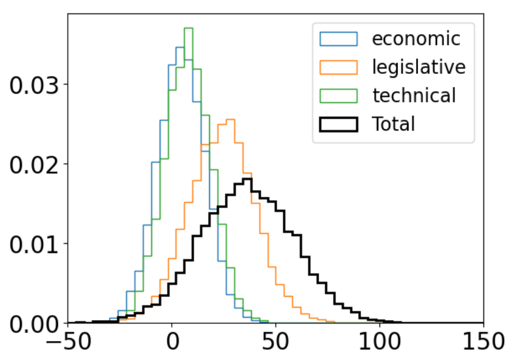

(a) Total (Rev+Outlays)

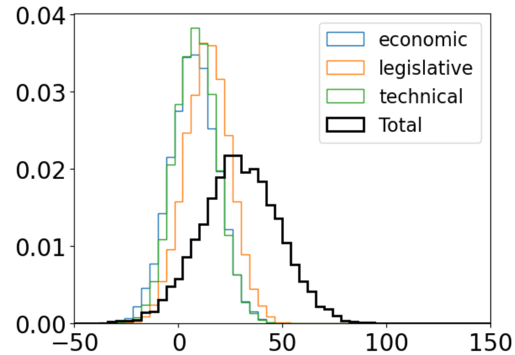

(b) Revenues

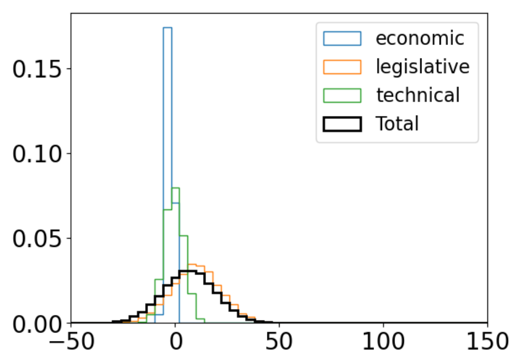

(c) Outlays

Note: X-axis is debt as share of GDP. Panels 12a and 12c show the three sources of revision separately by revenues and outlays.