Hutchins Center Working Paper #46

October 2018

Optimizing the Maturity Structure of U.S. Treasury Debt: A Model-Based Framework

Terry Belton
JPMorgan Chase

Huachen Li
JPMorgan Chase

Kristopher Dawsey
The D. E. Shaw Group

Srini Ramaswamy
JPMorgan Chase

David Greenlaw
Morgan Stanley

Brian Sack
The D. E. Shaw Group

ABSTRACT

This paper describes a model that can be used to inform debt management choices facing the U.S. Treasury. The model captures the dynamic interactions between economic conditions, budget deficits, and interest rates at various maturities, and it uses that structure to assess the likely outcomes for government funding costs under any assumed debt issuance strategy. The model—which was developed by several members of the Treasury Borrowing Advisory Committee (TBAC) and has appeared in several TBAC presentations—can be used to assess the tradeoffs between various issuance strategies and to explore the sensitivity of those tradeoffs to different assumptions. The results under a baseline set of assumptions suggest that issuing at intermediate maturities is appealing, as those securities do not involve high expected costs and yet are effective at smoothing variation of interest expense and deficits. In addition, the results show that debt managers can achieve considerable gains by employing a reaction function that varies issuance patterns in response to a small set of economic and financial variables.

Belton, Greenlaw, and Sack are members of the Treasury Borrowing Advisory Committee — a committee composed of senior representatives from investment funds and banks that meets quarterly with the Treasury Department. This paper was prepared to further discussion of potential changes to the model that informs the U.S. Treasury's debt management choices. The authors are employees of investment funds and banks, however, they did not receive financial support for their work on this paper.

The authors would like to thank participants from an authors' workshop held at The Hutchins Center on Fiscal and Monetary Policy at Brookings in May 2018 for their helpful feedback, and Colin Kim for his helpful comments. The authors would also like to thank Jason Cummins for suggesting this project, and Walter Mueller for providing guidance in its early stages. All views expressed in this paper are those of the authors and do not represent the views of JPMorgan Chase, The D. E. Shaw Group, Morgan Stanley, or the Brookings Institution.

1. Introduction

Government borrowing costs constitute a large share of public expenditures in many countries, and hence it is an important policy priority to achieve the lowest possible borrowing cost over time with the lowest possible risk. A government debt manager can decide how to allocate his or her issuance over a wide range of maturities in order to achieve this outcome. These issuance points will typically involve a tradeoff between expected funding costs and fiscal risks. On the one hand, short-term financing may be relatively inexpensive but can complicate budget planning by raising the variability of near-term interest expense. On the other hand, longer-term borrowings mitigate the volatility of financing costs, but typically have a higher expected cost because term premia tend to be higher as maturity lengthens. A primary responsibility of debt managers is to strike a balance between these competing considerations and determine the optimal issuance profile across all maturities.[1]

Striking this balance is no easy task. Issuers have to take into account the considerable variability of interest rates and budget deficits — and the correlations between them – over relatively long horizons. And they have to consider a variety of potential issuance approaches. Assessing these approaches is a challenging exercise that requires an analytical framework.

Our objective in this paper is to put forward such an analytical framework that can be used to optimize the maturity structure of U.S. Treasury debt. The model proposed is similar in nature to those that have been employed in some other countries, including Canada, the UK, Sweden, Brazil, and Turkey (see Appendix A for more discussion). Many differences exist in the exact structures of these models and in their specification of costs and risks. But while there are important differences in the exact model specifications, the overall approach used in different countries appears to be broadly similar and generally involves an underlying model driving the evolution of economic variables and interest rates, and an optimization module in which the debt manager is assumed to minimize expected issuance costs through time given constraints on risk and other variables. Some of the important similarities and differences across these models are summarized in the appendix.

To date, there has not been a particular model that has been applied to assess debt management decisions for the U.S. Treasury. Most discussion of the U.S. Treasury's debt management strategy has focused on maintaining a "regular and predictable" pattern of issuance. We fully agree that this approach has served the Treasury well, as it has contributed to the substantial liquidity of the Treasury market and has promoted lower yields and hence more favorable borrowing conditions for the U.S. Treasury (see Garbade, 2007). However, being "regular and predictable" does not specify the ultimate maturity structure that best serves the U.S. Treasury. This paper is focused on assessing the optimal maturity structure, while maintaining the view that any adjustments to issuance to achieve this structure should be consistent with the regular and predictable approach.

In pushing towards this objective, this paper makes some important contributions to the current literature on debt structure. The existing literature mostly casts the question of optimal debt issuance in a static manner, solving for a single "optimal mix" or static issuance fractions. We believe this is too

. . .

1. The objective function in this paper reflects the traditional view of optimal debt management. Greenwood et al (2016) note that an optimizing government may also have other objectives, including using the structure of the debt to affect macroeconomic outcomes, as was done in the Federal Reserve's quantitative easing program, and promoting financial stability.

limiting, in that it offers little practical guidance as to how a debt manager should react to market conditions. For instance, very low yields or very flat curves can offer attractive opportunities to extend debt maturities, but models that solve for static solutions say little in this regard. It is arguably the case that a dynamic strategy that is responsive to issuance conditions will outperform a static strategy of targeting any particular static mix. We create a framework that addresses this limitation, solving for an optimal response function rather than a static optimal mix.

The paper first describes the underlying model of economic conditions, fiscal outcomes, and interest rates (section 2). It then uses simulations of that model to assess the tradeoffs that the debt manager faces under static strategies of setting a given distribution of issuance across maturities (section 3). The full optimization problem, with a focus on dynamic issuance strategies that shift issuance patterns in response to economic and financial variables, is then explored (section 4). These results are then used to assess some recent patterns in debt issuance (section 5) and the sensitivities to alternative model assumptions (section 6), with some concluding thoughts then offered (section 7).

Our hope is that this framework both provides some useful insights into current debt management decisions and spurs additional research that can advance our understanding of the issues involved.

2. Description of the model

In specifying a model to be used for assessing the debt structure, our goal was not to innovate in describing the dynamics of the economy and the yield curve. Instead, we wanted to use well-established modeling frameworks and apply them to the debt management problem. We prefer to use “off the shelf” estimates for macroeconomic relationships where possible.

The core of our simulation model has four blocks: the macroeconomic block, the rates block, the fiscal block, and the debt dynamics block. The purpose of starting with a macroeconomic block is to ensure that the model is rich enough to capture important stylized facts, such as recessions being associated on the one hand with higher funding needs, but on the other hand with lower funding cost for short-term debt. We think that these correlations are important to consider in thinking about optimal debt structure.

The following figure shows a high-level overview of the four blocks and their interlinkages. We will offer a more detailed description of each block in subsequent subsections.

Figure 1. Model overview

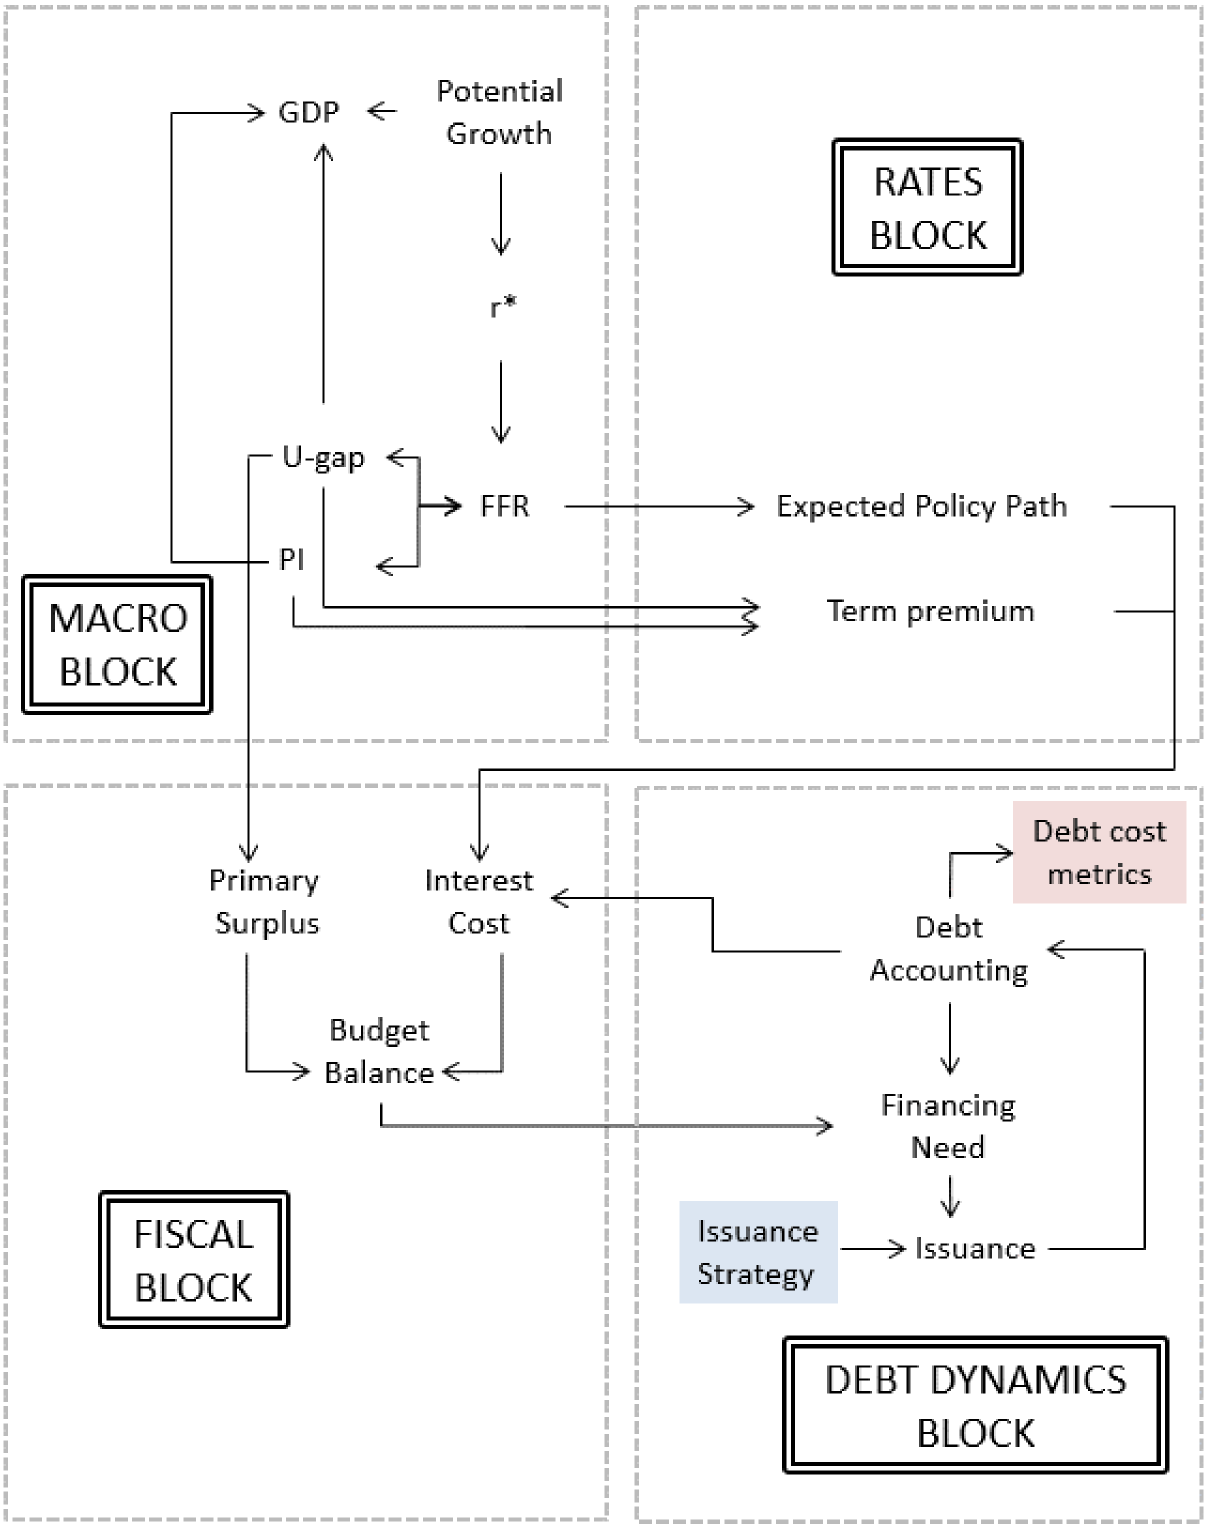

2.1 Macroeconomic block

The macroeconomic block is grounded in a standard three-equation macro model: an I/S curve relating the unemployment gap (*UGAP*) to its own lags and the stance of monetary policy; a Phillips curve relating inflation to its own lags, inflation expectations, and the unemployment gap; and a monetary policy rule relating the stance of monetary policy to unemployment and inflation gaps. Specifically, we use the I/S curve and Phillips curve equations reported in Rudebusch and Williams (2014), and reproduced below, where *FF* is the federal funds rate, *PI* is actual inflation, PI^E is expected inflation and R^* is the natural rate of interest.[2]

$$UGAP_t = 1.57 * UGAP_{t-1} - 0.62 * UGAP_{t-2} + 0.028 * [0.5 * (FF_{t-1} - PI_{t-1} - R^*_{t-1}) + 0.5 * (FF_{t-2} - PI_{t-2} - R^*_{t-2})] + \varepsilon_{UGAP,t}$$

$$PI_t = 0.58 * PI_{t-1} + 0.26 * PI_{t-2} + 0.16 * PI_t^E - 0.133 * UGAP_{t-1} + \varepsilon_{PI,t}$$

The unemployment gap is measured as the unemployment rate relative to its full-employment level. Inflation in our model refers to the quarterly annualized percent change in the core PCE price index. The standard deviation of shocks to the I/S curve is 0.24 and to the Phillips curve is 0.79.

We use a monetary policy rule that has been referred to as a "balanced approach" rule, with an inertial coefficient of 0.85. This policy rule – often described as an inertial version of the Taylor (1999) rule – is a fairly standard assumption to use; for example, Laforte and Roberts (2014) note its use in the Federal Reserve's FRB/US model. We do not allow an additive error term ("policy shocks") in this equation.

$$FF_t = 0.85 * FF_{t-1} + 0.15 * [R_t^* + PI_t + 0.5 * (PI_t - 2) - 2 * UGAP_t]$$

R* is a stochastic variable in our model, which we think is important to allow sufficient uncertainty around long-run steady-state outcomes.[3] Specifically, we start with the assumption that R* is the sum of potential growth and a residual "Z":

$$R_t^* = G_t + Z_t$$

We model G as a random walk and Z as an AR(1) process, reflecting temporary headwinds. We calibrate the quarterly standard deviation of shocks to the G process to be 0.0624, from drift in the twenty-year moving average of real GDP growth. We calibrate the parameters of the Z process so that it averages -50 basis points in the long run, deviations from the long-run steady state have a half-life of two years, and the overall volatility of R*matches that of the 20-year moving average of the real federal funds rate. Specifically:

$$Z_t = (1 - d) * Z_{SS} + d * Z_{t-1} + \varepsilon_{Z,t}$$

where d = 0.917 and the standard deviation of shocks to the Z process is 0.018.

. . .

2. Rudebusch and Williams equations are estimated on quarterly data since 1960.

3. The Rudebusch and Williams I/S curve which we use in our model is not estimated with a time-varying R*, which we acknowledge introduces a slight inconsistency.

We also track GDP in the macroeconomic block, as we will use it to scale the debt burden and debt cost in the results below. The level of GDP can be calculated from the path of potential GDP and the unemployment gap, using a standard Okun's law assumption that the output gap is twice the unemployment gap. We do not allow for any wedge between the PCE price index and the GDP deflator.

Inflation expectations are fully anchored at the 2% target for all simulations in this paper, although our model can also incorporate other assumptions in which expectations are less anchored, such as adaptive expectations. We impose an effective lower bound on the federal funds rate of 0.125%. The monetary policy rule is not adjusted in any way for unconventional monetary policy measures.

2.2 Rates block

For modeling the yield curve, we decompose nominal interest rates into the expected short-term rate and a term premium and explicitly model each component. Our model generates interest rates for bills and for 2-, 3-, 5-, 7-, 10-, 20-, 30-, and 50-year coupon-bearing securities.

We assume that the expected short-rate component embedded in longer-term rates is consistent with the dynamics of the model. At each time t in the simulation, we calculate an expected path for the federal funds rate extending out fifty years consistent with the equations in the macro block, taking account of starting conditions and assuming an absence of shocks. The expected short-rate component of the term rate of k years is the average of this federal funds rate path from t+1 quarter to t+ 4*k quarters.

For the term premium, we rely on the Adrian, Crump, and Moench (2013) term premium model that is commonly used by researchers and market participants. Moreover, we decided to anchor our term premium modeling on two points in the yield curve—the 2-year and 10-year maturities. By considering two separate points, we allow richer dynamics of the term premium across the curve.

Specifically, we first estimate equations for the 2-year and 10-year term premia (TP) by regressing them on rate volatility, inflation expectations, and the unemployment gap.[4] We introduce positive correlation in shocks to term premia between the two equations, in order to allow for the effects of other unobserved factors, by including the 10-year term premium in the specification for the 2-year term premium. The specifications are shown below:

$$TP_t^{10} = -3.59 + 0.207 * UGAP_t + 1.22 * PI_t^E + 1.75 * RV_t + \varepsilon_{TP10,t}$$

$$TP_t^{2} = -1.32 - 0.014 * UGAP_t + 0.477 * PI_t^E - 0.030 * RV_t + 0.420 * TP_t^{10} + \varepsilon_{TP2,t}$$

We assume that interest rate volatility, RV, persists at its average level from 2011 to the present. While we could use a measure of realized rate volatility generated endogenously within the model, we want to allow for the possibility that the typical level of volatility may change over time, and hence we use the more recent average, which has been lower than the full sample average. With this structure, the steady-state level of the term premium ends up being -5 basis points at the two-year maturity and 51-basis-points at the ten-year maturity.

Note that these equations do not assume any feedback effects from the supply of Treasury debt to term premia. One might suspect that overall Treasury debt supply could affect the term premium broadly, as investors might require a higher expected return to hold duration risk when they have to bear

. . .

4. In this estimation PI_E is in CPI rather the PCE price index terms. We make an adjustment for the spread between these measures.

more of it. In addition, one might wonder about temporary effects at particular maturity points when issuance of a particular security is increased sharply. Neither of these effects are included in the baseline model specification (and would be captured in the residuals).

The residuals from the term premium equations show persistence, as they capture other factors that are important for driving these measures. To capture this, we estimate that the residuals follow AR(1) processes with quarterly autocorrelation of 0.63 at the 2-year point and 0.73 at the 10-year point. The standard deviations of innovations to the error processes for term premia are 9 basis points and 41 basis points, respectively.

We also need to fill in the entire term structure of term premia, in order to determine the cost of issuance at all maturity points. To do so, we regress the Adrian, Crump, and Moench (ACM) term premium at each maturity point on the 2- and 10-year term premium measures, which allows us to interpolate (in a non-linear way) between them. For bills, we impose that the term premium is zero. For maturities beyond 30-years (for example if we want to consider a 50-year security), we extrapolate from our measures using the assumption that the convexity-adjusted term premium is linear in duration beyond the 5-year point—an assumption that seems consistent with the near-linear shape of the convexity-adjusted yield curve at longer maturities.[5]

After calculating the expected rate paths and term premia, we make a series of small adjustments to the rate levels to account for: (1) a bills / federal funds rate basis and (2) an on-the-run / off-the-run spread. The first adjustment allows us to include an extra premium on bills, which we set at eight basis points. The second adjustment is needed because the ACM term premium estimates are inferred from yields estimated from a fitted curve excluding on-the-run securities, while new issuance in contrast reflects an on-the-run premium.

2.3 Fiscal block

The main variable to forecast in the fiscal block is the primary budget balance as a share of GDP (*PRI*). This variable is strongly related to the business cycle, given the cyclicality of both tax receipts and expenditures in the presence of "automatic stabilizers." That relationship can be captured with a very simple equation:

$$PRI_t = 0.34 - 1.50 * UGAP_t + \varepsilon_{PRI,t}$$

We explicitly model persistence in the errors of the budget equation, with a quarterly autocorrelation of 0.92 and a standard deviation of innovations to the error process of 0.35. The primary budget balance is expressed at an annual rate.

We adjust the primary budget balance equation to approximate the CBO's projections for debt/GDP over a ten-year window by including add-factors. Beyond that horizon, we impose sufficient adjustment of the primary budget balance to create a gradual stabilization in debt/GDP in our baseline projections, in

. . .

5. We first add the convexity adjustment to our term premium measures, and then we extrapolate using the slope between the 5- and 10-year measures to arrive at the convexity-adjusted term premium for longer securities. We then subtract the convexity adjustment to get the actual term premium measures. The convexity adjustments are median values from a survey of a handful of market participants. Specifically, the adjustments are 1, 10, 30, 51, and 81 basis points at the 5-, 10-, 20-, and 30-, and 50-year points, respectively.

order to avoid an explosive fiscal baseline which would add an additional layer of complexity to interpreting the long-term model results.[6]

2.4 Debt structure dynamics and issuance patterns

The remaining equations are accounting identities: calculating the total budget balance from the primary balance plus the interest cost on the stock of debt; calculating the financing need from the total budget balance and rollovers of maturing debt; and updating the stock of debt with new issuance – which occurs at market interest rates determined in the rates block – in each period.

We do not make any adjustment for non-deficit financing needs, and we assume that methods of financing the deficit other than issuance of securities are trivial for the purposes of our analysis. We do not constrain gross issuance to be positive in the model, which means that we are implicitly allowing buybacks of debt. However, this outcome is infrequent in our simulations, as it requires very positive surprises on budget surpluses around the baseline assumptions. We do not make any explicit adjustment for the interaction between the Federal Reserve's balance sheet policy and the Treasury Department's effective financing cost through interest income remittances.

It is crucial for the model to accurately track debt dynamics — not just in terms of aggregate debt amounts, but in terms of the distribution of debt maturities at any given time. The model does so by tracking this distribution at quarterly maturities out to the longest issuance point. This accounting allows us to track quantities such as the weighted average coupon (WAC) and weighted average maturity (WAM) of debt over time, as well as debt sustainability ratios such as debt/GDP.

Of course, the pattern of debt that emerges will depend on the maturity distribution of issuance that is assumed in the model. Much of the remainder of the paper will involve changing the assumed issuance distribution and running simulations of the model to see the implications for borrowing costs and the evolution of debt, in order to assess various issuance strategies. These exercises will be explained in more detail in subsequent sections. At this point, however, it is worth mentioning some simplifications that we use to consider various issuance patterns.

The model could potentially allow the Treasury to set issuance at every maturity node, subject to an overall constraint that total issuance has to be sufficient to meet overall financing needs in each quarter. The maturity points that we consider in the model are one-year bills and 2-, 3-, 5-, 7-, 10- and 30-year coupon securities. Note that, in this version of the model, we are not considering TIPS. We also do not consider floating-rate notes, but in effect we can lump their issuance into the bills category since they have effectively zero duration.

In calculating the optimal issuance strategies for the simulations considered below, we could allow Treasury to choose its issuance at every single node. However, that would result in an excessive number of free parameters, with the Treasury effectively choosing 6 issuance levels in every quarter (the 7 issuance nodes, less 1 degree of freedom for the overall funding need constraint). If we were to extend the model to other types of securities, such as TIPS or 50-year securities, the dimensionality problem would become even more severe.

Moreover, it is not clear that we need to allow that degree of flexibility across issuance points, as the gross issuance profile rarely exhibits highly discontinuous variation across tenors, perhaps because yield levels and term premia tend to be relatively continuous across maturity. Therefore, we may reasonably

. . .

6. At the time of this writing, the CBO had not yet updated its ten-year projections to reflect the tax legislation enacted at the end of 2017 or the spending legislation enacted at the start of 2018, although updated projections later became available.

expect that the optimal issuance amounts will also show some structure across maturities, and not produce random issuance profiles.

To reduce the dimensionality of the issuance choices facing the Treasury, we require the actual issuance at each time step to be constructed as a linear combination of several elementary issuance profiles. We refer to these elementary issuance profiles as *kernels*, and the weights that must be attached to each kernel in order to produce actual issuance notional amounts as kernel loadings.

The characteristics of the kernels in turn are shown in Figure 2. We chose kernel 1 to reflect a reasonable baseline issuance profile, which is calibrated roughly from average issuance levels in recent years. Kernels 2, 3 and 4 were chosen to correspond to reasonable notions of shifting issuance into bills, into intermediate maturities (which we loosely call the "belly" of the curve), and into the long end, respectively.

Under this approach, a wide range of issuance patterns can be described as a simple linear combination of these kernels, with the issuance decision now described by four parameters—the weights on the four kernels. Figure 3 illustrates this idea by decomposing actual issuance into a linear combination of kernels we have chosen. Note that kernels 2 to 4 are defined to sum to zero across maturity points. As a result, we collapse the constraint of meeting the Treasury's funding requirement into the first kernel. With that kernel loading pinned down, the debt management decision essentially involves three parameters at each point in time—representing how debt managers want to deviate from the baseline issuance profile by shifting issuance towards bills, the belly, or the long end.

Figure 2. Describing Treasury issuance with four kernels

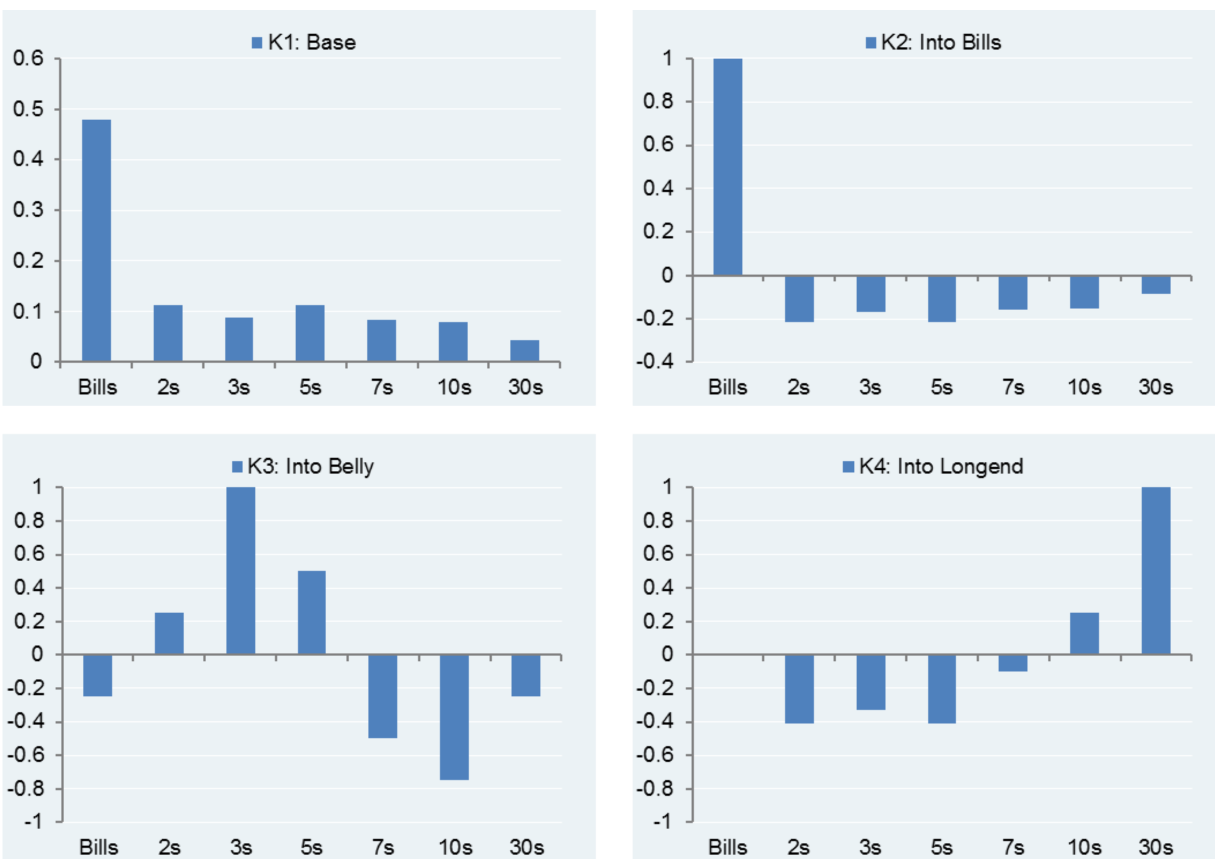

Figure 3. Representative issuance pattern described by kernels

Tenor	Final Issuance Profile			Baseline Profile		Shift into Bills		Shift into Belly		Shift into long bonds
<1Y	87.5			0.48		1.00		(0.25)		-
2Y	38.9			0.11		(0.21)		0.25		(0.41)
3Y	38.9			0.09		(0.17)		1.00		(0.33)
5Y	41.6	=	250 x	0.11	- 30 x	(0.22)	+10 x	0.50	-5 x	(0.41)
7Y	20.9			0.08		(0.16)		(0.50)		(0.10)
10Y	15.9			0.08		(0.15)		(0.75)		0.25
30Y	6.1			0.04		(0.09)		(0.25)		1.00
Sum	250.0			**1.0**		**0.0**		**0.0**		**0.0**

2.5 Baseline projections

We are now armed with a model that can be used to generate simulations of the economy, the yield curve, and the amount and characteristics of outstanding Treasury debt under any assumed parameters for the debt issuance strategy. We will use this model to assess various issuance strategies in detail over the next two sections. However, before leaving this section, we show some baseline projections from the model in order to highlight a few of its characteristics.

Figure 4. Model simulations of economic and financial variables over next 20 years

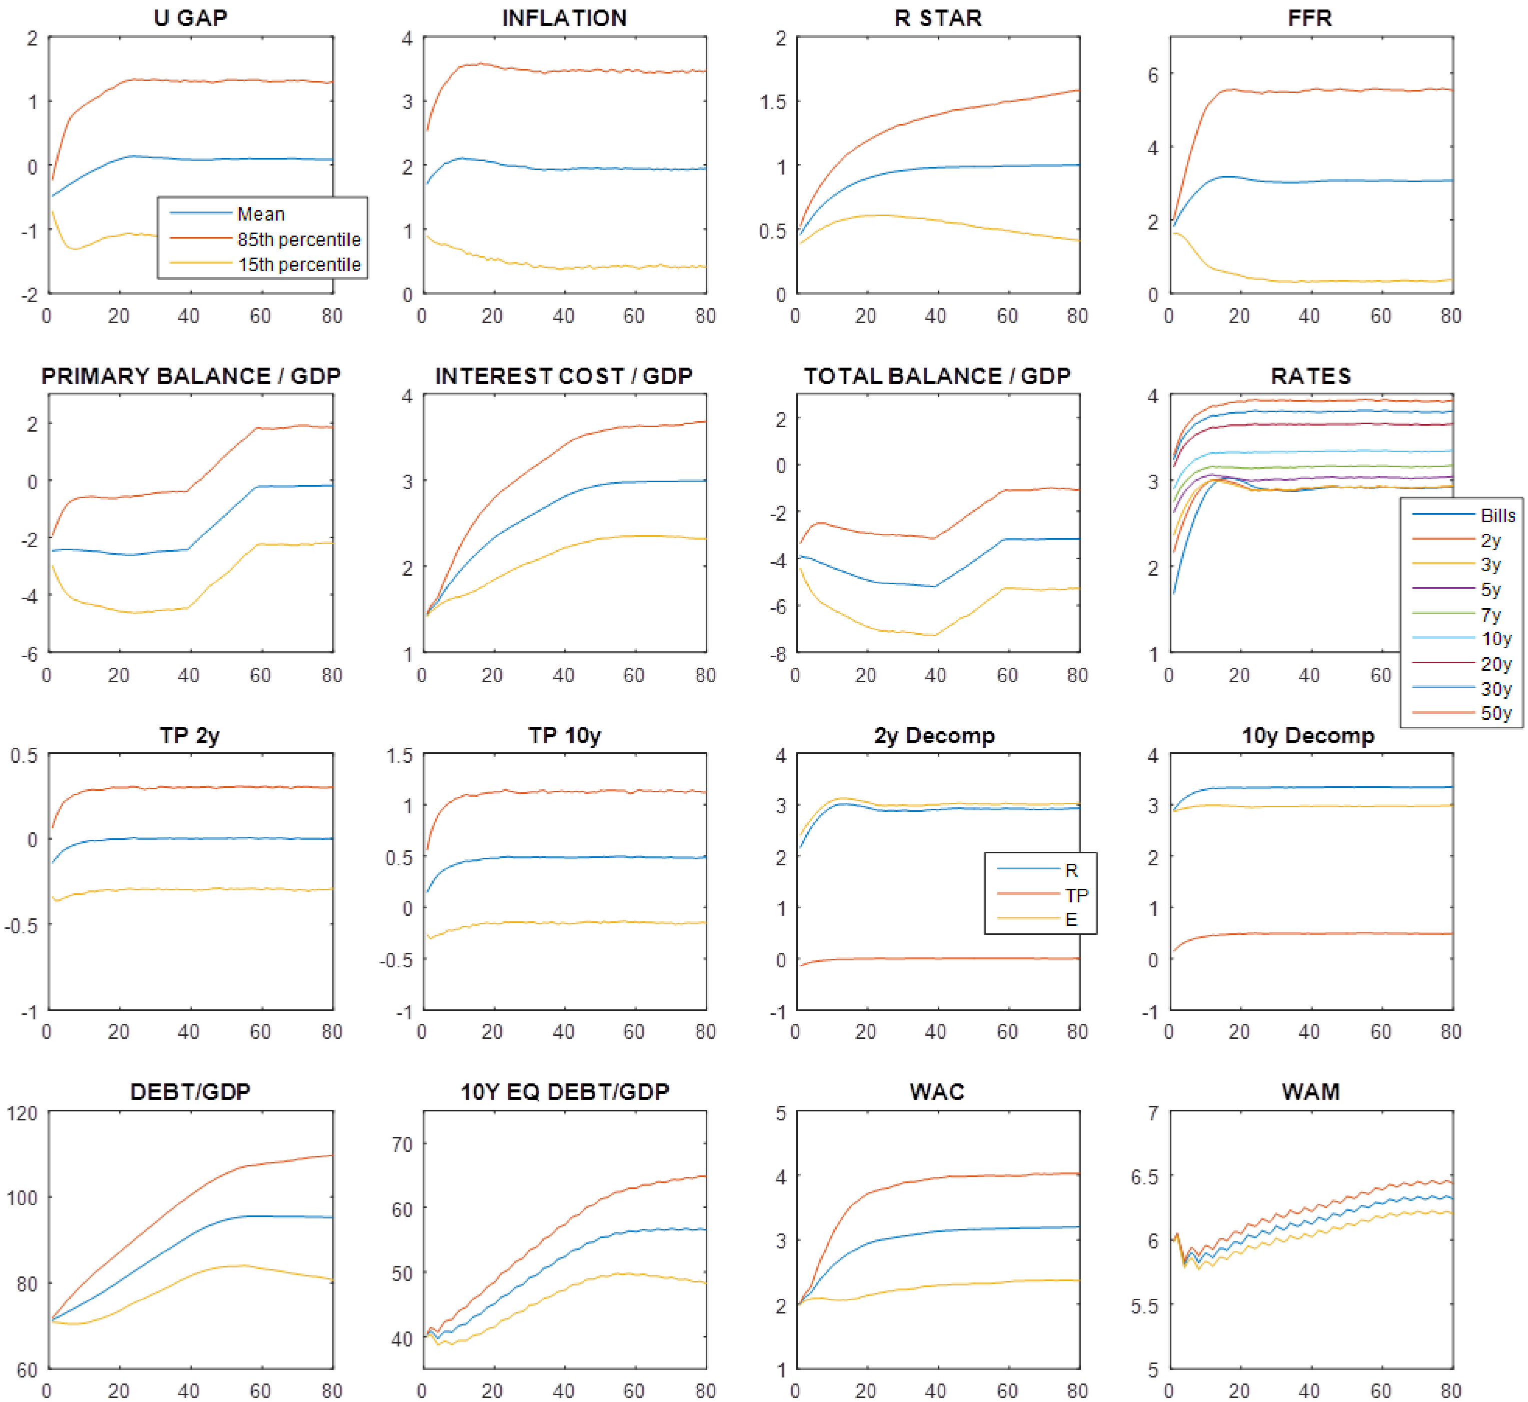

Figure 4 shows baseline results from simulating our model across 10,000 alternative future paths, conditioned on information available as of early 2018. We set starting residuals in the term premium equations to line up the model-implied yield curve roughly with the observed yield curve in March 2018, and we assume an issuance strategy that is similar to historical issuance, using the baseline kernel (K1) described above. The horizontal axis of the charts represents quarters from the present. 85th/15th percentile bands – roughly corresponding to a one-standard deviation interval – are shown to illustrate the variance of simulated values around their means.

A few points are worth highlighting about the baseline results. First, the macro variables in the model, including the unemployment gap and the inflation rate, strongly revert towards their long-term values. That is, cyclical dynamics largely play out over five years and are almost completely absent after ten years, which in part reflects the effectiveness of the monetary policy response in bringing those variables back to their targeted levels.[7] Second, although term premia mean-revert fairly quickly, significant long-run variation remains. Third, as mentioned above, we impose fiscal sustainability in the long-run by forcing the primary budget balance toward zero after the ten-year window, which stabilizes the debt-to-GDP ratio at about 85%. Given the average level of yields, the Treasury ends up paying interest expense that is just over 3 percent of GDP in steady state.

3. Optimization under static issuance strategies

In this section we begin to look at the properties of the model and to consider how the debt manager could optimize his or her issuance strategy. In doing so, we restrict ourselves to "static" issuance strategies, in which the debt manager is determining a single distribution of issuance across maturities that will be maintained throughout time, with the amount of issuance being scaled up or down proportionally to meet funding needs. The actual optimization problem for the debt manager would also involve their ability to change issuance patterns in response to economic or financial conditions—an approach that we consider in the next section.

3.1 Defining the frontier of potential outcomes

Even static issuance strategies can tell us a lot about the trade-offs that a debt manager faces. In particular, we can vary the issuance strategy in our simulations and see the effects on different metrics that are of interest to debt managers.

To begin, in order to illustrate the properties of the model, we consider a set of approaches that are even simpler than the issuance kernels described in the previous section. Specifically, we start by assuming single-issuance strategies, where all of the Treasury's issuance is concentrated at a single maturity point, without regard to the operational challenges this would involve or the market consequences it would have. These strategies are clearly unrealistic, as debt managers universally acknowledge the benefits of maintaining liquid benchmarks across the curve, but the simulations serve to illustrate some important conceptual points about the model.

Two metrics that we think are intuitively important in comparing issuance strategies are the average ratio of debt service cost to GDP and the standard deviation of debt service cost to GDP. The left chart in Figure 5 shows a scatter of simulated cost/variance outcomes by issuance strategy, where the outcomes are measured at the twenty-year forward horizon. Each point in the chart represents a simulation of ten thousand paths. Moving south in the chart represents better outcomes in terms of minimizing average

. . .

7. Because of the constraint imposed by the effective lower bound, the unemployment gap is slightly positive, on average, even in the longer-term, while the inflation gap is slightly negative.

cost, while moving to the left represents more favorable outcomes in terms of minimizing the variance of funding costs.

Figure 5. Average debt cost/variability trade off under single-maturity issuance strategies

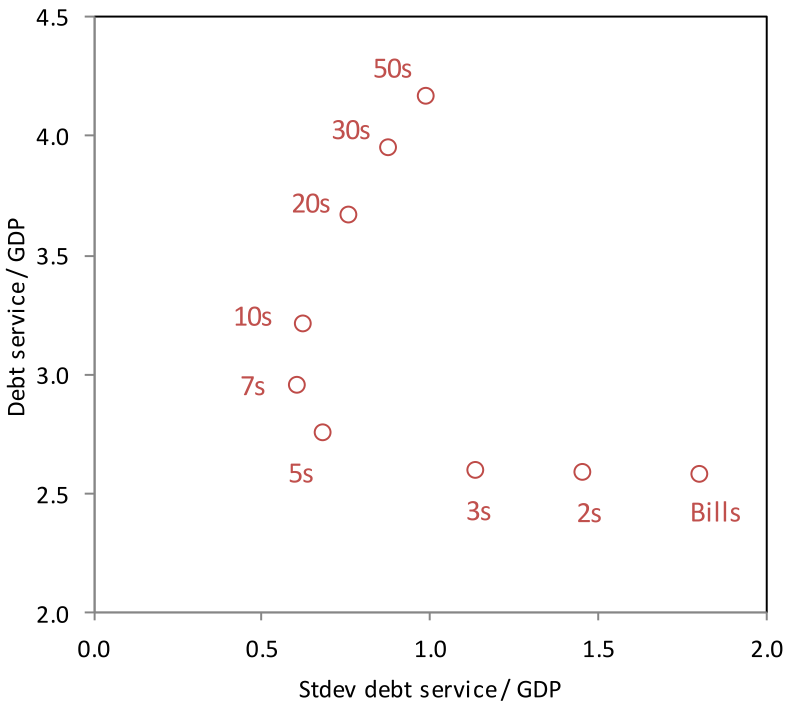

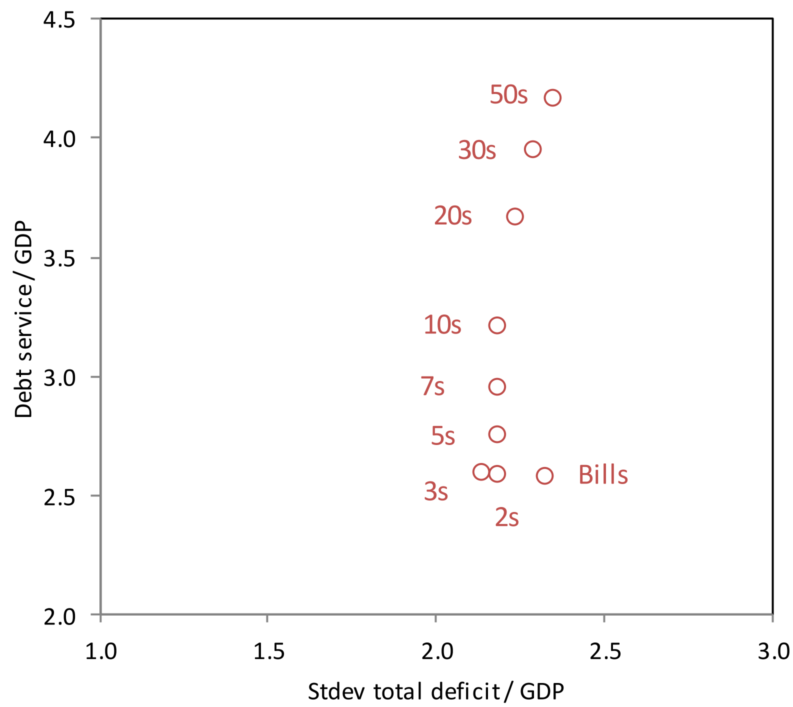

From the chart, it is clear that bills, two-year, and three-year issuance strategies have similar average debt service costs. This is intuitive as our empirical model has a near-zero steady-state term premium for two-year securities. Average debt cost then begins to rise gradually as issuance is pushed into longer maturities, reflecting the positive term premium the debt manager is paying, as well as the larger debt stock that accumulates from incurring that cost.

The variance of debt service falls considerably when issuance moves from bills to the five-year maturity point, as this shift reduces the sensitivity of the debt stock to interest rate changes resulting from the various shocks in the model. The reduction in volatility is substantial when measured in this way. However, the variance of debt service costs stops falling quickly beyond the five-year maturity point, and it begins to increase at maturities beyond ten years. There are a few contributors to the higher variance of debt service at longer maturities. Most importantly, the higher term premium associated with issuing at longer maturities results in a greater debt service burden over time and hence higher average debt levels, which results in greater variation in debt service costs in the face of shocks. Secondarily, the term premium varies considerably, which becomes a greater factor for issuance strategies tilted toward longer maturities.

One important consideration is that the debt manager may care not about minimizing the variance of the debt service itself, but the variance of the total deficit (debt service plus the primary deficit). That is because it is the total deficit that determines the amount of new borrowing needed, or the magnitude of the discretionary fiscal policy adjustments needed to avoid this borrowing.[8] During recessions, the primary budget balance tends to deteriorate, while the Federal Reserve tends to cut interest rates, introducing negative correlation between the average issuance rate for short-dated borrowing strategies and the size of the borrowing requirement. This correlation provides a benefit to issuing at shorter maturities, as can be seen in the right chart in Figure 5, which replaces variance in debt service with variance in the total deficit on the horizontal axis. In contrast to the results shown to the left, issuing at the two- and three-year maturity points now provides the lowest variation, with the bills-only strategy performing only modestly worse.

Figure 6. Average debt cost/variability trade off under kernel-based issuance strategies

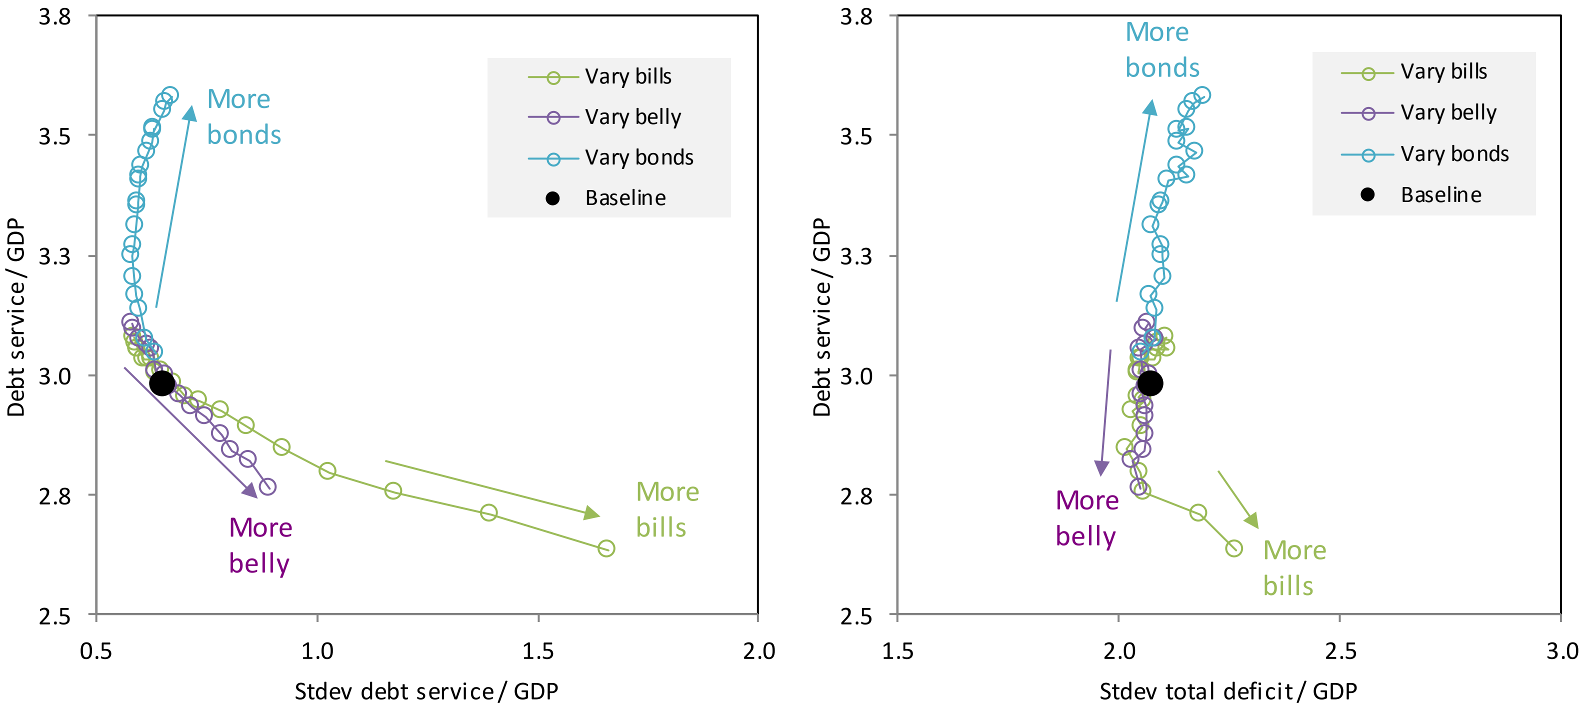

Of course, none of these single-issuance-point strategies would be practically feasible, so we turn to more realistic issuance strategies based on the issuance kernels described in the previous section. The results are shown in Figure 6. The black point represents the "baseline" issuance strategy captured by the first kernel. The green points vary the bills factor from -0.4 to +0.5, with the arrow representing the direction of greater bill share. The purple points vary the belly factor from -.08 to +0.1, and the teal points vary the bonds factor from 0 to +0.25 – again, with the arrows indicating the direction of a greater belly and bond share. The bounds on the factors are chosen to ensure that issuance amounts are strictly positive.

The key point from the single-issuance strategy simulations also appears in the more realistic simulations: Issuing debt at intermediate maturities, particularly 2-, 3-, and 5-year maturity points (as captured by the "more belly" kernel), looks particularly attractive in the trade-off between average debt

. . .

8. Indeed, this broader measure of risk better aligns with the debt manager's objective of improving social welfare by smoothing tax payments rather than debt servicing costs. See Barro (1979), Lucas and Stokey (1983) and Greenwood et al (2015).

cost and the variability of debt service or deficits. Skewing issuance too much towards longer-dated securities appears to raise costs without substantial benefits in terms of reducing variability, while skewing issuance too much towards bills leaves the debt manager exposed to too many shocks to funding costs.

We want to more explicitly consider the optimization problem that the debt manager might perform in these circumstances. To do so, we need to know the full set of possible tradeoffs between average debt cost and variability. We can of course consider any particular issuance strategy that we want and assess how it performs in the model simulations. In Figure 7, we add to the previous results a set of strategies that blend together the issuance factors (orange points) in order to fill in a cost/variance frontier more completely. The lower envelope of these points (gray line) represents the lowest feasible debt service cost attainable for a given variance in debt service.[9]

Figure 7. The frontier for average debt cost/variability tradeoff under static issuance strategies

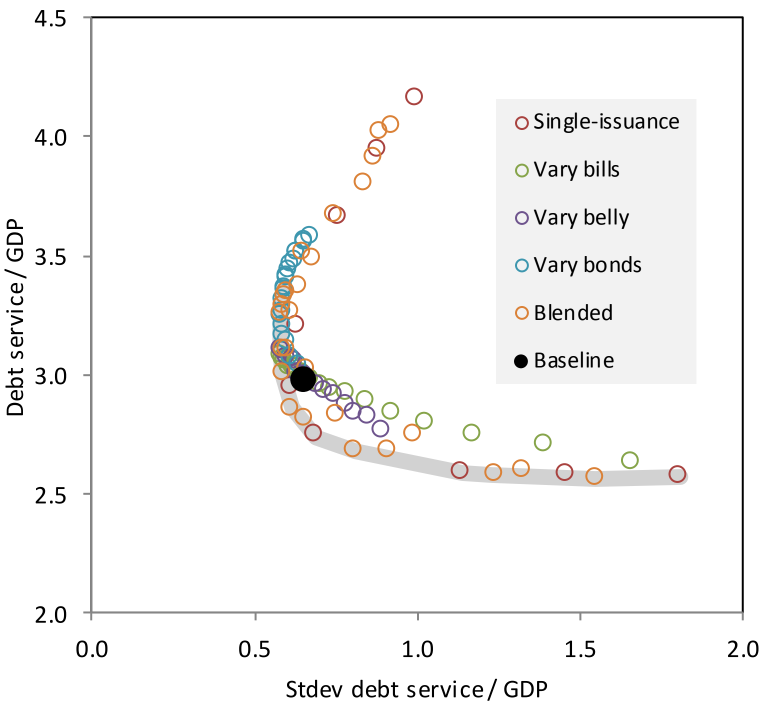

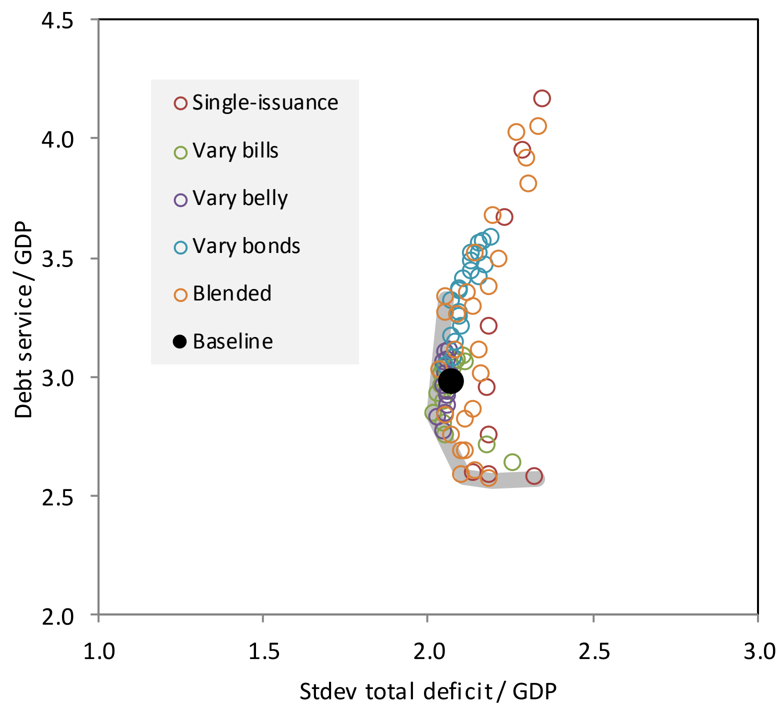

3.2 Maximizing the objective function of the debt manager

Armed with that frontier, let's consider the optimization problem for the debt manager. As noted above, the debt manager likely focuses on minimizing his or her expected funding cost and limiting the variation of either this funding cost or the budget. There could also be additional terms in the objective function. The Treasury has in the past discussed the desire to maintain a liquid market across a range of maturity

. . .

9. Again, we think that some of the points on the efficient frontier are likely not sufficiently diversified to represent realistic debt management choices in light of other considerations, but our analysis highlights some important tradeoffs and stylized facts for debt managers to consider.

points, and providing those benchmarks to the market could be important. In addition, we could incorporate the desire to be regular and predictable directly into the objective function, by using a cost term associated with changes in issuance patterns. For the moment, we omit these additional terms and simply focus on the cost terms in the objective function.

We assume that the debt manager has a loss function that is linear in expected debt service cost as a share of GDP and the variation of that cost (or of the total deficit) as a share of GDP, where both are measured at a given point in the future. Specifically, we assume that the loss function of the debt manager is given by $E_t[DebtService_{t+n}/GDP_{t+n}] + lambda * StDev_t[DebtService_{t+n}/GDP_{t+n}]$ or $E_t[DebtService_{t+n}/GDP_{t+n}] + lambda * StDev_t[Deficit_{t+n}/GDP_{t+n}]$, where the parameter lambda captures the degree of risk aversion of the debt manager, and n is the horizon considered. Those preferences can be represented as linear indifference curves, as shown in Figure 8. The debt manager will then choose the point on this frontier at which the indifference curve is tangent to the efficient frontier. A risk-neutral debt manager (with horizontal indifference curves) will tend to skew issuance more towards the front end. A risk-averse debt manager that is concerned about variation (with sloped indifference curves) will end up issuing a greater amount of intermediate maturities, in order to achieve more effective risk reduction.

Figure 8. Optimal issuance patterns under different preferences for the debt manager

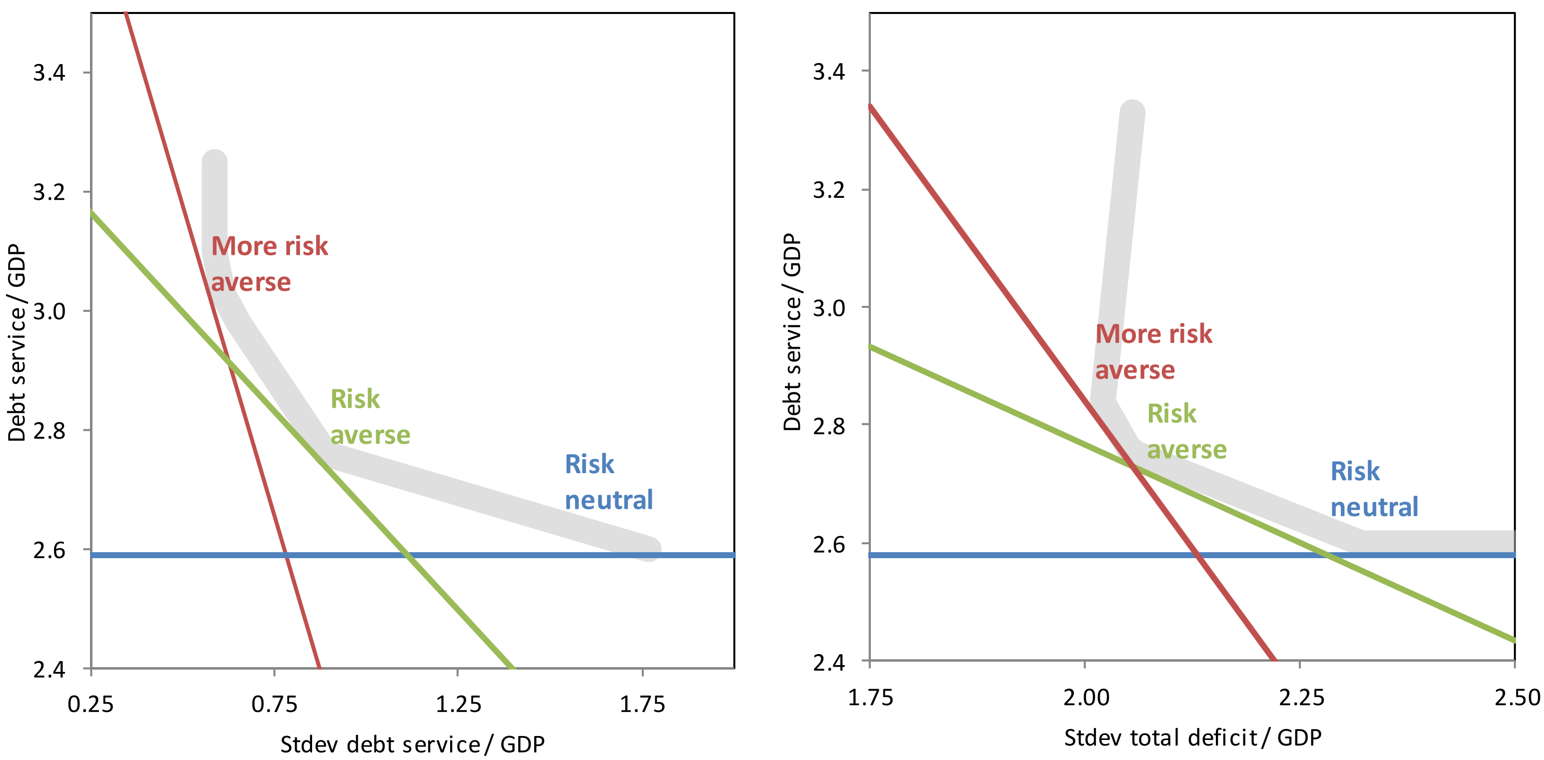

Figure 9 provides some results on the optimal strategy implied by the model, using a horizon of 20 years and making various assumptions about the risk aversion parameter of the debt manager. Under an assumption of being risk neutral (lambda=0), the debt manager would choose an issuance pattern that relies almost entirely on short- and intermediate-term maturities, for the reasons discussed above. Even under a moderate degree of risk aversion (lambda=2/3), the debt manager skews issuance heavily in that direction, producing a larger concentration of debt with maturities of 5 years or less than observed in the current debt stock.[10] As can be seen, the WAM and truncated WAM (TWAM) of the debt is meaningfully lower under the optimal strategy than the recent issuance strategy.[11] The only situation that produces a longer WAM is when the debt manager is focused on debt service cost as the risk measure and has a high degree of risk aversion.

Figure 9. Issuance patterns and debt statistics under different preferences for the debt manager

Risk aversion	Debt metrics		
Lambda	WAM debt	TWAM debt	<5y share
Risk measured by debt service:			
0	0.81	0.75	98.8%
2/3	4.18	2.97	84.1%
2	7.26	4.90	56.6%
Risk measured by total deficit:			
0	0.81	0.75	98.8%
2/3	4.18	2.97	84.1%
2	4.86	3.37	73.6%
Memo: Current debt	5.83	3.89	67.5%

Before moving forward, it is worth making a few additional points about the risk measure from the optimization problem. As we have noted throughout the discussion, the decision of whether to measure risk in terms of the variation of funding costs or the variation in the budget has a meaningful impact on the optimal debt strategy. In addition, the literature on debt management also allows for different ways to measure risk for a given choice of the variable, which further complicates this issue. For example, rather than focusing on the standard deviation, some papers assume that the debt manager is only concerned about variation that is unpredictable from a year in advance of the horizon considered (the conditional variation) or about particularly bad budget outcomes (a measure corresponding to the tail of the distribution). However, our results do not appear highly sensitive to the different ways of measuring risk for a given variable, as the measures are all quite correlated. (See appendix B for more details.) Instead, it is the choice of the variable itself (debt service cost or budget) that is more relevant for shaping the results.

. . .

10. In these results, the optimization settles on the same set of kernels under the two risk measures when lambda is set to 0 or -2/3, but that will not be the case at other values of lambda.

11. We are characterizing the debt with a wider set of statistics than just the WAM, as suggested in a 2017Q4 TBAC presentation (https://www.treasury.gov/resource-center/data-chart-center/quarterly-refunding/Documents/Q42017CombinedChargesforArchives.pdf) . The truncated WAM measure treats all debt with maturity greater than 10 years as having a 10-year maturity.

All of the above results assume a static issuance strategy, in which the distribution of issuance across maturities is set at the beginning of the simulation and does not vary with economic conditions. It would not be surprising if the debt manger could perform even better by incorporating some responsiveness into the debt management strategy. In the next section we turn to a more complete optimization problem for debt managers that allows for a dynamic reaction function in which the pattern of issuance varies in response to a small set of economic conditions.

4. Optimization under dynamic issuance strategies

In this section we turn to a more complete optimization problem for debt managers. In particular, we solve for optimal issuance strategies that involve a dynamic reaction function, in which the pattern of issuance varies in response to a small set of economic conditions. We believe that solving this optimization problem is a significant advance in the literature on debt management strategies.

4.1 Dynamic reaction function for the debt manager

A debt manager tasked with managing the average cost of the debt while also limiting fiscal risk might reasonably be expected to respond to market conditions, by increasing long-maturity issuance in periods where term premium is very low, for instance. To be sure, the U.S. Treasury's commitment to a "regular and predictable" issuance strategy might limit the extent to which it can respond to market conditions. But it is reasonable to suppose that a regular and predictable strategy does not mean a complete disregard for market conditions, and hence we can consider a response function that implements gradual changes to issuance strategies over time.

The desirability of a dynamic issuance strategy of course reflects the stochastic nature of the economic and financial environment. By choosing a formulation that solves for an optimal response function, we allow actual issuance patterns to adjust in response to this variation in conditions. This response function can be optimized to achieve the proper balance between expected funding costs and deficit risks.

This type of dynamic formulation could substantially improve upon the static issuance strategies considered in the previous section. However, it also greatly increases the computational complexity of the optimization problem. Here we provide a qualitative description of the optimization problem. (Appendix C reports more details of the formulation of the problem.)

We allow the debt manager response function to depend on three macroeconomic variables (MEVs) that are observable in the model: the level of real short-term interest rates, the term premium, and the size of the budget deficit. The response function, which does not vary with time or across paths, determines the loadings on issuance kernels 2 through 4 as functions of these variables. That is, after observing those three variables, the debt manager gets to decide his or her kernel loadings—the amount by which to shift issuance into (or out of) bills, the belly, or the long end—while still ensuring enough issuance to meet financing needs.

The parameters governing those responses can be optimized for an assumed objective function. As in the optimization problem from the previous section, we assume that the debt manager is focused on minimizing debt service costs with a penalty on a risk term. For the baseline set of results, we will use the standard deviation of the budget as the risk measure, and we set the risk parameter to reflect a low-to-moderate degree of risk aversion.[12]

The optimization problem assumes a linear reaction function with 12 parameters, governing the responses of each of the three issuance kernels to each of the three macroeconomic conditions (and a constant term for each kernel). We model this problem as a Linear Programming problem, detailed in Appendix C. While the number of reaction function parameters is small, the complexity in optimization stems from propagating this into each path and time step.[13]

The optimal reaction function coefficients resulting from a baseline run of our model are shown in Figure 10. The numbers in the table correspond to the coefficients of linear functions that relate kernel loadings for K2 to K4 (the rows) to the variables (corresponding to each column). As can be seen, and will be discussed in more detail below, the reaction function involves meaningful responsiveness of the debt manager to each of the three MEVs.

Figure 10. Optimal response function of kernels to MEVs ($bn)

Kernel	Constant	Real 2y Yield	10y Term Premium	Budget Deficit
K1: Base	Gross Financing Requirement			
K2: Into Bills	(279.4)	255.2	418.1	(0.6)
K3: Into Belly	70.6	(6.2)	70.3	0.3
K4: Into Long End	(42.3)	25.0	(5.9)	(0.1)

In the results below, we will compare this reaction function to the optimal static issuance strategy. One can think of the optimal static issuance strategy as a special case of the above reaction function in which the kernel loadings on the MEVs are restricted to be 0 (so that issuance is independent of the macroeconomic variables), in which case the optimization involves finding only the constant terms for each kernel.

4.2 Results under dynamic issuance strategy

We can look at what this reaction function implies for issuance, as presented in Figure 11. Issuance fractions implied by the dynamic response function are shown both for recent MEV values (as of 4Q17)

. . .

12. There are differences between the implementations of the static issuance optimization and the dynamic issuance optimization, which were mostly done to preserve linearity for reasons of computational tractability. Since standard deviations are nonlinear, the model uses linear proxies to implement risk constraints. When the model's internal risk penalty is recast as a preference between "traditional" risk and reward, it corresponds to a lambda parameter between 0 and 0.5 in the discussion above. Second, kernel loadings represent dollar amounts rather than percentages (which would become nonlinear); consequently, loadings are internally scaled by a factor that grows with time to account for the generally increasing amounts of gross issuance with time. This is detailed in the appendix. Third, the dynamic strategies used only 50 paths for both the optimization and the calculation of expected cost/risk tradeoff, whereas the static issuance strategies were run on 10,000 paths.

13. One aspect of our set-up is that the stochastic paths for the economy, budget needs, and the yield curve that are generated by the macroeconomic model discussed in section 2 are independent from the debt manager's reaction function. This recursive structure is critical for making the optimization problem feasible. If we were to introduce feedback channels from the reaction function back into the behavior of rates and the economy (as we consider later in the paper), the optimization problem would no longer be tractable.

and for the steady-state MEV values from the model. The static issuance fractions that are obtained under the assumed objective function are also shown for comparison.

The dynamic optimization results suggest an issuance mix for the current (2017Q4) MEV values that is similar to that derived from the static optimization, but this outcome is somewhat coincidental. If the MEVs were to move over time towards their longer-run levels, the issuance allocation would shift away from intermediate maturities and into a heavier allocation to bills.

Figure 11. Optimal issuance pattern and debt characteristics under dynamic reaction function

	Dynamic Response Function		
Gross Issuance	4Q2017 MEVs	Long run MEVs	Static
Bills	20%	39%	25%
2s	20%	16%	19%
3s	22%	20%	22%
5s	22%	18%	21%
7s	9%	5%	7%
10s	5%	2%	4%
30s	1%	0%	2%

The *effective* sensitivity of optimal issuance patterns to the three MEVs is shown in Figure 12. The coefficients in Figure 10 may be thought of as partial sensitivities of kernel loadings to the MEVs. But the overall impact of a move in (say) real short rates is the sum of the impacts from all kernels. To provide some intuition about the overall impacts of moves in the MEVs, Figure 12 lays out the change in issuance stemming from a one-standard-deviation move in each MEV.

Figure 12. Optimal response of issuance to MEVs (one-sigma move in MEVs); $bn

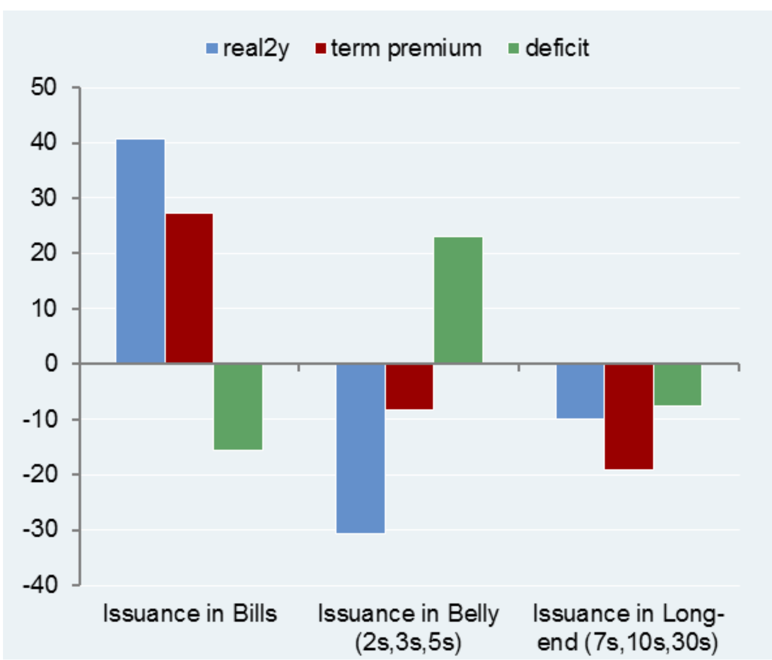

As the macroeconomic environment changes over time, optimal issuance will change according to these sensitivities. Not surprisingly, increases in term premium favor a rotation from all coupons into bills; increases in real 2-year rates generally favor moving from the belly of the curve (2s/3s/5s) as well as the long end into bills; and rising deficits result in increases in issuance in the belly offset by a reduction in bills and the long end. Collectively, these effective sensitivities to the MEVs clarify the reason for the increase in the reliance on bills as the MEVs move to their steady state levels, as both the term premiums and the real 2y yield are moving higher.

One should expect that the dynamic reaction function would allow debt managers to achieve lower expected cost for a given volatility level relative to a static issuance strategy. After all, the optimization problem could always choose to have no sensitivity to the economic and financial variables included in the reaction function. In Figure 13, we compare the average cost/risk outcome from the dynamic debt management strategy (pink diamond) to the static issuance strategies. The average cost under the dynamic strategy lies well below the efficient frontier generated by static optimization, indicating significant benefits to allowing the issuance strategy to vary with the macro environment.

Figure 13. Expected cost/risk tradeoff of dynamic issuance strategies relative to the static issuance frontier

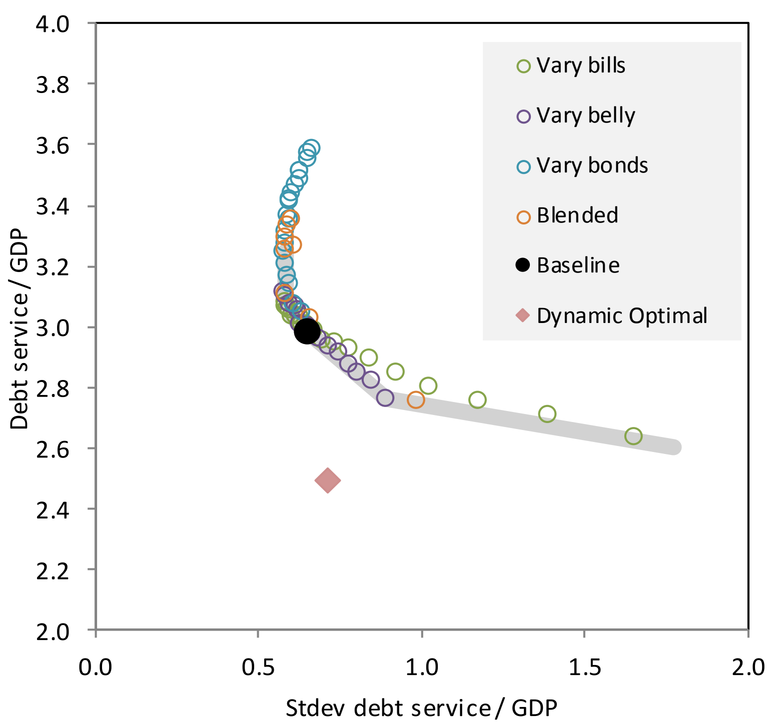

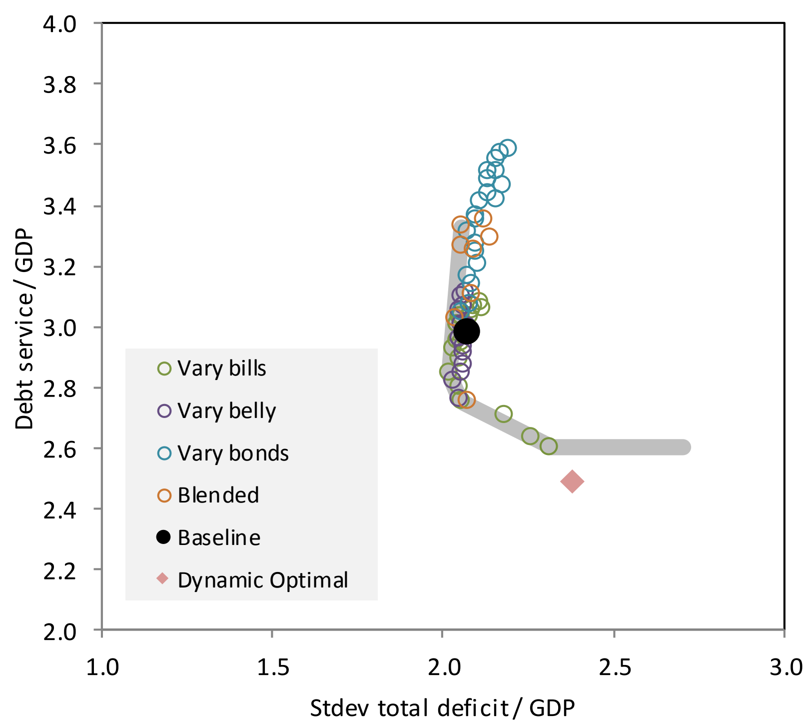

Lastly, we consider how the optimal issuance mix recommended today is affected by two choices related to the risk measure. The first issue is whether risk is measured in terms of variability of debt service costs or the budget deficit, as discussed above. The second issue is the size of the penalty placed on the risk term. The optimal issuance patterns that are obtained by varying these factors are shown in Figure 14.

As can be seen by comparing the two panels, the decision about whether to measure risk by debt service costs or the budget matters considerably, consistent with the results discussed under the static issuance section. Measuring uncertainty based on the budget pushes issuance from long maturities into short and intermediate maturities, given that they benefit more extensively from the correlation properties discussed above. Moreover, the differences get exaggerated as the budget penalty is ramped up.

Figure 14. Optimal issuance mix under two different measures of risk, varying the penalty on risk

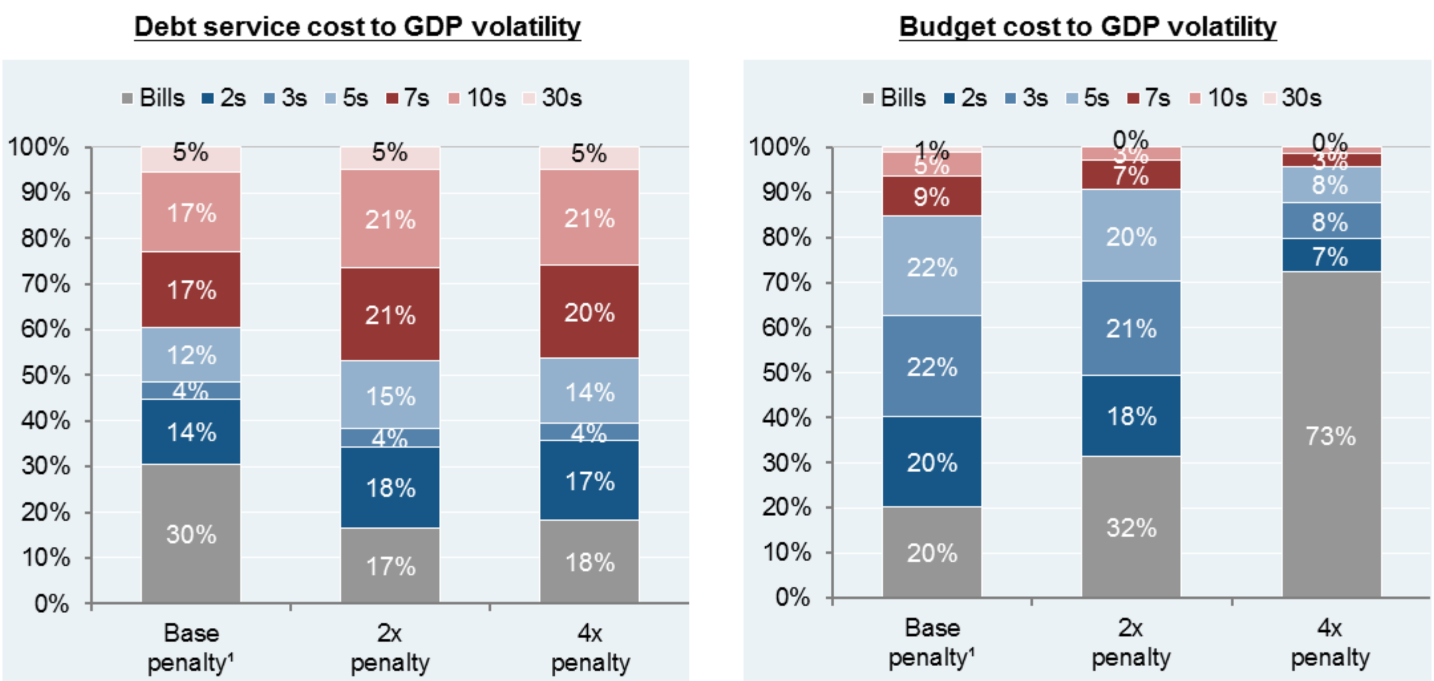

5. Assessment of recent debt issuance patterns

We can use the model to assess debt management decisions that were implemented over the past 10 years. In particular, we conduct two exercises in this section. The first looks at the specific decision to extend the weighted average maturity (WAM) of the debt that was made in 2009. The second looks at the general pattern of issuance over the past decade, taking into account the evolution of the economic and financial variables in the debt management reaction function.

5.1 Decision to extend the WAM

In 2009, the TBAC recommended lengthening the WAM of the debt, given that it had fallen to less than 5 years, the lower end of its historical range. Some lengthening would have occurred under existing issuance patterns, with the WAM likely to have extended to about 6 years over time.[14] However, the set of debt management decisions made from 2007 to 2015 raised projected WAM from 6 to 7 years. We can assess that shift in the context of the model. In addition, we can consider additional potential issuance changes that would raise the WAM by an additional year or by an additional 2.5 years. These further changes are assumed to be accomplished mainly through raising issuance levels in long maturities.

The results are shown in Figure 15. The maturity extension decision cost the Treasury about five basis points of GDP in terms of expected debt service cost in this model, or an average annual cost of around $10 billion. When variation is measured by the standard deviation of debt service costs, this change moved the debt structure closer to the frontier, and this trade-off would be attractive under a wide range of utility functions (basically all except completely risk neutral preferences). When variation is instead measured by the standard deviation of the deficit, the extension achieved basically no risk reduction, and hence the additional cost was incurred with no substantial benefit.

Figure 15. WAM extension in the context of the model

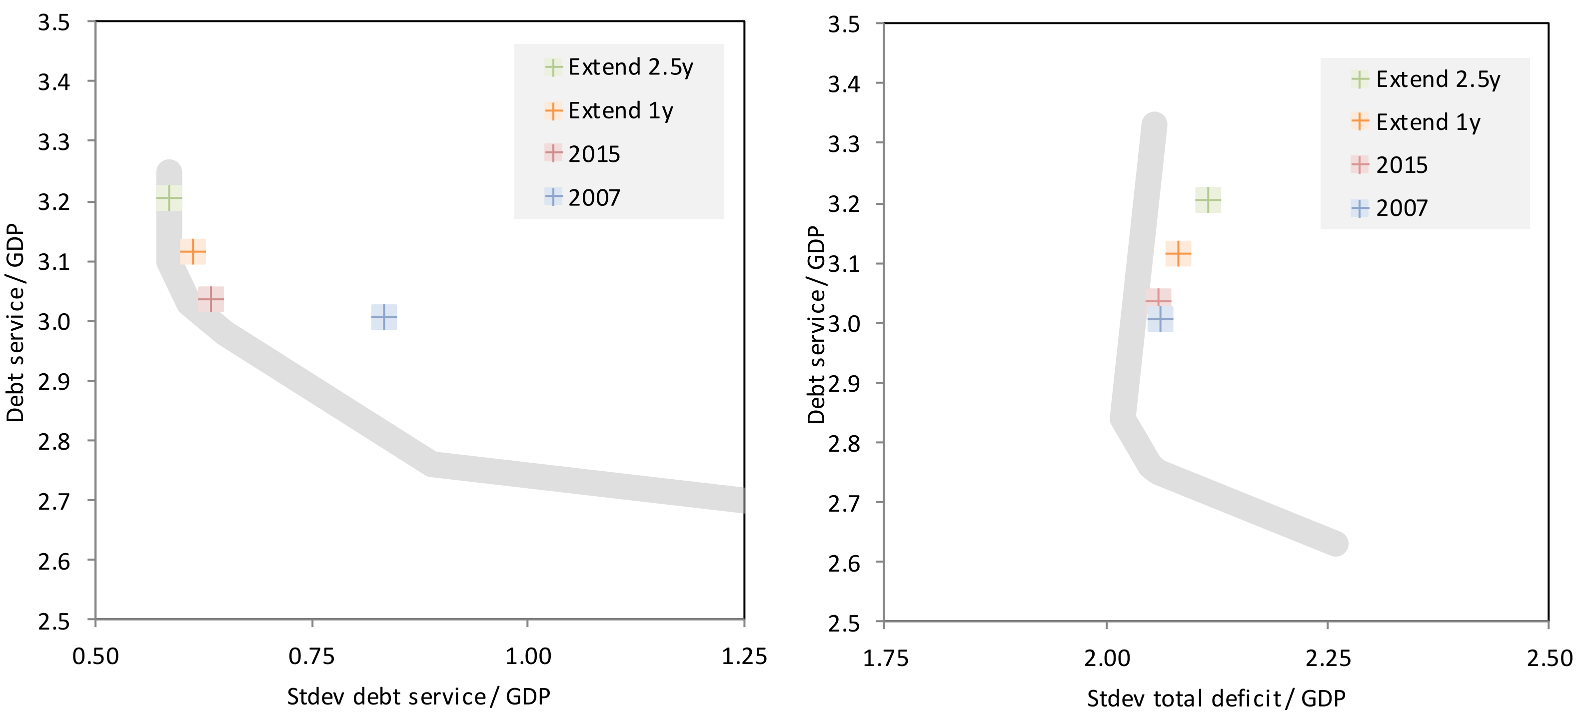

The charts also show the effects of further WAM extension from 2015 levels. Under the both measures, further WAM extension represents an increasingly less favorable trade-off. The cost continues to increase, and the variation goes down by less (left chart) or even begins to rise (right chart).

. . .

14. In this calculation, we use 2007 issuance patterns, to avoid some of the large shifts that occurred from unexpected funding needs during the financial crisis.

The model can also be used to calculate the cost of being off the efficiency frontier. Based on 2007 issuance, the Treasury could have achieved the same variation in debt service costs with an issuance pattern that would have reduced expected debt service cost by just over 0.2 percent of GDP, or about $40 billion per year based on current GDP. When variation is measured in terms of the deficit, the gap is slightly wider. This calculation illustrates the potential value that can be realized by developing a set of analytics to better optimize debt structure.

5.2 General issuance patterns over past decade

We can also use the model to make a broader assessment of debt management decisions over the past decade. Specifically, we compare the actual Treasury issuance patterns to those that would have been obtained from the optimal reaction function using historical values of the MEVs. The results, shown in Figure 16, highlight that the model has generally favored larger issuance at intermediate maturities and smaller issuance in bills and bonds, relative to actual Treasury issuance. That pattern was noted above based on the current MEVs, but this exercise shows that it has held over most of this sample.

The deviation from realized issuance patterns is most pronounced in recent years, as the combination of falling term premium, declining real front end rates, and (more recently) rising deficits implies significant benefit from reducing bill issuance in favor of intermediates. Indeed, today's optimal issuance mix suggests close to the largest allocation to intermediates of the last decade with the model recommending over 60% of annual issuance in 2s/3s/5s and less than 25% in bills.

Figure 16. Actual issuance patterns compared to the optimal dynamic issuance strategy

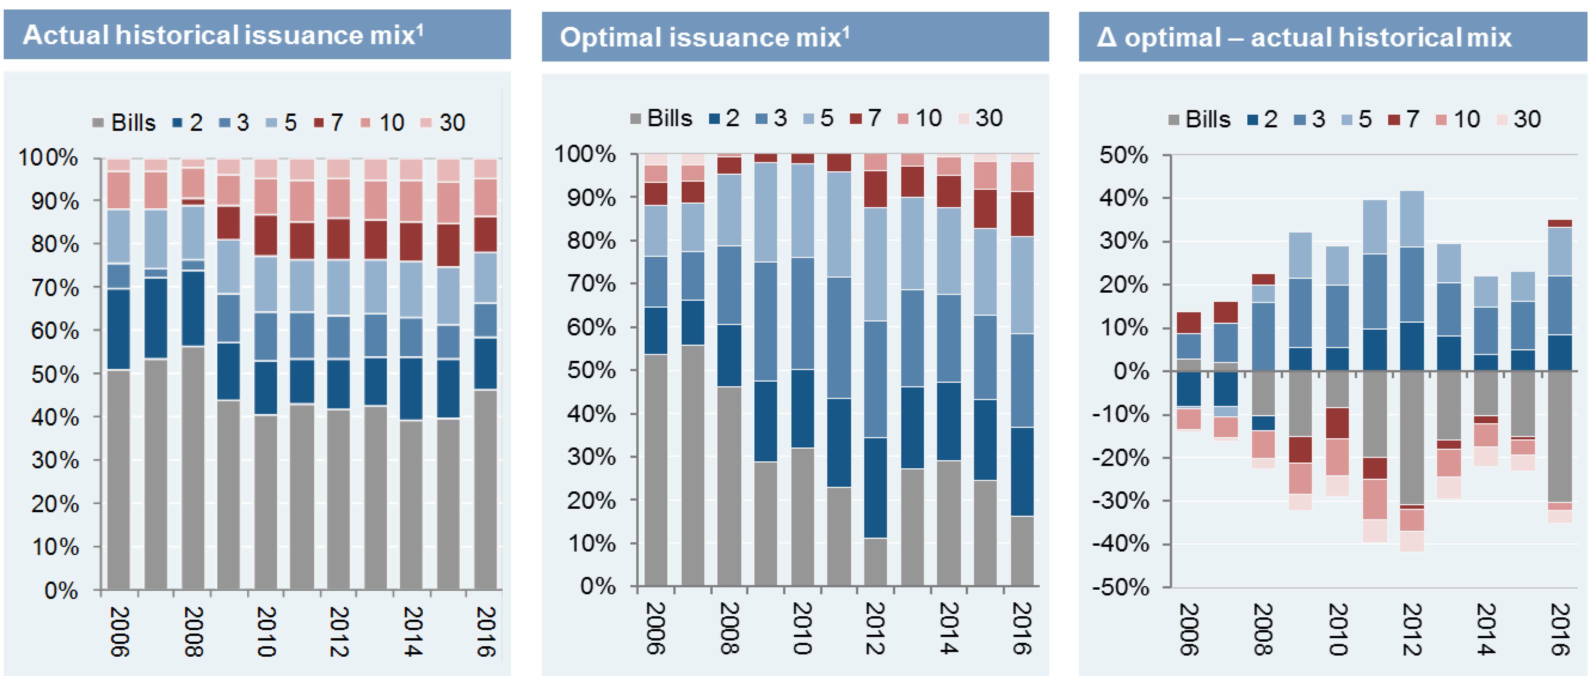

1. Bill issuance across varying tenors are all scaled to 1-year (52-week) tenor, i.e. 100B of 26-week Bills scales to 50B of 1-year equivalent Bills

6. Sensitivity of results to alternative model assumptions

One of the advantages of having a model is that we can change particular assumptions and assess how it affects the recommended debt structure. In this section we consider several variants of the model.

6.1 Alternative term premium assumptions

We begin by exploring the sensitivity of our results to different assumptions about the term premium, along two dimensions. First, we add feedback from the level of debt outstanding to the term premium. Second, we investigate what happens if the steady state levels of the term premium remain at their current low levels.

To add supply feedback to the model, we assume that the 10-year term premium rises by six basis points per percentage point of GDP increase in ten-year equivalent (TYE) debt outstanding. This assumption is consistent with estimates from Gagnon et al (2010) and Dawsey (2013), and is around the midpoint of the range of estimates collected in a literature review from Gagnon (2016). While there is considerable uncertainty about the magnitude of supply effects, the basic conclusions of our analysis would hold assuming alternative smaller or larger sensitivities.

To incorporate supply effects, we assume that current yield levels reflect expectations for our baseline fiscal path, and so assess ten-year equivalent debt outstanding as a share of GDP relative to its path in our baseline scenario.[15] In other words, a debt level the is 10pp above our baseline fiscal scenario (in terms of TYE) at the twenty-year forward horizon would raise the ten-year term premium by 60 bps at that time, relative to the steady-state level that would otherwise occur. We assume that the supply impact on the two-year term premium is proportional in duration to the impact on the ten-year term premium.

Figure 17 illustrates the results of adding supply feedback to the model, compared with our baseline simulations (previously shown in Figure 6). Strategies with increased belly share relative to bond share result in lower debt service costs than in our baseline simulations, because they result in lower ten-year equivalent amounts of debt outstanding and hence depress the term premium relative to the baseline steady state. Strategies which significantly increase the bond share of issuance result in explosive growth in the cost and volatility of debt service. A vicious cycle of higher debt outstanding feeds back to higher interest rates, which raises debt service costs and results in even greater debt levels in the future, which feeds back into even higher interest rates, and so on.

. . .

15. As a reminder, we do not assume an ever-increasing level of debt to GDP in our baseline simulations, and instead gradually bring the primary deficit back toward balance after ten years to stabilize debt dynamics by assumption.

Figure 17. Alternative frontiers incorporating supply feedback on the term premium

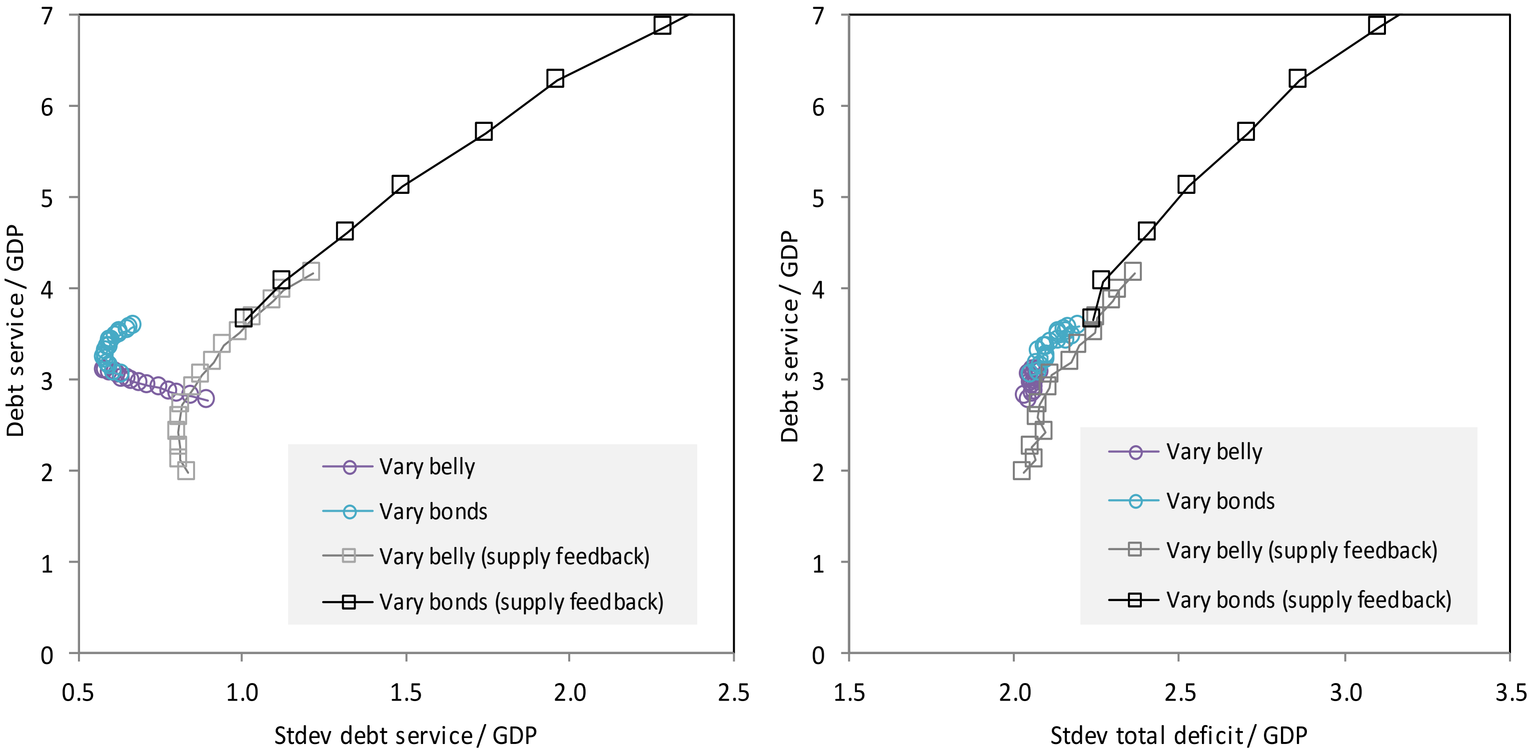

We next push our interest rate assumptions in the opposite direction, removing supply feedback and assuming that term premia remain at their current low levels, rather than reverting to more historically normal levels as assumed in our baseline results. Specifically, we adjust our term premium equations to result in steady state values for the two- and ten-year term premium of -35 bps and -39 bps, respectively. These assumptions line up with the Q1 2018 averages estimated from the ACM term premium model. This change introduces two important differences from our baseline results: first, the average level of term premia is much lower (and negative), and second, the term structure of term premia is slightly downward sloping from two- to ten-years, rather than upward sloping.

A priori, one would expect these alternative assumptions to result in lower debt service costs for all issuance strategies, and for longer-dated issuance to appear relatively more advantageous, compared with our baseline results. Figure 18 compares the two. First, the downward shift in the frontier is dramatic, as the lower average level of term premia presents a substantial benefit to the debt manager. Second, the assumption of a flat term structure of term premia significantly compresses the range of variation in average debt service costs across strategies. The most important point from these simulations, though, is that it is now more ambiguous whether the optimal strategy involves issuing more bonds or more belly. Issuing more bonds reduces the variation in debt service (in contrast to our baseline results), while issuing more in the belly reduces the variation in total deficits (in line with our baseline results).

Figure 18. Alternative frontiers assuming no term premium reversion

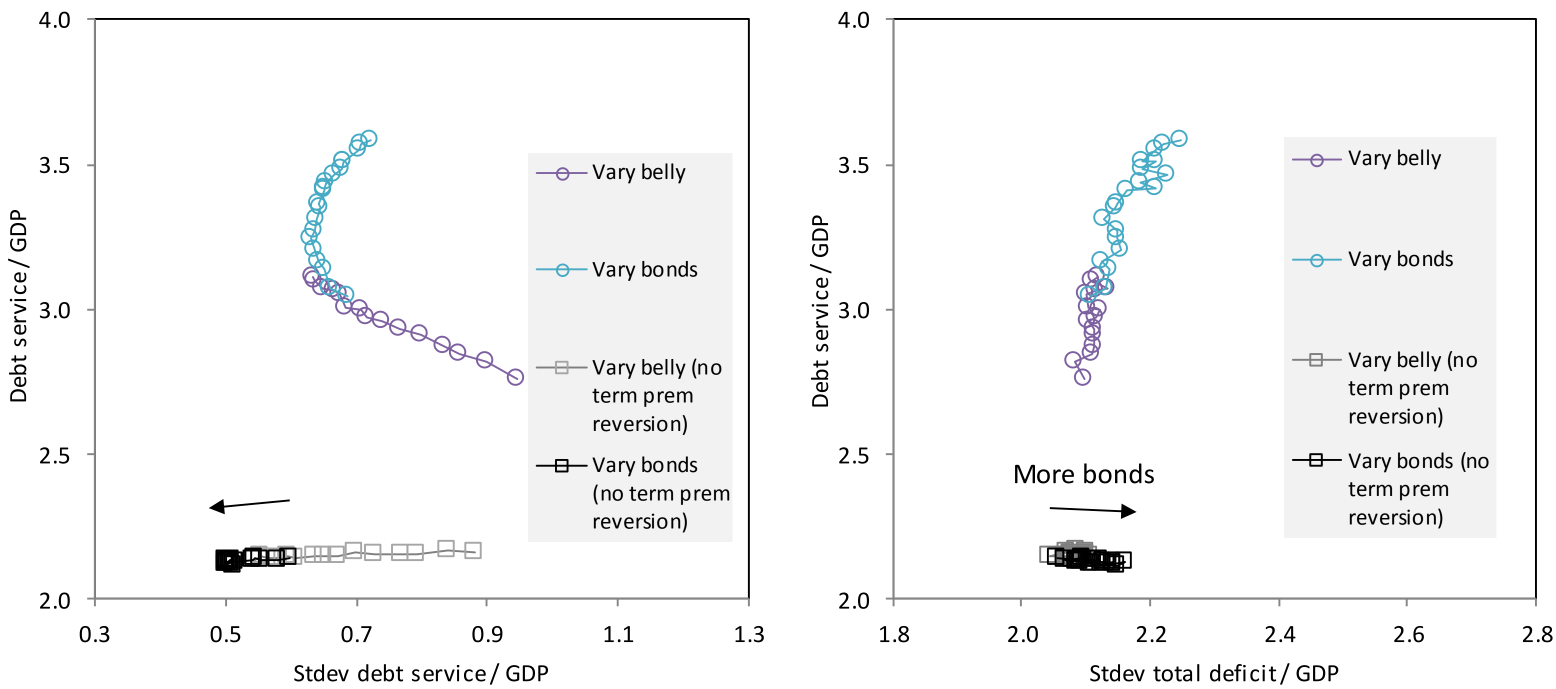

6.2 Drift in interest rates

One potential concern about the model is that it has a considerable degree of mean reversion in its structure, as noted earlier in section 2. It is therefore useful to consider the impact of changing the assumptions around the trajectory of interest rates in ways that allow for more drift. To examine this, we investigate a model that has a greater amount of uncertainty about the level of interest rates in the longer run. Specifically, we increase the shocks to the neutral policy rate by a factor of 3, which means that the uncertainty about the level of rates in 20 years is 2.3 percentage points higher than in our baseline model, as measured by the 15th/85th percentile interval. Because long rates are priced off of model-consistent expectations of short rates, they too will eventually reflect the uncertainty about interest rates over the long run. As can be seen in Figure 19, increasing the variance of R* shift the frontier strictly to the right, which is intuitive. More importantly, higher variance in R* reduces the relative disadvantage of longer-dated issuance strategies, in part because there is a greater advantage now associated with locking in funding costs. In the left-hand chart, the frontier does not start bending backwards as bond issuance increases, suggesting that an extremely risk averse debt manager might still find bond-heavy strategies to be optimal. In the right-hand chart, increasing the belly share traces out a path with a less favorable risk/cost trade-off than that seen in our earlier results.

Figure 19. Alternative assumptions about R* variance

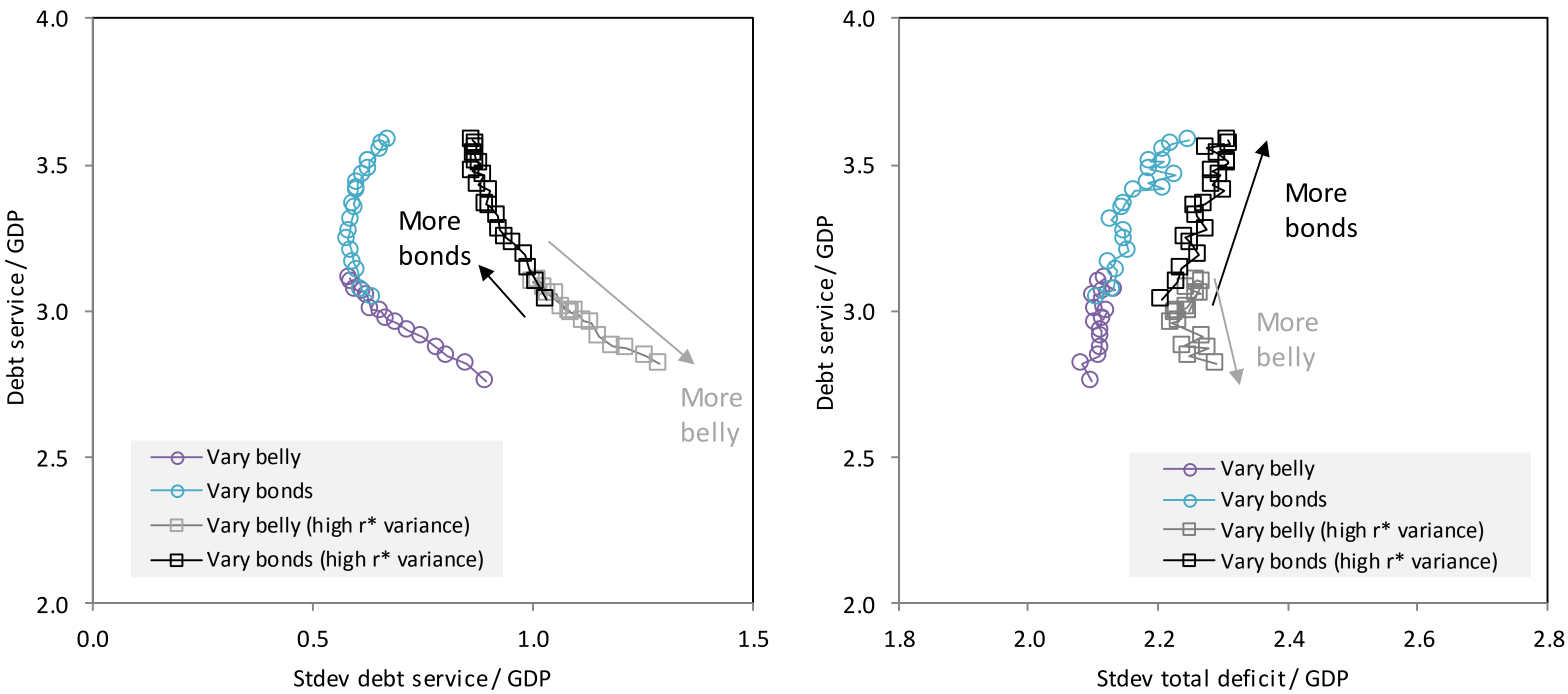

In general, we feel that a key advantage of having a model for assessing debt management is that it can be used to investigate how the debt management decisions are affected by changes in the model's assumptions. These results highlight that reasonable changes in assumptions can have meaningful effects on the trade-offs that the debt manager faces.

7. Conclusions

The goal of this paper is to develop an optimization model that helps US policymakers formulate an appropriate debt management strategy. Our findings yield some useful insights for framing debt management choices. First, issuing at intermediate maturities appears attractive in the model, providing the Treasury with significant reduction in the variation of funding costs with little additional cost. Second, issuing too much in the long end is not attractive. Even if it extends the WAM significantly, it does not reduce risk relative to issuing intermediate securities, and it has a high expected cost. Third, a willingness by the debt manager to adjust issuance patterns gradually over time in response to economic and financial conditions can provide it with significantly better performance.

We view this paper as an initial step and hope that it encourages further research on this topic. There are several obvious directions to explore. First, a shortcoming of the model is the omission of inflation-linked issuance (i.e., TIPS). Second, further work needs to be done on exploring alternative models of the term premium and the sensitivity of the optimization results to those assumptions. Third, the model needs to more carefully calibrate the feedback between changes in issuance volume and the cost of debt issued, perhaps through the introduction of some sort of issuance penalty. In this regard, there should be a more in-depth evaluation of the relative costs and benefits inherent in following a "regular and predictable" issuance pattern as opposed to a more dynamic model-based adjustment strategy.

From a broader perspective, optimization models such as the one presented in this paper can be an important part of the decision-making process for debt managers. But, it is only one of a variety of inputs available to policymakers, and Treasury should still seek a range of additional information about how market participants perceived the desirability of various maturity points and their estimates of issuance capacity across them. It is left to the Treasury to determine how to balance this judgmental information on market capacity with the model-based estimates on the debt structure that would be optimal.

REFERENCES

Adrian, Tobias, Richard K. Crump, and Emanuel Moench. (2013) “Pricing the Term Structure with Linear Regressions.” Federal Reserve Bank of New York Staff Report Number 340.

Balibek, Emre and Hamdi Alper Memis. (2012) "Turkish Treasury Simulation Model for Debt Strategy Analysis". World Bank Policy Research Working Paper 6091.

Barro, Robert J. (1979) “On the determination of the public debt.” Journal of Political Economy 87(5): 940-971.

Bergstrom, Pal, Anders Holmlund and Sara Lindberg. (2002) “The SNDO’s Simulation Model for Government Debt Analysis”, Swedish National Debt Office mimeo.

Blommestein, Hans J (editor). (2005) Advances in Risk Management of Government Debt, OECD.

Bolder, David J. (2008) “The Canadian Debt Strategy Model”, Bank of Canada Review.

Bolder, David J. (2003) “A Stochastic Simulation Framework for the Government of Canada’s Debt Strategy”, Bank of Canada Working Paper.

Bolder, David J and Simon Deeley. (2011) “The Canadian Debt Strategy Model: An Overview of the Principal Elements”, Bank of Canada Discussion Paper.

Dawsey, Kristopher. (2013) “Lower Deficits Won’t Keep Rates from Rising” *Goldman Sachs US Economics Analyst*. May 24, 2013.

Gagnon, Joseph. (2016) “Quantitative Easing: An Underappreciated Success.” *Peterson Institute for International Economics Policy Brief*. 16-4.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack. (2010) “Large-Scale Asset Purchases by the Federal Reserve: Did They Work?” *Federal Reserve Bank of New York Staff Reports*. #441.

Garbade, Kenneth D. (2007) “The Emergence of Regular and Predictable as a Treasury Debt Management Strategy”, FRBNY Economic Policy Review, Volume 13, Number 1.

Giavazzi, Francesco and Alessandro Missale. (2004) “Public Debt Management in Brazil”, NBER Working Paper No. 10394.

Greenwood, Robin, Samuel Gregory Hanson, Joshua S. Rudolph, and Lawrence Summers. (2015) "The Optimal Maturity of Government Debt." Chap. 1 in *The $13 Trillion Question: How America Manages Its Debt*, edited by David Wessel, 1–41. Brookings Institution Press.

Greenwood, Robin, Samuel Gregory Hanson, and Jeremy C. Stein. (2016) "The Federal Reserve's Balance Sheet as a Financial-Stability Tool." *Jackson Hole Economic Symposium Conference Proceedings (Federal Reserve Bank of Kansas City)*: 335–397.

Laforte, Jean-Philippe and John Roberts. (2014) “November 2014 update of the FRB/US model.” FEDS Notes. November 21, 2014.

Lucas, Robert Jr. and Nancy L. Stokey. (1983) “Optimal Fiscal and Monetary Policy in an Economy without Capital.” Journal of Monetary Economics 12(1): 55-93.

Pick, Andreas and Myrvin Anthony. (2006) “A simulation model for the analysis of the UK’s sovereign debt strategy”, UK Debt Management Office Research Paper.

Rudebusch, Glenn D. and John C. Williams. (2014) “A Wedge in the Dual Mandate: Monetary Policy and Long-Term Unemployment.” Federal Reserve Bank of San Francisco Working Paper 2014-14.

Taylor, John B. (1999) "A Historical Analysis of Monetary Policy Rules," Chapter 7 in: Monetary Policy Rules, National Bureau of Economic Research.

APPENDIX A: DEBT MANAGEMENT MODELS FOR OTHER COUNTRIES

Quantitative modelling and simulation of the risk-cost trade-off associated with different debt management strategies has been used by many countries. Debt managers in countries including Canada, the UK, Sweden, Brazil, and Turkey have published detailed working papers highlighting the key components of their models (see Bolder & Deeley, 2011, Pick & Anthony, 2006, Bergstrom, Holmlund & Lindberg, 2002 and Balibek & Memis, 2012). Other countries, including Denmark, Austria, Portugal and Belgium, have indicated that they utilize stochastic simulation of alternative debt management strategies (see Blommestein, 2005).

Many differences exist in the exact specification of costs and risks in the published quantitative models. Debt costs can be measured in dollars, as a percent of outstanding debt, or as a percent of GDP. Moreover, they can be discounted or undiscounted. Risk can be measured with respect to debt costs or the overall fiscal balance, and measured either using a standard deviation approach or a metric like tail risk VaR. Also, risk can be calculated at a single point in time or averaged over all time periods. For example, the Canadian model is one of the few that provides actual simulation results showing how the cost-risk tradeoff is impacted by alternative risk specifications. Some countries take credit risk into consideration. These are typically countries with elevated public indebtedness or emerging market economies, and they tend to view debt management as part of an integrated asset/liability framework.

While there are important differences in the exact model specifications, the overall approach used in different countries appears to be broadly similar and generally involves four key components: 1) a basic macroeconomic model that can be used to generate stochastic simulations of different economic and interest rate environments, 2) a term structure model that relates the yield curve to short term rates, 3) an objective function that typically involves minimizing expected issuance costs through time given constraints on risk and other variables and 4) an optimization module that identifies low cost strategies given alternative risk and issuance constraints. In general, the models are used to quantify the tradeoffs between cost and risk (i.e., an efficient frontier) rather than to identify a single optimal strategy.

A survey of this literature highlights some important similarities and differences in these four key components of the quantitative models. For example, the macro models that are used tend to be quite similar and typically involve equations for the output gap, inflation, short term interest rates and the primary deficit. In contrast, a variety of yield curve models are utilized to generate paths for interest rates as a function of the macroeconomic environment. There are also significant differences in the treatment of supply effects. In some cases this is ignored but other models include a penalty function that generates higher yields (i.e., greater cost) as issue size increases. Meanwhile, most of the models include an objective function that minimizes the expected costs of debt issuance over a long term horizon. Expectations are generally taken over multiple paths for interest rates, deficits and inflation. Choice variables include the allocation amounts across different points on the yield curve usually constructed to be constant weights through time.

Risk constraints seem to differ quite a bit across countries and can involve a limitation on the variability of either interest expense or the fiscal balance. In many cases, constraints based on fiscal balances incorporate the correlation between interest rates and primary deficits. This approach is well aligned with academic literature that highlight the social welfare benefits from tax smoothing (see Barro, 1979). The measurement of risk can be either a standard volatility measure (e.g., standard deviation of debt expense or budget) or a VaR-type limit (e.g., 95% confidence interval). Risk measures can also be incorporated directly into the objective function in a manner that is mathematically equivalent to having them as a constraint. For example, a constraint involving the desire to maintain regular and predictable issuance appears in many models. This tends to be accomplished by including a limitation on the change in issuance for each tenor from period to period. Other observed constraints relating to risk include: a limitation on the overall weighted average maturity (WAM) and maintaining a specified minimum volume of issuance at various points on the curve in order to support market liquidity and/or regulatory objectives.

Recognizing that the model specifications differ from country to country, a survey of the findings of the literature reveals the following: 1) optimization models often show the most attractive risk reduction per unit of cost by extending from bills to intermediate maturities (e.g., 5-year notes) while the risk reduction tends to be lower when extending further out the curve, 2) including a constraint that specifies a high volume of bill issuance (e.g., in order to meet market needs for liquidity or avoid operational disruption) appears to have the largest impact on reducing 10-year and longer issuance with smaller effects on the intermediate maturities, 3) setting risk constraints on the volatility of the fiscal balance rather than debt cost volatility generally results in higher levels of short term debt in the optimization reflecting the negative correlation between interest rates and primary deficits (the negative correlation means that it is relatively cheap to fund short-term when the primary deficit is cyclically large), 4) inflation-linked bonds tend to provide a diversification benefit and reduced cost volatility especially when constraints on short term budget volatility are present, 5) if issuance is constrained by a target WAM, the optimal solution tends to include a heavy reliance on 30-year or longer issuance reflecting the fact that the marginal impact of such debt on WAM is larger per unit of cost relative to say 5-year and 10-year maturities, and 6) the inclusion of issuance penalties that impose additional costs on heavy volume in a specific maturity reduces the incidence of corner solutions and tends to spread the optimal mix more evenly across the curve.

APPENDIX B: ALTERNATIVES MEASURES OF RISK FOR THE DEBT MANAGER

The literature on optimal debt management uses a number of different measures of risk. In the main part of this paper, we show that the decision of whether to measure risk based on the standard deviation of the debt service burden or the deficit had a meaningful impact on the results. Here we explore an even wider set of measures.

Even if we were to settle on those two variables for capturing risk, it is still the case that the risk associated for each variable could be measured differently. In the above results, we focus on the standard deviation of each variable. However, one could also measure risk based on the probability of particularly poor outcomes (a tail risk measure). In addition, the literature also considers some measures of conditional uncertainty—for example, not the standard deviation of the measure in 20 years, but the standard deviation of what would not be predictable a year in advance of that horizon. See Bolder (2008) for more extensive discussion.

We also believe that it is reasonable to focus on variables other than the two mentioned above. Those measures focus on outcomes at a particular point in the future, say 20 years ahead. Hence, they do not capture the benefits of intertemporal diversification. For example, a strategy of issuing only bills might look highly uncertain based on a snapshot taken for any one year, but it is also likely that issuing bills involves less risk when averaged over time. To account for this, we consider the uncertainty surrounding a third variable, which is the stock of debt as a share of GDP (in addition to uncertainty about debt service costs or the deficit). Since this measure is uncommon in the literature, we did not include it in the main text.

The following figure shows the frontier (under static issuance strategies) as we change the risk measure along these two dimensions. Moving across the charts horizontally, we are changing the variable that we are measuring risk on—debt service (first column), total deficit (second column), or amount of debt (third column), all expressed as a share of GDP. Moving down the page, we are changing the way that we measure risk on that variable—the standard deviation (first row), the cost of the 5% tail (second column), and the conditional standard deviation (third column).

For a given variable, the shape of the frontier is not strongly affected by whether we measure risk by the standard deviation, tail risk, or conditional volatility (the results do not change dramatically moving down the page in a given column). However, the shape of the frontier—and hence the resulting approach that appears optimal—varies considerably as we change the variable being measured (the results change a lot across columns). We already discuss in the paper how the shape of the frontier shifts as we move from debt service to total deficit (because of the correlation properties on shorter rates). As you can see, the results shift even more if we use the amount of debt as the variable to measure risk on. Basically, the front end looks even more appealing, because the variation it creates gets smoothed out over the business cycle (the level of debt can largely be thought of as being driven by the average deficit over the 20 years.

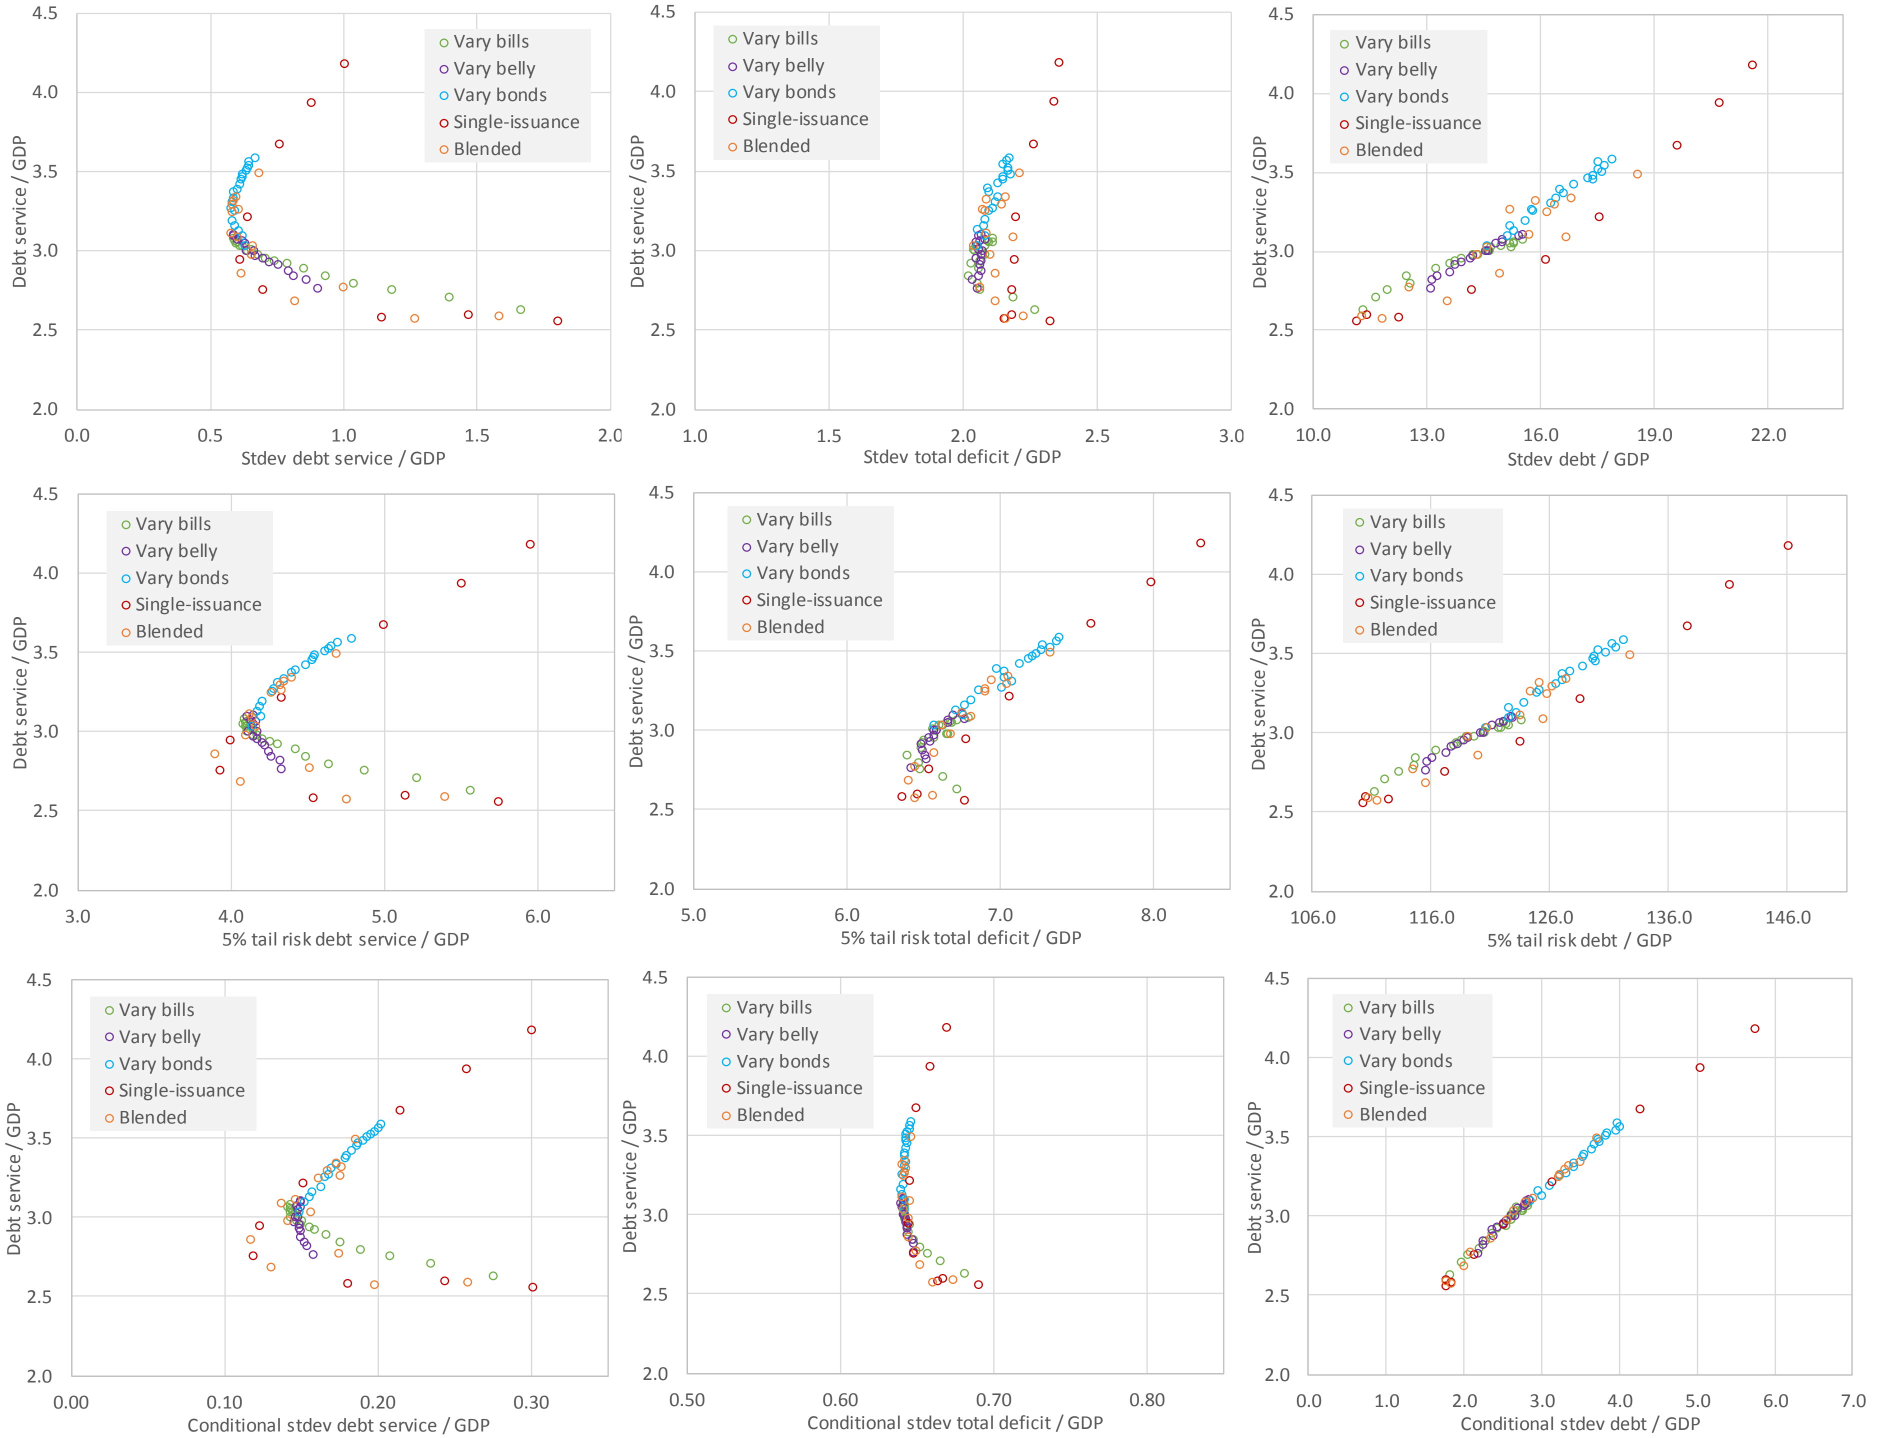

Vary bills
Vary belly
Vary bonds
Single-issuance
Blended
Debt service / GDP
Stdev debt service / GDP
Stdev total deficit / GDP
Stdev debt / GDP
5% tail risk debt service / GDP
5% tail risk total deficit / GDP
5% tail risk debt / GDP
Conditional stdev debt service / GDP
Conditional stdev total deficit / GDP
Conditional stdev debt / GDP

APPENDIX C: THE OPTIMIZATION PROBLEM FOR DYNAMIC ISSUANCE STRATEGIES

The text of the paper described the broad outlines of formulating the optimal dynamic issuance strategy. In this appendix we formally define our notation and write down the mathematical formulation.

Notation

Here is some of the key notation used in the optimization.

Definition of key sets:

- Let **T** ≡ {0, t_1, t_2, … T} denote the set of time steps in the problem. We will use the subscript letter *t* as an index that runs over this set. In our implementation, we measure time in months and use annual time steps with a 50-year horizon – thus, the set of time steps would be {0, 12, 24, … 600} for our specific implementation. In the interest of promoting readability, we will slightly abuse notation to refer to time steps by their indices rather than values – thus, for instance, it will henceforth be understood that time k actually refers to time t_k for some index k.
- Let **M** ≡ {m_1, m_2, …} denote the set of maturity points where issuance is permitted. Since we exclude TIPS and represent the entire Bill sector with the 1Y point, this is basically the set of permitted instruments. For our implementation, we use the specification **M** = {12, 24, 36, 60, 84, 120, 360}. We will use the letter *m* as a subscript to index that runs over this set. Similar to above, we will often refer to maturity tenors simply by the index that references the appropriate value in this set.
- Let **P** ≡ {1, … , P} denote the set of simulated paths being included in the optimization. This will often be indexed by the subscript *p*.
- Let **K** ≡ {1, … , K} denote the set of issuance kernels being included in the optimization. This will often be indexed by the subscript *k*.
- Let **N** ≡ {1, … , N} denote the set of macroeconomic variables (MEVs for short) being included in the optimization. This will often be indexed by the subscript *n*. The optimal solution produced by this optimization model will be a response function that is linear with respect to these MEVs.

Key input quantities:

- y_{pmt} denotes the Treasury yield in path p, at time step t, and for maturity point m. In the interest of brevity, unless otherwise noted it will be assumed that p ∊ **P**, m ∊ **M** and t ∊ **T.**
- D_{pt} denotes the primary deficit for the time interval (t-1, t] in path p. We use the convention that positive values denote deficits, and negative values denote surpluses.

- V_{npt} denotes the value of macro-economic variable n, in path p at time step t. (Note: the inclusion of a constant or "intercept" term in the response function is not automatic, but can be engineered by having an MEV that is identically equal to 1 at all time steps and in all paths).
- $\hat{R}_t$ denotes the amount of debt maturing at time t from the original stock of debt that was in existence at inception. Note that the redemption profile of existing debt is independent of paths (as no callable bonds issued by Treasury are in existence, and we exclude TIPS), while the total amount of redemptions at time t (denoted by R_{pt}) will in general vary from path to path.
- $\hat{c}_t$ denotes the debt service cost incurred in the time period (t-1,t] that accrues from the original stock of debt in existence at problem inception.
- Each kernel is specified by a set of coefficients f_{km} that denote the amount of issuance that will occur maturity m per unit amount of the kernel k. We will require that $\sum_m f_{km} = 1$ for k = 1, and $\sum_m f_{km} = 0$ for all other k.
- θ_t denotes the length of the time period ending in t, expressed in years. For instance, if quarterly time steps are used and all time intervals are evenly spaced, then θ_t will be 0.25 for all t.
- z_{pt} denotes the present value (at time step 0) of a unit cash flow occurring at time step t in path p.
- ϵ^{RP} denotes the percentage increase or shrinkage in issuance amounts in a given tenor that are permissible, from the standpoint of Treasury's desire to maintain a "regular and predictable" issuance strategy.

Decision variables:

- N_{pimt} denotes the notional amount of Treasuries of original-issue-maturity m, issued at time i, in existence at time t, in path p. It should be noted that $i \in \mathbf{T}$, and $i \leq t$.
- G_{pmt} denotes the gross issuance amount of Treasuries of maturity m, in path p, issued at time t.
- W_{pkt} denotes what we call a kernel loading – i.e., the weighting associated with each kernel k, at time step t in path p. The variables W_{pkt} collectively (over all k) define the actual gross issuance amounts G_{pmt} at time step t for each maturity m, as will be seen later.
- h_t denotes growth index reflecting growth of nominal stock of debt over time
- β_{kn} denote the coefficients of the response function, helping to determine the values of the kernel loadings at each time step in each path. These are the ultimate targets of the optimization, as we will see later.
- R_{pt} denotes the amount of debt maturing in the time interval (t-1,t].
- c_{pt} denotes the total debt service cost incurred for the period ending in time step t. In general, this varies across paths because issuance amounts can vary across paths.
- Ω_p denotes the present value of debt service costs accrued along path p

Constraints

We can now specify each constraint in the problem.

Reconstructing gross notional amounts

We begin with a definitional constraint that links gross issuance variables to kernel loadings. This is the constraint that forces issuance amounts in any solution to be obtainable as a linear combination of our chosen kernels.

$$G_{pmt} = w_{p,1,t} f_{1,m} + \sum_{k=2 \ to \ K} w_{pkt} f_{km} h_t \quad \forall p, m, t$$

Kernel loadings as a function of MEVs

A second definitional constraint forces the kernel loadings themselves to be obtainable as linear combinations of MEVs. Note that this constraint, as written, doesn't account for an "intercept" term in the response function. That is easily addressed, if desired, by including an MEV that takes the value 1 in each path and at each time step.

$$w_{pkt} = \sum_{n} \beta_{kn} V_{pnt} \quad \forall p, k, t$$

The above two sets of constraints make it clear that the β_{kn} variables are the ones that ultimately determine issuance amounts. Note that these variables have no subscript that indexes path or time – these coefficients specify an optimal response function that is constant through time and across paths, even though the realized values will of course vary across time and across paths.

Tracking redemptions

In each path, we will need to track the amount of debt maturing at each time step – i.e., the variables R_{pt}. These are intermediate variables that will be necessary in order to write a flow-balance constraint.

$$R_{pt} = \hat{R}_t + \sum_{m} G_{p,m,t-m} \quad \forall p, t$$

Flow balance

The family of flow balance constraints ensure that the total amount of gross issuance at each time step and in each path is exactly equal to the amount necessary to cover redemptions, fund the primary deficit in that period and also cover the debt service cost in that period. It is worth noting that this constraint is written as an equality constraint, rather than an inequality constraint that might allow for overfunding in some periods. This is because in Treasury's case, overfunding is forbidden under current law. Relaxing this assumption to explore the potential benefits of opportunistic overfunding is a potential area of future research.

$$\sum_{m} G_{pmt} = R_{pt} + D_{pt} + c_{pt} \ \ \forall\, p, t$$

Debt service costs

We need to add a family of constraints to ensure that the debt service cost variables c_{pt} take on their appropriate values. We do that by adding the following constraints.

$$c_{pt} = \sum_{i,m: i+m \geq t} G_{pmi}\, y_{pmi}\, \theta_t \ \ \forall\, p, t$$

Path cost

We are now ready to define the debt service cost of an entire path. We do this by adding the following definitional constraint.

$$\Omega_p = \sum_{t} c_{pt}\, z_{pt}$$

Regular and predictable issuance

Treasury has historically sought to maintain a “regular and predictable” issuance strategy. We model this preference via the parameter ϵ^{RP}, which controls how much issuance amounts can vary from one time step to the next in a given tenor and in a given path.

$$G_{p,m,t} \leq (1 + \epsilon^{RP}) G_{p,m,t-1} \ \ \forall\, p, m \text{ and } \forall\, t > 1$$

$$G_{p,m,t} \geq (1 - \epsilon^{RP}) G_{p,m,t-1} \ \ \forall\, p, m \text{ and } \forall\, t > 1$$

In our exploration of the model where the desire is to learn about the dynamics of the response function, this constraint is set loosely and does not impact the results. Tightening this constraint and exploring the trade-off between cost and the merits of being “regular and predictable” is an area of future research.

Optimization

The optimization problem can now be specified as follows.

Objective function

From a public policy perspective, there are several possible choices of what a sovereign debt manager's objectives should be. For our purposes, we assume that Treasury is focused on two objectives – minimizing debt service costs, and mitigating the variability of budget costs. Consistent with this, we define the objective function which is to be minimized as the sum of two parts:

- The expected present value of debt service costs across all paths, and …
- … a linearized proxy for budget cost volatility. This proxy is just the absolute deviation of budget cost 20 years forward across all paths.

The two cost components are blended together with a weight parameter that can be tweaked to alter the required balance between costs and volatility. Thus, our objective function is specified as

$$Min \sum_{p} \Omega_p + \lambda \cdot \epsilon_p^+ + \lambda \cdot \epsilon_p^-$$

where

$$\epsilon_p^+ - \epsilon_p^- = \left(\frac{(c_{pt} + D_{pt}) z_{pt} - \widehat{\Omega}}{GDP_{pt}} \right), \ \ where\ t = 20\ years\ forward$$

and

$$\widehat{\Omega} = \sum_{p} \frac{(c_{pt} + D_{pt}) z_{pt}}{number\ of\ paths}, \ \ where\ t = 20\ years\ forward$$

We also require ϵ_p^+ and ϵ_p^- to be non-negative. Note that these two variables have a complementary existence, and at least one of them will be 0 in the optimal solution. Thanks to this property (itself the result of a penalty in the objective function), $\epsilon_p^+ + \epsilon_p^-$ will effectively represent the absolute value of the right hand side term in the equation above, which is penalized in the objective function. This problem is solved using 50 simulated paths for the economy, funding needs, and the yield curve over a 40-year period. It is implemented in MATLAB and solved using MATLAB's native Linear Programming solver. The results from this optimization problem are shown in the main text of the paper.

Hutchins Center on Fiscal & Monetary Policy at BROOKINGS

The mission of the Hutchins Center on Fiscal and Monetary Policy is to improve the quality and efficacy of fiscal and monetary policies and public understanding of them.

Questions about the research? Email communications@brookings.edu. Be sure to include the title of this paper in your inquiry.