Paul Glasserman, Amit Sirohi, and Allen Zhang

The Effect of "Regular and Predictable" Issuance on Treasury Bill Financing

- To meet the nation's financing needs at the lowest cost over time, the U.S. Treasury issues its Treasury securities using a "regular and predictable" approach.
- But by doing so, does it forgo the short-term gains that might be achieved by issuing debt "opportunistically"—that is, when market conditions are most advantageous?
- This study compares financing costs under a strict cost-minimization strategy with those of alternative strategies that focus instead on "smoothness" considerations—interpreted here as variations of the "regular and predictable" principle.
- The additional cost of such strategies in terms of average auction yield is likely less than one basis point. Adding the flexibility to use cash management bills narrows the gap further.

1. Introduction

In a speech in 2002, Peter Fisher, then under-secretary of the Treasury for domestic finance, stated that "the overarching objective for the management of the Treasury's marketable debt is to achieve the lowest borrowing cost, over time, for the federal government's financing needs" (Fisher 2002). Treasury officials have followed Fisher's agenda ever since.

In pursuit of financing at least cost over time, the Treasury adheres to a "regular and predictable" issuance program. As reported in Garbade (2007), the Treasury initially moved toward regular issuance of short-term notes in 1972 and fully embraced the practice in 1975 after rapid growth of the deficit. In 1982, Mark Stalnecker, then Treasury deputy assistant secretary for federal finance, testified that "regularity of debt management removes a major source of market uncertainty, and assures that Treasury debt can be sold at the lowest possible interest rate consistent with market conditions at the time of sale." In 1998, Gary Gensler, at the time the Treasury assistant secretary for financial markets, reinforced that principal, stating that "Treasury does not seek to time markets; that is, we do not act opportunistically to issue debt when market conditions appear favorable."

This article is based in part on work done while Paul Glasserman, the Jack R. Anderson Professor of Business at Columbia University, was a consultant to the U.S. Department of the Treasury through the Sapient Corporation; Amit Sirohi was a senior associate at Sapient; and Allen Zhang was deputy director of the Office of Debt Management at the U.S. Treasury Department.

Correspondence: Allen Zhang, xpzhang2001@gmail.com

The authors thank Chris Cameron, Dave Chung, James Clark, John Dolan, Lev Dynkin, Colin Kim, Fred Pietrangeli, Matthew Rutherford, Linda Xie, and Ernest Zhu for helpful discussions and two anonymous referees for their detailed comments. Any remaining errors or omissions are the sole responsibility of the authors. The views expressed in this article are those of the authors and do not necessarily reflect the position of the U.S. Department of the Treasury, the Federal Reserve Bank of New York, or the Federal Reserve System.

To view the authors' disclosure statements, visit https://www.newyorkfed.org/research/author_disclosure/ad_epr_2017_treasury_glasserman.html.

Chart 1

Historical Issuance of Treasury Bills

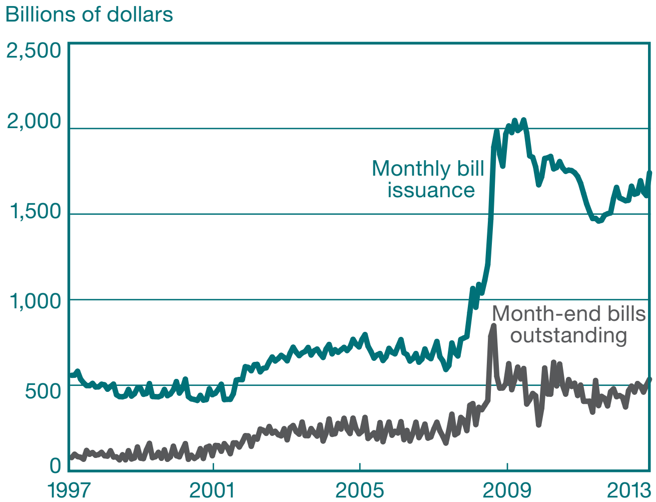

Source: U.S. Department of the Treasury Office of Debt Management (ODM).

Note: The chart reflects issuance and bills outstanding for privately held Treasury bills (excluding System Open Market Account holdings and Supplementary Financing Program bills) from January 1997 to February 2013.

In practice, regular and predictable issuance entails prior announcement of the issuance schedule and gradual adjustment of issuance sizes. Of course, taking a regular and predictable approach does not mean that debt management practices never vary. Borrowing requirements change frequently and the Treasury constantly reevaluates issuance strategies and occasionally revises them to best serve the debt management mission. The process requires the definition of objectives and constraints, recognizing that, given multiple ways of satisfying financing needs, some approaches are better than others.

This article focuses on the potential impact of regular and predictable issuance on the short-run cost of issuing Treasury bills. As an issuer of both Treasury bills and coupon-bearing securities (including fixed-rate and inflated-protected securities), and given a coupon issuance schedule, the Treasury uses the bills in part for short-term financing and in part for cash management. The overriding constraint is to raise enough cash to satisfy the government's financing needs. In addition, cash balances need to be in an appropriate range—large enough to provide the Treasury with a buffer against unexpected events, but not so large as to create inefficiencies through over-borrowing. In addition, since Treasury bills are used extensively in the global financial system, it is desirable to maintain a steady supply for investors.

The historical bill issuance and amounts outstanding during the past fifteen years are shown in Chart 1. The figures reflect private issues only, and exclude rollovers in the Federal Reserve's System Open Market Account (SOMA) and sales of Supplementary Financing Program (SFP) bills.[1] Because of the short maturity of bills, the gross auction amount is astonishingly large, reaching a peak of almost $6.7 trillion in fiscal 2009 amid the turbulence of the financial crisis. Issuance subsequently decreased when the Treasury moved to extend the weighted average maturity of debt to reduce "rollover risk"—the risk of facing unfavorable interest rates when rolling over matured debt in the future—and to take advantage of historically low term premia.

A key question—which is simple, yet has important policy implications—is whether regular and predictable issuance raises the Treasury's borrowing costs. Relevant studies in the literature are scarce. Garbade (2007) relies on a "natural experiment" in which he compares nominal coupon issuance in 1971-75 (when bills were sold on a "tactical" basis) with that in 1981-86 (when they were offered on a "regular and predictable" schedule). Using the root-mean-square change in yields over the interval from the close of business one business day before an auction announcement to the close of business one business day after the announcement, Garbade finds that most changes in yield are statistically significant in the tactical period while all changes in yield are insignificant in the regular period. He concludes that "the move to regular and predictable issuance helps to reduce market uncertainty, facilitate investor planning, and lower the Treasury's borrowing cost." However, the drawbacks of a natural experiment are that it is difficult to control for differences in environment, and it does not allow for counterfactual policy analysis or scenario analysis.

In this article, we quantify the potential cost of a regular and predictable approach to bill issuance by examining alternative issuance strategies in an optimization framework.[2] An issuance strategy describes offerings over a period of time; throughout

[1] The Supplementary Financing Program was initiated in September 2008 for the purpose of draining reserves from the banking system. Proceeds from those auctions were maintained at the Federal Reserve Bank of New York. SFP balances declined to zero in July 2011 and the Treasury has not yet resumed auctioning SFP bills. In this article, because we focus on the net cash raised in the private market to satisfy the government's financing needs, SOMA and SFP holdings are excluded.

[2] Mathematical optimization, or simply, optimization, is commonly understood as the selection of a best element with regard to some criterion from some set of available alternatives. In our case, we minimize a real function (the "objective function") by systematically choosing input values from within an allowed set (the "constraints").

the article, we will refer to this period of time as the "projection horizon." In choosing a sequence of issuance sizes, we optimize alternative objectives subject to financing and issuance constraints. All of the objectives are quadratic functions of the issuance amount and all of the constraints are linear, so we can formulate the optimization as a constrained quadratic programming problem. Each optimal solution—corresponding to a particular objective function—is characterized by variability and cost metrics. The analysis ignores the ancillary benefits of reducing market uncertainty and facilitating investor planning, both of which could be expected to promote auction participation; the analysis may, therefore, understate the overall benefit of regular and predictable issuance.

Our benchmark is a pure short-run cost-minimization strategy, which assumes that the Treasury seeks only to keep short-run financing costs as low as possible given a forecast of future interest rates. Such an objective may lead to opportunistic issuance, with the Treasury possibly appearing to be "timing the market." Alternatively, on an ex ante basis, the Treasury can choose to optimize on some regular and predictable behavior, rather than on cost-minimization. For example, it may try to smooth issuance by minimizing changes in offering sizes, resulting in higher short-run issuance costs as compared with the benchmark. The cost difference between the two approaches measures the trade-off between being regular and predictable and being "opportunistically" cost-minimizing. The Treasury could also employ other proxies for regular and predictable behavior, such as maintaining a low cash balance, not deviating from a baseline strategy, or targeting low gross issuance overall. By comparing the cost of an alternative strategy with the cost of the benchmark strategy, we determine what we give up to be regular and predictable.

We also extend our basic framework to include the option to issue cash management bills (CMBs). CMBs are securities with flexible (usually very short) maturities whose proceeds are used by the Treasury to meet temporary shortfalls. Modeling the decision to issue CMBs requires the introduction of binary variables, which substantially complicates the optimization problem. To mitigate this increased complexity, we develop a heuristic rule based on the shadow prices (the Lagrange multipliers) associated with the financing constraints to identify the timing of CMBs.[3] Use of such a rule allows us to bypass what would otherwise have been a far more complex calculation and saves significant computing time. We detail this methodology in Section 4.2.

[3] The multiplier measures the change in the objective function owing to a marginal change in the constraint. A high multiplier indicates a possibly high benefit of using a CMB.

Our examples indicate that cost tends to decrease with higher levels of week-to-week variability—that is, when allowing larger changes in consecutive issuances—until it reaches the global minimum cost (GMC). By definition, the GMC is the benchmark strategy. Alternative optimization objectives that are based on regular and predictable issuance lead to higher financing costs than the GMC strategy, though the flexibility to use a limited number of CMBs may help reduce those costs.[4]

Other optimization problems arising in national debt management have been addressed in recent work. These include Adamo et al. (2004) in Italy; Balibek and Köksalan (2010) in Turkey; Bolder (2008) in Canada; Date, Canepa, and Abdel-Jawad (2011) in the United Kingdom; and Hahm and Kim (2003) in Korea. However, none of these models considers the question of regular and predictable issuance.

The rest of the article proceeds as follows. In Section 2, we introduce an issuance optimization model for Treasury bills. In the third section, we solve for the "efficient frontier" (the combination of securities that offers the lowest risk for a given level of return or the best return for a given amount of risk) and illustrate differences in variability–cost profiles associated with different objectives. Section 4 considers the use of CMBs. In Section 5, we present our conclusions. The appendix offers a robustness check of our main results.

[4] Our optimization framework uses projected cash needs to set auction sizes. All of our strategies are ex ante in nature: (1) the optimization assumes that funding needs will unfold exactly as projected at the beginning of the planning period; and (2) our performance evaluation implicitly assumes that the issuance plan, once selected, is followed strictly throughout the planning period. In reality, both assumptions are questionable: Forecasts of funding needs are revised every week, and the Treasury will adapt to updates in funding needs and revise future issuance plans accordingly. Hence, the realized (ex post) issuance strategy likely differs from the ex ante optimal one, regardless of which objective function is used.

To capture the effect of new information over time, we developed a step-through simulation procedure to evaluate the ex post performance of an issuance strategy. In the simulation, the Treasury optimizes the issuance plan over the full planning horizon but locks in the auction sizes for the first week only. We then advance the simulation by a week, and revise the projected cash needs from a statistical model of fiscal revisions. The Treasury re-optimizes the issuance plan based on the new projections and the process repeats. We simulate a large number of such paths to compare the realized performance of alternative rules in the face of forecast revisions and unanticipated changes in fiscal flows. Technically, we solve an "open-loop" problem, but then we implement it "closed-loop" because we lock in only the first step and then re-solve the problem to respond to the new environment—a procedure known as "model predictive control" in the control literature.

We do not detail the step-through procedure in this article because it is not essential to understanding the trade-off between variability and cost. It may suffice to note that the actual realized issuance based on the step-through optimization may result in further narrowing of the cost difference. In other words, part of the short-term cost advantage of the GMC is lost once we recognize changes in projected cash needs over time.

2. Issuance Optimization

The Treasury has a fixed auction calendar of bills, nominal fixed-rate notes and bonds, floating-rate notes, and inflation-protected notes and bonds; a detailed description of Treasury securities and the auction process can be found in Garbade and Ingber (2005). In its quarterly refunding announcement, the Treasury specifies the intended auction amounts of coupon-bearing securities for the upcoming quarter. Given a starting cash balance, fiscal flow forecasts, and cash flows from coupon securities, the Treasury issues bills to fund cash requirements and maintain a proper cash balance.

As opposed to coupon securities, which make semiannual (or quarterly, in the case of floating-rate notes) payments until maturity, bills are single-payment securities that are sold at a discount and pay a specified face value at maturity; their yields are floored at zero in primary market auctions. Regular offerings include four maturities: four weeks, thirteen weeks, twenty-six weeks, and fifty-two weeks, with the first three maturities auctioned weekly and the fourth every four weeks. CMBs, with maturities ranging from one day to a few months, are offered as needed.

Given the debt management objective of minimizing cost and the principle of regular and predictable issuance, it is natural to examine bill issuance strategies in an optimization framework. In particular, we want to solve for the optimal issuance program over a specified projection horizon, such that the net cash flow from bill issuance and redemption is sufficient to cover financing needs (resulting from net fiscal flows, coupon payments, and principal redemptions of coupon-bearing securities) and maintain an appropriate cash balance. (In Section 3, we take the projection horizon to be twenty-seven weeks, mainly because fiscal forecasts become less reliable beyond half a year. Additionally, the Treasury is only committed to the issuance sizes of coupon securities in the next quarter, so we may not be able to take coupon issuance amounts as given for a longer term.) Bills are the residual financing instrument in this short-term issuance model.

2.1 The Issuance Program

We outline here our optimization model. Let the N-dimensional column vector X denote the sequence of issuance amounts of regular bills before the end of the projection horizon, where N is the total number of offerings over the horizon. The components of X are the choice variables of the optimization problem. We denote the issuance amount in week i of a bill with a term to maturity of j weeks as $x_{i,j}$ for $i = 1, 2, \ldots, T$ (where $T = 27$ weeks) and $j = 4, 13, 26$, and 52. The vector X is formed by stacking up the $x_{i,j}$ terms.

2.2 Exogenous Inputs and Related Constructions

There are four exogenous inputs:

1. a T-dimensional column vector f of weekly cash needs resulting from net fiscal flows, coupon payments on notes and bonds, redemptions of previously issued bills, and issuance and redemption of coupon-bearing securities;
2. the Treasury's initial cash balance, denoted c_0;
3. the Treasury's last issues of 4-, 13-, 26-, and 52-week bills prior to the start of the projection horizon, denoted $x_{0,j}$ for $j = 4, 13, 26$, and 52; and
4. future bill auction rates.

We use these four inputs to compute bill auction prices, the Treasury's cumulative net cash requirements, the trajectory of Treasury cash balances associated with a given issuance program, and week-to-week changes in issuance.

Let $r_{i,j}$ denote the auction rate in week i for a bill with a term to maturity of j weeks, so that the issuance price of the bill per dollar payable at maturity, $p_{i,j}$, is given by $p_{i,j} = 1 - (j / 52) \times r_{i,j}$.[5] We assume that bill rates over the projection horizon are equal to the forward, or expected, rates implied by the on-the-run[6] curve at the beginning of the projection horizon. (We are not claiming that forward rates are the best forecasts of future spot rates; they have the obvious drawback of excluding term premia. They do, however, provide a reasonable forecast of future interest rates that allows us to focus on the consequences of varying the issuance objective.)

[5] Issuance decisions could feed back to auction rates. For example, Li and Wei (2012) suggest that the total supply of Treasury securities may affect Treasury yields through a term premium channel. In addition, as evidenced by the bid elasticity curve observed from bid-level data, deviation from the issuance size would affect the auction stop-out rate (the lowest accepted bid rate), suggesting a "funding mix effect." In this study, we ignore supply effects since we only consider bill issuance in the short term, which does not carry significant information about total debt outstanding. We also bypass the funding mix effect by imposing hard constraints on issuance sizes and changes in issuance sizes, allowing only marginal shifts.

[6] The on-the-run Treasury curve is derived from on-the-run securities—which currently refer to the most recently issued Treasury notes (2-year, 3-year, 5-year, 7-year, and 10-year) and bonds (30-year)—as opposed to "off-the-run" securities that were issued before the most recent issue and are still outstanding. On-the-run securities comprise more than half of total daily trading volumes, and are mainly traded in the interdealer market. It is commonly believed that on-the-run securities have better liquidity than off-the-run, and the on-the-run curve is the primary benchmark used in pricing fixed-income securities.

Denote the sequence of *cumulative* net cash requirements as the T-dimensional column vector b, where the ith element of b is weekly cash needs prior to and in week i, less the Treasury's initial cash balance:

$$b_i = \sum_{k=1}^{i} f_k - c_0 \qquad i = 1, 2, ..., T$$

Denote the contribution to weekly Treasury cash balances from a bill with a term to maturity of j weeks issued in week i, per dollar payable at maturity, as the T-dimensional column vector $a_{i,j}$. If $i + j$ is less than or equal to T, the first $i - 1$ elements of $a_{i,j}$ are equal to zero (because the bill is not issued until week i), the next j elements are equal to $p_{i,j}$ (because the bill is outstanding during those weeks), and the remaining elements are equal to $p_{i,j} - 1$ (because the bill is redeemed in week $i + j$). If $i + j$ is greater than T, the first $i - 1$ elements of $a_{i,j}$ are equal to zero and the remaining elements are equal to $p_{i,j}$. Ordering the $a_{i,j}$ column vectors into a T-by-N matrix A to match the order of the elements of X, we can express the sequence of weekly Treasury cash balances resulting from the bill issuance program X as $AX - b$.

Finally, consider the "gradient" of issuance—that is, the changes in issuance from one week to the next. The gradient of the first four 4-week issuances, when the last known issuance is $x_{0,4}$, is given by

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \end{bmatrix} \begin{bmatrix} x_{1,4} \\ x_{2,4} \\ x_{3,4} \\ x_{4,4} \end{bmatrix} - \begin{bmatrix} x_{0,4} \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

The indexation becomes more complicated when the issuance program X includes bills with a variety of maturities and different issuance frequencies, but the gradient is nevertheless linear in X and can be written as $DX - d$, where D is an N-by-N matrix and d is an N-dimensional column vector with, at most, four nonzero elements, identified as $x_{0,4}$, $x_{0,13}$, $x_{0,26}$, and $x_{0,52}$.

2.3 Two Metrics

Given an issuance strategy X, we propose two metrics: one to assess short-run financing costs and the other to assess variability, or changes in issuance size from week to week.

The cost metric is straightforward. Bills are offered on a discount basis, so the cost of issuing a bill with a term to maturity of j weeks in week i, per dollar payable at maturity, is $(j / 52) \times r_{i,j}$ if $i + j$ is less than or equal to T, and a prorated amount of that quantity otherwise. Ordering these cost terms in an N-dimensional row vector h, we can express the total (undiscounted) financing cost over the projection horizon, denoted FC, as

$$FC = hX.$$

We define the variability metric AG as the root mean of the squared changes in consecutive issuances (weekly for 4-week, 13-week, and 26-week bills and once every four weeks for 52-week bills):

$$AG = \sqrt{\tfrac{1}{N}\,(DX - d)'\,(DX - d)}.$$

All else equal, if issuance size needs to be increased, the Treasury prefers a sequence of small changes to a single large increase in order to minimize disruption in the market.

2.4 Constraints and Objectives

We study the trade-off between short-run financing costs and variability in an optimization framework. Simply put, if being regular and predictable means lower variability, does that always lead to higher short-run costs? And if so, by how much?

Our choice variable is X, the issuance program. We impose three constraints on the choice of X:

(1) $$\underline{c} \le AX - b \le \bar{c}$$

(2) $$L \le X \le U$$

(3) $$-\delta \le DX - d \le \delta$$

Constraint (1) is a financing constraint: After satisfying weekly cash needs, the weekly cash balances, $AX - b$, must be within a specified range $[\underline{c}, \bar{c}]$. The Treasury picks the range to maintain a cash buffer sufficient to safeguard against forecasting errors and unanticipated, sudden loss of market access. Both the floor ($\underline{c}$) and ceiling ($\bar{c}$) are vectors of dimension T, so the allowable range of cash balances can vary from week to week.

The next two constraints relate to issuance. Constraint (2) sets lower and upper bounds L and U, respectively, on offering amounts. These are vectors of dimension N, so the bounds can vary over time and across bill maturities. Constraint (3) limits the change in offering amounts between consecutive auctions. The choices of δ, L, and U

follow from the Treasury's understanding of market reaction to changes in offering amounts and the need to maintain a deep and liquid market.

We consider five alternative objective functions:

(4a) $$minCost: hX$$

(4b) $$minCB: (AX - b)'(AX - b)$$

(4c) $$minGrossIss: (AX - b)'(AX - b) + \omega X'X$$

(4d) $$minDevBase: (AX - b)'(AX - b) + \omega(X - \widetilde{X})'(X - \widetilde{X})$$

(4e) $$minGrad: (AX - b)'(AX - b) + \omega(DX - d)'(DX - d)$$

where the scaling factor ω is a scalar.

The *minCost* objective function minimizes financing costs and is used to identify the benchmark strategy, global minimum cost.

None of the other four functions explicitly minimizes cost. Rather, each reflects the Treasury's degree of aversion to a high level of cash balances and preference for "smoothness" in issuance—which involves the predictability of issuance given the auction calendar. The *minCB* objective reflects only a concern with high cash balances, owing to the "negative carry" that usually occurs in funding such balances.[7] Moreover, maintaining a relatively constant cash balance at the lowest possible level produces a smooth cash-balance profile, which may be welcomed by the market as evidence of "predictable" behavior.

Objectives (4c), (4d), and (4e) continue to express an aversion to higher cash balances but add a second component to capture issuance predictability. The *minGrossIss* objective seeks to limit the gross issuance. Because shorter-maturity bills require more frequent rollover, *minGrossIss* tends to favor the use of longer-term securities. The *minDevBase* objective function tilts the strategy toward a desired baseline $\widetilde{X}$, an input to the optimization process, reflecting the Treasury's understanding of market conditions. And *minGrad* is meant to explicitly control the variability metric AG. The scaling factor ω controls the relative importance of the two components in each objective.

[7] By issuing Treasury securities to fund the cash balance, the Treasury incurs interest costs while not earning any yield, thus the "negative carry." Recently, domestic banks and foreign banking organizations became eligible to earn interest on excess reserves (IOER) on reserve balances held in their Federal Reserve accounts since 2008 (as part of an effort to stabilize the federal funds rate). When banks purchase Treasury securities and reduce reserve balances, they earn yields from holding the Treasury securities but receive no IOER. Hence, the "negative carry" to the Treasury could be partially offset by IOER in this broad context.

Table 1
Model Parameters and Exogenous Inputs

Parameters	
T	Projection horizon in weeks
N	Number of new bill issuances over the course of the projection horizon
$\underline{c}$	(Possibly time-varying) cash balance floor, a T-dimensional vector
$\overline{c}$	(Possibly time-varying) cash balance cap, a T-dimensional vector
δ	(Possibly time-varying and term-specific) maximum change between consecutive auctions, an N-dimensional vector
L	(Possibly time-varying and term-specific) lower boundary on issuance size, an N-dimensional vector
U	(Possibly time-varying and term-specific) upper boundary on issuance size, an N-dimensional vector
$\widetilde{X}$	The Treasury's desired or baseline issuance strategy, an N-dimensional vector
ω	Relative weight when there are two items in the objective function, a scalar
Exogenous Inputs	
$r_{i,j}$	Auction discount rate in week i for a bill with a term of j weeks
f	Projected weekly cash needs vector, a T-dimensional vector
c_0	Initial cash balance
$x_{0,j}$	Last issuance size of a j-week bill before the beginning of the projection horizon

Notes: The choice set consists of the issuance size in week i of a bill with a term of j weeks, $x_{i,j}$ for $i = 1, 2, ..., T$ and $j = 4, 13, 26$, and 52. The N-dimensional vector X is formed by stacking up the $x_{i,j}$ terms.

We summarize the parameters and inputs needed to formulate and solve the optimization problem in Table 1.

3. An Efficient Frontier and Alternative Strategies

In this section, we introduce an efficient frontier describing the trade-off between the variability and cost metrics defined in section 2.3 given the constraints—equations (1), (2), and (3)—on issuance. We then analyze the five alternative strategies and contrast their variability and cost measures.

3.1 Model Parameters

For purposes of illustration, we use a twenty-seven-week horizon starting Thursday, November 12, 2009, (following issuance on that day of the 4-week, 13-week, and 26-week bills) to examine various bill issuance strategies. During this period, there were eighty-eight bill offerings. The Office of Fiscal Projections at the Treasury provides daily fiscal forecasts up to twenty-seven weeks in the future (at the time of this analysis). Adding to the net fiscal flows the scheduled coupon payments and issuances and redemptions of coupon-bearing securities, we have the weekly cash-need vector *f*. The Treasury ranged from needing $132 billion to having $85 billion over the twenty-seven-week period, as shown in Chart 2. The starting cash balance (c_0) is $66 billion and the latest issuance sizes are $30 billion, $30 billion, $31 billion, and $27 billion, for 4-week, 13-week, 26-week, and 52-week bills, respectively. For convenience, we also assume that the baseline issuance strategy $\widetilde{X}$(required to formulate the *minDevBase* objective) maintains the latest issuance sizes.

We need additional parameters to complete the model formulation. Regular bills follow a fixed issuance calendar: every Thursday for 4-week, 13-week, and 26-week bills and once every four Thursdays for 52-week bills. We assume that all have a minimum issuance size of $10 billion, and maximum sizes of $45 billion, $40 billion, $40 billion, and $30 billion, respectively; this defines *L* and *U* in constraint (2). The maximum absolute change in issuance size for consecutive issuances is set at $8 billion, $2 billion, $2 billion, and $1 billion, respectively; this is the δ in constraint (3). As such, we are willing to accept larger changes in 4-week issuances than in longer-maturity bills. The cash balance has a floor ($\underline{c}$) of $25 billion and a cap ($\bar{c}$) of $1,000 billion, a total that essentially leaves the weekly balance uncapped.

For the three dual objectives in (4c), (4d), and (4e), the choice of the weight ω is relevant. All scenarios examined in this article have a setting of $\omega = 10$. In the extreme, we can set ω to be a very large number, making the cash balance component irrelevant. In fact, we experimented with values of ω from 10 to 100 and found no material change in our main conclusion that the potential short-run cost of a regular and predictable approach is not likely to be significant.

3.2 Efficient Frontier

There are many issuance programs that satisfy constraints (1), (2), and (3). Among such "feasible" programs, we are interested in the subset of "efficient" programs that

Chart 2

Weekly Net Cash Need

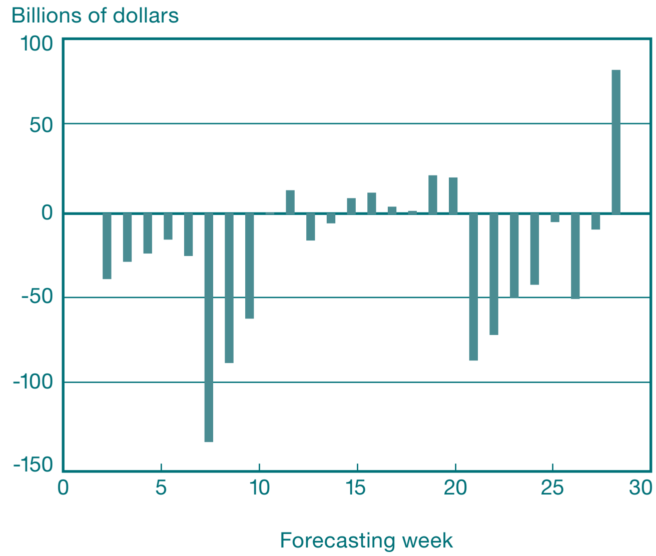

Sources: U.S. Department of the Treasury Office of Fiscal Projections (OFP); authors' calculations.

Notes: The OFP provides fiscal flow forecasts and calculates total cash flows assuming projected issuances of bills and coupon-bearing securities. By backing out the bill issuances projected by OFP, we calculate the cash needs that must be satisfied from new bill issuances. A negative number indicates the net cash that must be raised through new bills (excluding the minimum cash balance requirement). A positive number indicates the net cash flow from fiscal flows, coupon securities, and previously issued bills, even with zero new bill issuances. The chart shows the weekly cash needs in the twenty-seven weeks starting November 12, 2009.

minimize short-run funding costs (FC) for a given level of variability (AG) and minimize variability for a given level of funding costs.

The first step in determining the set of efficient issuance programs is identifying the global minimum cost (GMC) program—the issuance strategy that minimizes short-run funding costs subject to constraints (1), (2), and (3) or that solves the problem:

$$(5) \qquad \min_{x} \; hX$$
$$\text{subject to (1), (2), and (3).}$$

Given the inputs described in Section 3.1, the GMC issuance program has funding costs of $0.481 billion and variability of $3.82 billion.

Consider next the interval, $[\underline{G},\overline{G}]$, where the gradient of issuance has a floor ($\underline{G}$) of $0.0 and a cap ($\overline{G}$) of $3.82 billion. A feasible strategy on the efficient frontier solves the problem

$$(6) \quad \min_{x} \; hX$$
$$\text{subject to (1), (2), and (3) and } \sqrt{\frac{1}{N}(DX-d)'(DX-d)} = g$$

for some $g \in [\underline{G}, \overline{G}]$. We can trace the efficient frontier by solving (6) for different values of $g \in [\underline{G}, \overline{G}]$.[8] We illustrate the result in Chart 3, where the x-axis is the average gradient and the y-axis is the financing cost (both in billions of dollars).

The efficient frontier is downward sloping: The minimum attainable short-run funding cost falls when we allow larger week-to-week variations in issuance. At one extreme, when setting $g = 0$, we follow the latest realized issuance pattern, which in this example results in a large cost as a result of high cash balance.[9] As we allow the issuance strategy to become more flexible, we are able to identify programs with lower costs.

Chart 3
The Efficient Frontier

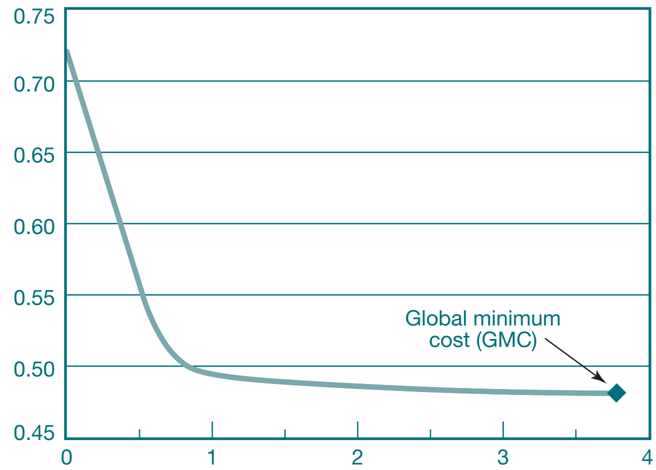

Source: Authors' calculations.

3.3 Alternative Strategies

The frontier clearly indicates a trade-off between short-run cost and variability. If an issuance strategy leads to a variability–cost pair that lies above the frontier, the indication is that, for the given level of variability, the strategy yields a cost that is higher than that of an efficient strategy. Hence, the vertical distance between the variability–cost pair and the frontier reflects the cost impact of including factors other than cost in the optimization objective. To find the optimal strategies associated with the five objectives (4a) to (4e), we solve the following quadratic programming problems:

$$(7) \quad \min_{x} \; objective$$
$$\text{subject to (1), (2), and (3),}$$

where *objective* is one of the objective functions (4a) to (4e).

We summarize the results in Table 2, with the columns reflecting the five alternative criteria and the rows reflecting the objective function actually used in the optimization problem. Because the criteria expressed in (4b) to (4e) return "sum square dollar" numbers (the sum of squared terms), for ease of comparison we report the corresponding root mean values, which may be understood as weekly averages. Panel B reports gross and net issuance amounts associated with each strategy.

The first column in Panel A shows financing costs for each of the five objective functions. The smallest cost—obtained with the *minCost* objective—is about $481 million, on $1.892 trillion total issuance. That figure represents a savings of $49 million from the most expensive strategy, which turns out to be *minGrossIss* with a total issuance of $1.748 trillion. So, if our objective were to minimize a combination of excess cash and gross issuance instead of minimizing cost, the optimal strategy would incur about 10 percent more funding cost during the twenty-seven-week period. In terms of annualized average auction yield, *minCost* produces an issuance cost of 4.90 basis points [interest cost divided by total issuance, or $(0.4810 / 1{,}892) \times (52 / 27)$], while *minGrossIss* incurs an annualized yield of 5.84 basis points, a difference of 0.94 basis points. As before, we stress that this result should be

[8] To generate the efficient frontier, (6) cannot be solved with quadratic programming because of the quadratic constraint. We used a general nonlinear constrained optimization package ("fmincon" in Matlab). It is the main issuance optimization problem, as set up in (7), that is solved efficiently with quadratic or linear programming.

[9] It is possible that we could not satisfy the cash needs by sticking with the latest actual issuance sizes at the beginning of the projection period; in that case, $g = 0$ would not be attainable. As our objective is to understand the general shape of the efficient frontier, we use a case that can accommodate a wide range of possible gradients.

Table 2
Values of Objectives

	Panel A: Values of Alternative Objective Functions (Billions of Dollars)					Panel B: Issuance (Billions of Dollars)	
Objective function used	(1)	$\sqrt{\frac{(2)}{27}}$	$\sqrt{\frac{(3)}{27}}$	$\sqrt{\frac{(4)}{27}}$	$\sqrt{\frac{(5)}{27}}$	Gross	Net
minCost	**0.4810**	81.740	157.623	81.975	84.599	1,892	(440.12)
minCB	0.5239	**69.868**	146.242	84.881	73.066	1,838	(439.93)
minGrossIss	0.5303	74.654	**141.581**	88.743	77.641	1,748	(439.95)
minDevBase	0.4918	73.042	149.155	**74.971**	76.181	1,862	(440.07)
minGrad	0.5240	69.868	146.142	84.718	**73.065**	1,836	(439.94)

Source: Authors' calculations.

Notes: The table reports the normalized objective-function values (columns) and issuance amounts associated with the five issuance strategies (rows) that result from minimizing (1) short-run financing cost, (2) cash balance, (3) cash balance plus gross issuance, (4) cash balance plus deviation, and (5) cash balance plus gradient, respectively, as defined in objectives (4a) to (4e). The scaling factor ω is set at 10 for the three dual objectives *minGrossIss*, *minDevBase*, and *minGrad*. Panel A reports the values of the objectives; Panel B reports the gross and net issuance amounts.

viewed as an upper bound on the additional cost, because it disregards changes in bidder behavior that would result from a more erratic issuance policy.

The second column of Panel A contains the average weekly cash balance, which differs by about $12 billion between the strategies that produce the highest cash balance (*minCost*) and the lowest cash balance (*minCB*). This outcome is interesting: It says that in order to minimize cost, we must efficiently explore relative cheapness across maturity terms and use the cheapest funding source as much as possible, likely resulting in occasional overfunding that will push up the Treasury's cash balance.

The bottom three rows contrast strategies with dual objectives, combining the desire for issuance smoothness (expressed in multiple ways) with the desire for low cash balances. The objective *minGrossIss* seeks to limit overall issuance amount, the *minDevBase* strategy attempts to follow the baseline trajectory, and *minGrad* looks to limit changes in issuance size. As expected, each of these strategies turns out to be more expensive than the pure cost-minimizing strategy, since they try to account both for issuance smoothness and cash balance smoothness. As mentioned above, *minGrossIss* is the most expensive strategy, incurring 94 basis points more in annualized auction yield than *minCost*.

Panel B reports the gross issuance amount of each strategy. Although cost is closely associated with gross issuance, the composition of the issuance is more important than its sheer size. For example, the *minGrossIss* strategy produces a total issuance of $1.748 trillion, or $88 billion less than the strategy with the second smallest gross issuance (*minGrad*) and $144 billion less than the one with the largest gross issuance (*minCost*). However, *minGrossIss* is also the most expensive strategy, with a financing cost of $530 million. On a net issuance basis, all five strategies result in roughly the same amounts and thus lead to a very similar end-of-period bill portfolio.

We show the variability–cost profiles of the five strategies relative to the efficient frontier in Chart 4. The *minCost* objective leads to the global minimum cost (GMC) represented in the chart by the blue diamond. The other four strategies turn out to be more costly, with lower issuance variability. As shown in the top panel, *minGrossIss* results in the highest cost and the lowest variability. The strategy has dual objectives, seeking to maintain a low (and by construction, smooth) cash-balance profile and low gross issuance. As such, it cannot effectively exploit relative pricing differentials across the yield curve; rather, it is forced to rely more on longer-term bills because they generate higher net cash than shorter-term bills for the same amount of gross issuance, owing to lower frequency of rollover. Relative to the pure cost-minimizing strategy, *minGrossIss* leads to a steadier issuance of longer-term bills, resulting in higher cost and lower variability.

The top panel also shows that when we try to stick with the baseline strategy with the objective *minDevBase*, we only incur $11 million more in cost than *minCost*, or less

Chart 4

Alternative Strategies and the Efficient Frontier

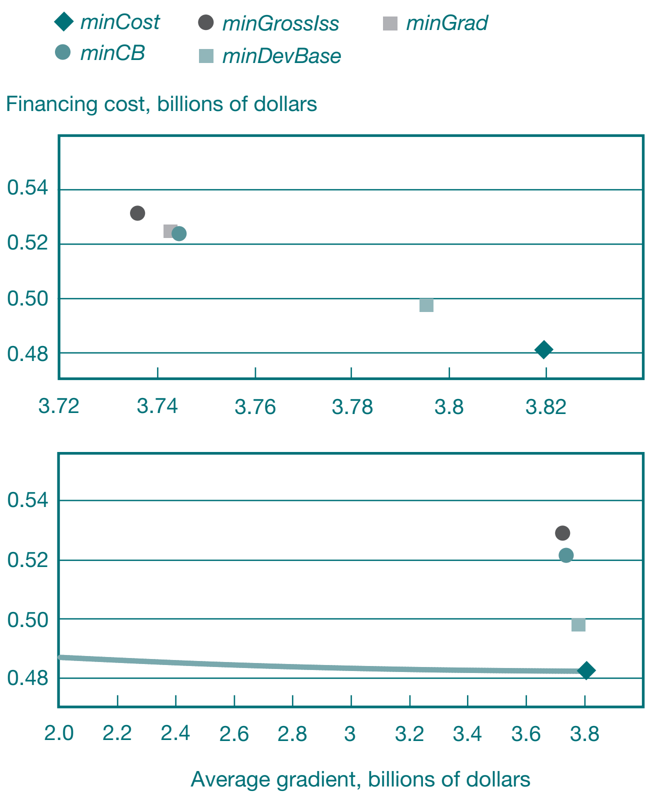

Source: Authors' calculations.

Notes: The top panel shows the variability–cost profiles of the five alternative strategies: minimizing financing cost (*minCost*), cash balance (*minCB*), cash balance plus gross issuance (*minGrossIss*), cash balance plus deviation (*minDevBase*), and cash balance plus gradient (*minGrad*). The bottom panel overlays these five points with a portion of the efficient frontier. Minimizing cash balance with or without consideration of gradient returns a very similar variability–cost profile (so *minGrad* partially overlaps *minCB*).

than 0.2 basis point in annualized auction yield. In fact, *minDevBase* is the least costly strategy among the four alternatives that target some form of predictability. This outcome suggests that stabilizing issuance at historical levels at each maturity point very nearly minimizes short-run costs.

The bottom panel shows the variability–cost profiles of the alternative strategies relative to a portion of the efficient frontier. Notice that these strategies all have fairly high *AG* values; compared with the feasible range of *AG*, the five alternatives are tightly clustered. (As shown at the end of Section 3.2, up to the GMC point, a higher *AG* is associated with a lower *FC*.) While the cost-minimizing strategy leads to the GMC point on the frontier, the other strategies are all above the frontier. The vertical distance between the alternative strategies and the efficient frontier thus measures the impact of factors other than cost in the objective function. In this example, a regular and predictable approach incurs an additional cost of $11 million to $49 million (or 0.19 to 0.94 basis point more in annualized auction yield) during a twenty-seven-week period. This estimate is subject to the same qualifications made earlier; it does not consider potential changes in bidder behavior that occur in response to changes in issuance strategies.

4. Improvement from Cash Management Bills

Various strategies using regular bills yield a spectrum of variability–cost trade-offs. In practice, the Treasury also has the option to use cash management bills (CMBs), which do not follow a fixed issuance calendar like other Treasury securities. Instead, the Treasury decides the auction time, issuance amount, and maturity for each CMB. In this section, we consider the advantages of using CMBs, propose a heuristic method to incorporate CMBs in issuance optimization, and examine the impact of CMBs on variability–cost trade-offs.

4.1 Rationale for Using CMBs

CMBs are useful for two reasons. First, the net cash flow from regular bills might be insufficient to meet the Treasury's cash need in a given week, even after the regular bill issuances are taken to the highest levels allowed by the size (2) and gradient (3) constraints. In this case, there would be no feasible issuance program using only regular bills, and we would need *feasibility* CMBs.

The other use of CMBs enables marginal improvements in a funding plan; we refer to these as *transitory* CMBs. As an illustration, the Treasury may anticipate a spike in cash needs in week i that is scheduled to be offset by a trough in week $i + 1$ through projected inflows. The Treasury may then choose to issue a 1-week CMB in week i to fill the gap, effectively shifting cash from the excess at week $i + 1$ to the shortfall at week i. CMBs command higher yields than regular bills; however, a short-term CMB may nevertheless have lower total cost, as measured by the product of yield and term. In addition, employing CMBs helps reduce the need to change regular issuance sizes and thus limits issuance variability. (We do not include CMBs in the calculation of AG variability because they are already penalized through a higher yield.)

4.2 A Heuristic CMB Algorithm

Before we include CMBs in the optimization procedure, we need to specify some model parameters in addition to those laid out in Section 3.1. Substantial evidence exists that CMBs are more expensive than regular bills of similar maturities. Simon (1991) finds that segmentation in the Treasury bill market is widespread and causes CMB yields to be higher than yields on adjacent-maturity bills. Seligman (2006) confirms the relative expensiveness of CMBs under the current uniform-price auction format and with off-cycle schedules (in other words, not necessarily conforming to the Thursday-to-Thursday issuance-to-maturity cycle). In our study, using CMBs in 2011, we find the average difference between the CMB yield and the implied yield based on the Treasury bill curve at the same maturity to be 4.57 basis points. Hence, to find the auction yield of a future CMB issuance, we use the bill curve to calculate the yield for the intended maturity, and then add 4.57 basis points. We also limit the term of a CMB to either one or two weeks, consistent with the actual usage of these securities since the introduction of 4-week bills. In addition, we limit the size of a CMB to between \$10 billion and \$40 billion.

Strictly speaking, incorporating CMBs in our issuance optimization turns the problem into a mixed-integer quadratic programming problem: For each week within the projection horizon, we need a binary variable that turns CMB issuance on and off. In fact, we need a separate binary variable for each potential CMB maturity. Unlike the original quadratic programming problem, this optimization problem quickly becomes prohibitively time-consuming. Instead, we determine the timing of CMBs through a simple heuristic.

In reality, the Treasury may consider CMBs if there is an unusually high cost associated with satisfying a particular weekly cash need with regular bills. To incorporate this intuition in issuance decision making, we use the Lagrange multiplier associated with the financing constraint (1) in solving the original problem (6). The multiplier, or shadow price, measures the change in the objective function owing to a marginal change in the constraint. It is calculated automatically as a byproduct of the quadratic programming algorithm. A high shadow price associated with the financing constraint in a given week indicates significant pressure on cash flow in that week, and thus a high benefit to issuing a transitory CMB. With this thought, we adopt a heuristic procedure to determine the timing of CMBs:

1. Add feasibility CMBs if necessary. Update the weekly cash-needs vector.
2. Run optimization (6) without transitory CMBs.
3. Identify the week in which the Lagrange multiplier associated with (1) is the largest. Run (6) with a 1-week CMB or a 2-week CMB in that particular week. Pick the CMB with the smaller objective function value.

Chart 5
The Impact of Cash Management Bills

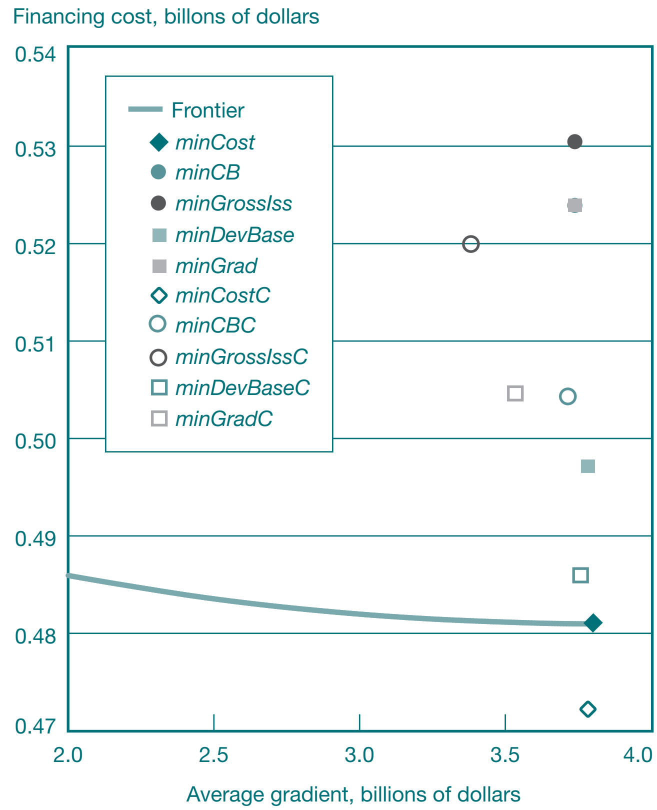

Source: Authors' calculations.

Notes: The chart shows three sets of results: the efficient frontier, five strategies using regular bills only, and five strategies using at most three cash management bills. The five strategies are minimizing financing cost (*minCost*), cash balance (*minCB*), cash balance plus gross issuance (*minGrossIss*), cash balance plus deviation (*minDevBase*), and cash balance plus gradient (*minGrad*). All CMB strategies exhibit better variability–cost trade-offs, positioned below and to the left of the same strategy using only regular bills. Note that *minGrad* overlaps *minCB*.

4. If multiple transitory CMBs are allowed, repeat step 3 until the maximum allowed number of CMBs is reached or until an additional CMB provides no material benefit.

Using CMBs improves the variability–cost trade-off, as shown in Chart 5 on page 11. We allow up to three CMBs in this example, erring on the conservative side of how the Treasury deploys CMBs in practice. The five objectives all return better variability–cost measures with the help of CMBs; they are below and to the left of those measures that use only regular Treasury bills. In addition, cost given a specific degree of variability no longer has a floor at the efficient frontier since we now have an expanded set of securities. Overall, using CMBs reduces cost by $8 million to $20 million, which equates to a 0.08-0.20 basis point reduction in annualized average auction yield.

5. Conclusion

Serving as collateral, hedging instruments, interest rate benchmarks, and safe and liquid investments, Treasury securities are essential to the functioning of the global financial system. In seeking to minimize its borrowing costs, the Treasury follows a principle of "regular and predictable" issuance of these securities. In this article, we attempt to quantify the short-term cost impact of forgoing opportunistic issuance of Treasury bills in favor of a regular and predictable approach. We do so by calculating the effect of including considerations other than short-term cost minimization when setting issuance sizes.

To overcome the practical obstacle of observing alternative strategies empirically, we use a model-based approach to compare tactical and regular issuance strategies. We quantify the cost impact of regular and predictable issuance as the cost difference, in dollars and in annualized auction yields, between a benchmark strategy focused on minimizing short-run costs and alternative strategies that include smoothness in the optimization objective function. To enable fast and efficient computation, we formulate the optimization problem as a quadratic program. We also examine how the inclusion of cash management bills would affect costs, using a heuristic approach based on shadow prices derived from the quadratic programming solution.

We find that taking a regular and predictable approach to issuance results in additional short-term costs. However, the additional cost is less than one basis point for most of the historical dates tested, and this increase is partially offset by using even a small number of 1-week or 2-week CMBs. Moreover, our analysis does not factor in changes in bidder behavior that would presumably result if a less regular and predictable issuance strategy were used; inclusion of those changes would likely further favor a regular and predictable approach. Thus, our overall conclusion is that the Treasury is not forgoing significant short-term gains by electing to follow a program of regular and predictable—rather than opportunistic—issuance of Treasury bills.

Appendix: Robustness Check

The analysis in this study, based on data as of November 12, 2009, suggests that various forms of regular and predictable issuance might effectively add 0.19-0.94 basis point in auction yield, but that introducing modest flexibility through limited issuance of cash management bills might drop the auction yield by 0.08-0.20 basis point, offsetting about a quarter of the cost impact of maintaining regular issuance. To check if these observations are robust, we examined a few dates since the beginning of the 2008 financial crisis. Each date roughly mirrors an important market development:

November 2007: pre-crisis
September 2008: crisis peak
March 2009: start of the first round of quantitative easing (QE1)
July 2009: total value of bills outstanding reaches a historic high
May 2010: a precipitous drop in bill amount outstanding
November 2010: QE2 begins
August 2011: fears of a breach in the debt limit and a downgrade of the U.S. government credit rating
October 2011: start of Operation Twist
March 2012: bill supply starts to increase
September 2012: QE4 begins

We use the fifteenth of each month (or the last business day before the fifteenth) for convenience.

We keep model parameters the same as described in Section 3.1. On each date, the yield curve and cash-need forecasts are the only differences. Interest rates changed substantially during the sample period; for example, the three-month rate dropped from around 3.4 percent in November 2007 to below 20 basis points in November 2008 and has stayed at the near-zero level ever since, reaching barely 10 basis points in September 2012. Financing needs varied as well; the gross bill issuance during the twenty-seven-week projection period was about $1.4 trillion in November 2007, climbed to $3.5 trillion in mid- and late 2010, and then ranged between $2.5 trillion and $3.1 trillion through 2011 and 2012.

We are interested in the cost of using objectives other than cost minimization and the advantage of using CMBs. The following table summarizes the cost and benefit, all in terms of effective auction yield changes in basis points. Column (I) is the maximum increase in auction yield (in basis points) from using objectives *minCB*, *minGrossIss*, *minDevBase*, and *minGrad*, while Column (II) is the minimum basis point decrease in auction yield when including CMBs with the same optimization objective. As before, for this exercise we only allow limited use of CMBs—at most three issuances in a twenty-seven-week period, each issuance having a maturity of one or two weeks.

Average Auction Yield Changes

Date	Maximum Increase in Yield Using *minCB*, *minGrossIss*, *minDevBase*, and *minGrad* (Basis Points)	Minimum Decrease in Yield Using *minCB*, *minGrossIss*, *minDevBase*, and *minGrad* with CMBs (Basis Points)
September 14, 2012	0.19	0.01
March 15, 2012	0.87	0.08
October 14, 2011	0.85	0.00
August 15, 2011	0.49	0.02
November 15, 2010	0.63	0.01
May 14, 2010	0.01	0.00
July 15, 2009	1.37	0.03
March 13, 2009	1.06	0.10
September 15, 2008	5.79	0.31
November 15, 2007	1.30	0.39

Source: Authors' calculations.
Notes: The table reports the changes in average auction yield attributable to using non-cost-minimizing objectives and CMBs (cash management bills) on ten historical dates around the time of the financial crisis. Relevant dates are given as the fifteenth of each month or the last business day before the fifteenth.

Consider September 14, 2012, for example. An issuance pattern that minimizes the cash balance while controlling for the gross issuance would effectively increase the annualized auction yield by about 0.19 basis point, or about $21 million in issuance cost, during the twenty-seven-week projection period. Other objectives also result in additional cost, though by smaller amounts. The flexibility of using CMBs generally leads to savings for the same objective, consistent with earlier observations that the option to use CMBs pushes the AG-FC trade-off down and to the left, as reflected in Chart 5. In this example, the smallest savings occur in *minGrad*, where the CMB option saves 0.01 basis point, or about $4 million.

Similar patterns hold for the other historical dates, though the maximum yield increase (additional issuance cost) or minimum yield decrease (issuance cost savings) do not always correspond to the same optimization objective. The magnitude of the additional cost or savings varies, with September 15, 2008, showing substantially larger values. The maximum yield increase, with *minGrad*, is about 5.79 basis points, or $120 million in additional issuance cost on a base of $5.74 billion. At the same time, using CMBs could save at least 0.31 basis point, or $38 million. These values are not inconsistent with the earlier results.

References

Adamo, M., A. L. Amadori, M. Bernaschi, C. La Chioma, A. Marigo, B. Piccoli, S. Sbaragli, et al. 2004. "Optimal Strategies for the Issuance of Public Debt Securities." International Journal of Theoretical and Applied Finance 7, no. 7 (November): 805-22.

Balibek, E., and M. Köksalan. 2010. "A Multi-Objective Multi-Period Stochastic Programming Model for Public Debt Management." European Journal of Operational Research 205, no. 1 (August): 205-17.

Bolder, D. J. 2008. "The Canadian Debt-Strategy Model." Bank of Canada Review, Summer, 3-16.

Date, P., A. Canepa, and M. Abdel-Jawad. 2011. "A Mixed Integer Linear Programming Model for Optimal Sovereign Debt Issuance." European Journal of Operational Research 214, no. 3 (November): 749-58.

Hahm, J., and J. Kim. 2003. "Cost-at-Risk and Benchmark Government Debt Portfolio in Korea." International Economic Journal 17, no. 2: 79-102.

Fisher, Peter. 2002. "Remarks before the Bond Market Association Legal and Compliance Conference." January 8, 2002.

Garbade, K. D. 2007. "The Emergence of 'Regular and Predictable' as a Treasury Debt Management Strategy." Federal Reserve Bank of New York Economic Policy Review 13, no. 1 (March): 53-71.

Garbade, K. D., and J. Ingber. 2005. "The Treasury Auction Process: Objectives, Structure, and Recent Adaptations." Federal Reserve Bank of New York Current Issues in Economics and Finance 11, no. 2 (February): 1-11.

Li, C., and M. Wei. 2012. "Term Structure Modeling with Supply Factors and the Federal Reserve's Large-Scale Asset Purchase Programs." Federal Reserve Board, Finance and Economics Discussion Series, 2012-37.

Seligman, J. 2006. "Does Urgency Affect Price at Market? An Analysis of U.S. Treasury Short-Term Finance." Journal of Money, Credit, and Banking 38, no. 4 (June): 989-1012.

Simon, D. 1991. "Segmentation in the Treasury Bill Market: Evidence from Cash Management Bills." Journal of Financial and Quantitative Analysis 26, no. 1 (March): 97-108.