Bank of Canada Banque du Canada

Working Paper 2003-10 / Document de travail 2003-10

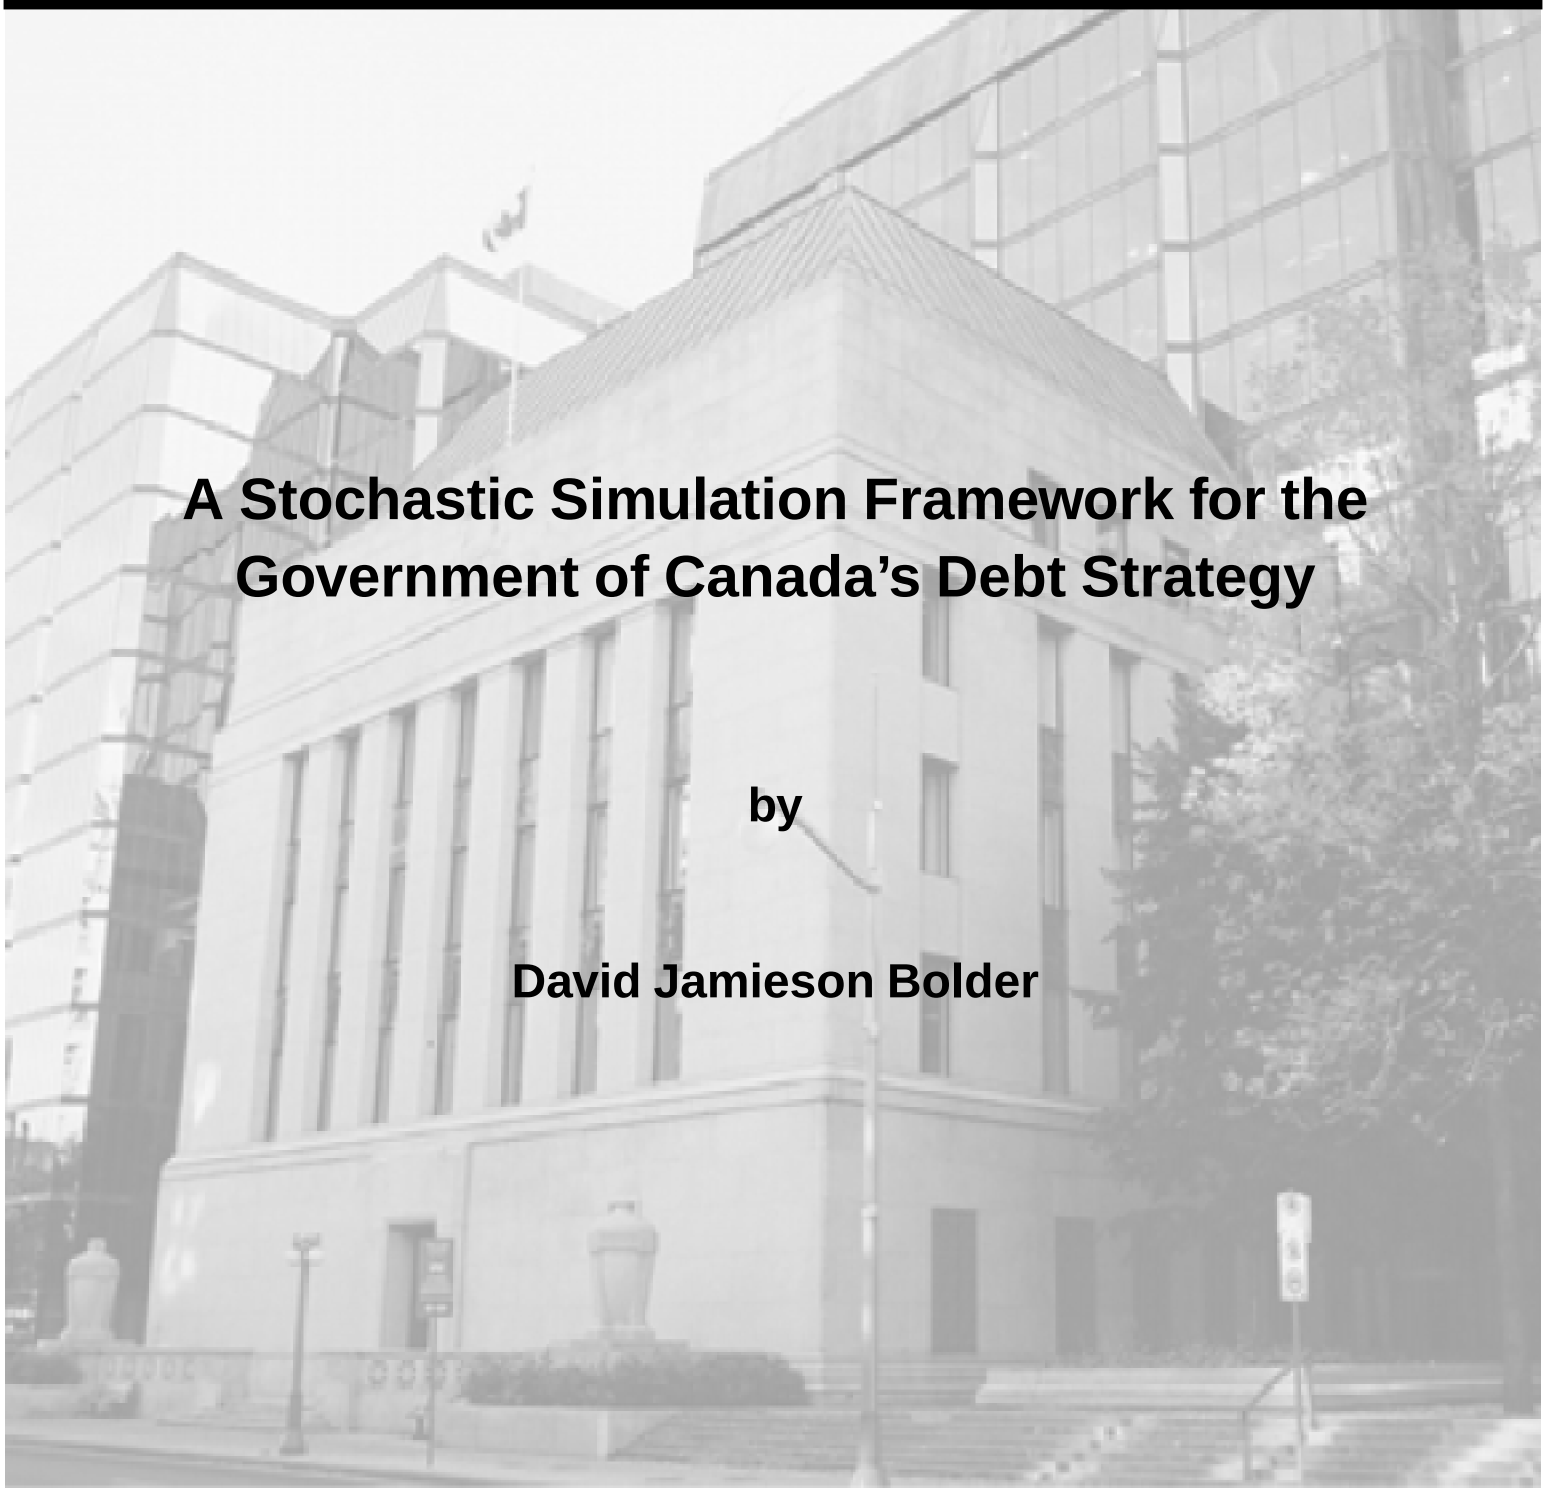

A Stochastic Simulation Framework for the Government of Canada's Debt Strategy

by

David Jamieson Bolder

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2003-10

April 2003

A Stochastic Simulation Framework for the Government of Canada's Debt Strategy

by

David Jamieson Bolder

Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
dbolder@bankofcanada.ca

Contents

Acknowledgements

I would like to thank, without implicating, Ron Morrow from the Bank of Canada and Rob Stewart from the Department of Finance Canada for making possible this fascinating project. I also appreciated the many fruitful (and always interesting) discussions on this topic with Pierre Gilbert from the Department of Finance Canada, Donna Howard and Philippe Muller of the Bank of Canada, Toni Gravelle of the International Monetary Fund, and Mark Reesor of the Applied Mathematics Department of the University of Western Ontario. Finally, I would like to particularly thank Etienne Lessard from the Bank of Canada for his extremely helpful contribution in assisting with the construction of the simulation model. As always, I alone am entirely responsible for any and all errors, omissions, and inconsistencies that may appear in this work.

Abstract

Debt strategy is defined as the manner in which a government finances an excess of government expenditures over revenues and any maturing debt issued in previous periods. The author gives a thorough qualitative description of the complexities of debt strategy analysis and then demonstrates that it is, in fact, a problem in stochastic optimal control. Although this formal definition is conceptually useful, the author recommends the use of simulation to help characterize the set of strategies that a government can use to fund its borrowing requirements. He then describes in detail a stochastic simulation framework, building from previous work in Bolder (2001, 2002); this framework forms one important element in the debt strategy decision-making process employed by the Government of Canada. The primary objective in constructing this stochastic simulation framework is to learn about the nature of the risk and cost trade-offs associated with different financing strategies. To this end, the paper includes a detailed description of the model; a set of possible debt cost and risk measures, including one potentially useful conditional risk measure; illustrative results under normal stochastic conditions; an analysis of the sensitivity of the results to various key model parameters; a novel approach to stress testing; and a possible framework for selecting a financing strategy, given assumptions about government risk preferences.

JEL classification: C0, C15, C52, H63

Bank classification: Debt management; Econometric and statistical methods; Economic models

Résumé

Cette étude en deux volets porte sur la stratégie de la dette, qui peut se définir comme la façon dont le secteur public finance l'excédent de ses dépenses par rapport à ses recettes, ainsi que sur le remboursement des dettes arrivant à échéance. Dans le premier volet, l'auteur fait ressortir la complexité de l'analyse de la stratégie de la dette, puis il montre que ce type d'analyse est par définition un problème de contrôle optimal stochastique. Puisqu'une telle définition n'est pas très utile en pratique, malgré les avantages qu'elle présente sur le plan conceptuel, l'auteur recommande d'utiliser des simulations pour déterminer les caractéristiques d'une stratégie globale de financement du secteur public. Dans le second volet, l'auteur reprend et expose en détail un cadre de simulation stochastique proposé dans deux de ses études précédentes (Bolder, 2001 et 2002); ce cadre est un élément important du processus de décision associé à la stratégie de la dette du gouvernement canadien. On trouve dans ce volet une description détaillée du modèle; un ensemble possible de mesures des risques et des coûts (notamment la possibilité d'une mesure conditionelle des risques et des coûts); des résultats, donnés à titre d'illustration, de calculs effectués dans des conditions stochastiques; une analyse de sensibilité des résultats à divers paramètres; une nouvelle méthode d'essai sous contraintes et enfin un cadre possible de choix d'une stratégie de financement, lequel est établi à la lumière de diverses hypothèses sur les préférences en matière de risques.

Classification JEL : C0, C15, C52, H63

Classification de la Banque : Gestion de la dette; Méthodes économétriques et statistiques; Modèles économiques

1 Introduction

Debt strategy is defined as the manner in which a government finances an excess of government expenditures over revenues and any maturing debt issued in previous periods. The question concerns the best way for the government to borrow these required funds. Should it, for example, use short-term debt, such as treasury bills or longer-term coupon bonds? Interestingly, an extensive academic literature on this subject does not exist. There are some exceptions. Missale (1994) studies the relevant aspects of debt management and demonstrates, in economic terms, that the composition of a government's debt portfolio actually matters but, unfortunately, offers little practical advice for debt managers.[1] Barro (1995) presents a *tax-smoothing* model as the basis for optimal debt management. This academic work notwithstanding, our perspective is more practical. Our analysis is based on the belief that a sustainable and prudent debt structure is critical for any sovereign nation. Moreover, we take the government's fiscal policy as given and attempt to characterize the set of financing strategies that have desirable risk-cost characteristics. Indeed, our primary objective in this work is to learn more about the nature of the risk and cost trade-offs associated with different financing strategies. The practitioner literature relating to better understanding this issue is found in publications from sovereign debt managers.[2]

Adopting this pragmatic perspective, we demonstrate that one can conceptualize the government's borrowing decision as an optimal-control problem in a stochastic setting. This problem has been extensively studied in the asset-pricing setting where an investor attempts to optimally select the proportion of risky and riskless assets that maximize their expected utility subject to appropriate wealth constraints (for example, Karatzas and Shreve 1998). In our situation, the government is attempting to optimally select the composition of its debt portfolio to minimize expected debt costs subject to risk and liquidity constraints. This useful mathematical definition of government debt strategy will be formalized later in the paper. Given practical complexities, however, it is not obvious how to use dynamic programming techniques to find a solution. Instead, we rely on simulation. In this sense, we extend the preliminary work begun in Black and Telmer (1999). We have also found that a simulation methodology termed *dynamic financial analysis* in the actuarial science literature is relevant for this task. Insurers are often faced with the problem of trying to set premiums and capital reserves, given stochastically evolving claims and investment returns. Structurally, the techniques used in dynamic financial analysis are relevant for our work in debt strategy analysis.[3]

[1]Under certain conditions on tax payments, one can argue that the division of a government's financing between taxation and debt is irrelevant.

[2]Excellent references in this area include Hörngren (1999), Danish Nationalbank (1998), Holmlund (1999), Holmlund and Lindberg (2002), Bergström and Holmlund (2000), Danmarks Nationalbank (2000, 2001), and the IMF and the World Bank (2001).

[3]Specific references that have been constructive in our work include Kaufmann, Gadmer, and Klett (2001), Cumberworth et al. (1999), and the Dynamic Financial Analysis Committee of the Casual Actuary Society (1999).

Because our approach to debt strategy analysis involves stochastic simulation, our secondary objective is to present the details of this simulation framework. Moreover, it is our view that management of the government's borrowing program is an important and difficult task requiring a combination of judgment and comprehensive analytical tools. The debt strategy simulation framework suggested in this paper is one such tool. This paper *neither* describes an optimal debt strategy for the government *nor* outlines the Government of Canada's current debt strategy. Instead, its focus is on the methodology of the proposed simulation framework that is used to compare and contrast alternative financing strategies. The more realistic and flexible the tool, the more successful we will be in achieving our primary objective of better understanding the risk and cost trade-offs of different financing strategies. By placing a discussion of this methodology in the public domain, we hope to engender debate and commentary from others about the defensibility of our assumptions and approaches. This should help us to construct a more complete and realistic model.

This paper is divided into three main sections. The first section describes the nature of the debt strategy problem in both qualitative and quantitative terms, to introduce the various aspects of the government's borrowing program and provide a formal mathematical definition of the problem. We proceed, in the second section, to describe in detail our simulation framework. This includes a brief discussion of the stochastic environment constructed in Bolder (2001, 2002), an overview of how the issuance of debt is modelled over time, and a formal presentation of a set of measures of cost and risk used to distinguish between various financing strategies. In the third and final section, we examine some illustrative applications of our model to characterize the risk and cost trade-off associated with a variety of alternative financing strategies. The section includes an examination of results under stochastic conditions with our base parameterization, an analysis of the sensitivity of the results to various key model parameters, a novel approach to stress testing, and a comparison of a large number of possible financing strategies. These results assist in the formulation of a suggested approach for selecting a financing strategy, given assumptions about government risk preferences.

2 Background

Before we can discuss our debt strategy model, it is important to understand the principal ideas involved in the formulation of government debt strategy. This section describes the concept of debt strategy in a clear, straightforward, and conversational manner while working from first principles.

2.1 The key issues

The idea behind debt strategy analysis is deceptively simple, but in fact subtle. It starts with some basic ingredients and poses a question. These ingredients include an existing *stock of debt* denominated in a variety of instruments (treasury bills, coupon bonds, real-return bonds, retail savings bonds, etc.) and a sequence of *financial requirements.*[4] Not surprisingly, this sequence of financial requirements is a forecast the reliability of which decreases as we move further from the current point in time. Generally, this forecast continues out over a five- to ten-year time horizon; given the nature of the budgetary process, we typically focus on an annual frequency. By combining the forecast annual financial requirements with the associated annual *rollover* or *refinancing* of the existing debt, we determine the annual *borrowing requirement.* Here, we begin to see some of the complexity of the problem: the actual debt refinancing in future periods will depend on the financing strategy employed in previous periods. For example, consider a two-year scenario, where the first year's financing requirement is funded entirely with either one-year treasury bills or 10-year coupon bonds. The implications for debt issuance in the second year are quite different under these two financing strategies. In particular, the one-year treasury-bill strategy requires all of the first year's issuance to be refinanced in the second year, whereas the 10-year coupon bond strategy does not. This is one indication of how the time dimension adds complexity to the problem.

These basic ingredients in hand, the question is: how should the government finance this series of annual borrowing requirements? Clearly, there are an infinite number of ways to accomplish this task. We need some way to distinguish between various strategies for financing the sequence of borrowing requirements. One obvious approach is to select a *financing strategy* that has the lowest cost for the government and, ultimately, for the taxpayer.[5] To compute the cost of a given financing strategy is nevertheless more complex than it might initially appear, because the set of financial instruments used to meet the government's financing requirements are fixed-income securities and thus depend importantly on future interest rates. This is a difficulty, given that we have no a priori knowledge of future interest rates. We now begin to see a second element of complexity in this analysis. In particular, there is a continuum of future term structures of interest rates taking an infinite number of possible future forms. Moreover, each of them has some positive probability of occurrence. Fortunately, we can construct models of the term structure of interest rates that reflect some of its more important empirical properties and are also consistent with some key financial relationships, such as the absence of arbitrage opportunities. That is the good news. In fact, a previous work (Bolder 2001) describes how the class of affine term-structure models can be applied to the debt strategy problem. The bad news is that modelling the dynamics of the term structure of interest rates is a tricky business. One needs to be cautious on the theoretical side as well as with the statistical analysis required to find parameters for these models. The consequence is that Bolder (2001) is not the final word on this subject of interest rate modelling, but rather a reasonable starting point.

[4]Of late, in many countries, including Canada, these requirements have been negative, implying a government surplus position; deficits, of course, are also possible.

[5]We realize that we have not, as yet, formally defined the concept of a financing strategy, but we subsequently discuss it in substantial detail.

Some shortcomings are associated with assessing the desirability of a given financing strategy by focusing on a single dimension of cost. As stated earlier, we do not know with certainty the future evolution of the term structure of interest rates. As a consequence, we cannot make definitive statements about the annual interest cost of a given financing strategy. At best, we can quantify—conditional on our model of the term structure of interest rates—what we *expect* the cost of a given financing strategy to be over some time interval. Essentially, this expected (or average) cost reflects our *uncertainty* regarding future time periods. Uncertainty in government budgeting and planning, however, is also problematic. In fact, the stability of a government's ability to fund itself must be a critical consideration in any complete analysis of a financing strategy. For example, we may be able to say that the expected cost of a given financing strategy is low; as stated earlier, this is an obviously desirable characteristic.[6] If, at the same time, we state that our uncertainty about this expectation is high, then we must consider this financing strategy as being somehow *less* desirable. How *much* less desirable is this financing strategy? The answer will depend, of course, on how uncertain we are about our estimate of the cost and how this uncertainty compares with other financing strategies. This discussion implies that in some future states of the world (which have a positive probability of occurrence) the *actual realized cost* of that financing strategy will be much higher than what we expect to occur on average. Thus, it is essential to consider the *variance* of our cost estimate when investigating a potential financing strategy. Indeed, the concept of variance is often termed, in an operational sense, *risk*. Therefore, a particular financing strategy is *risky* when its expected cost has a high degree of uncertainty or, alternatively stated, a large variance. As a consequence, in our analysis we need to consider not only the expected cost of a given financing strategy but also its risk.

To complicate matters further, there tends to be a *tension* between the expected cost of a financing strategy and its risk, which stems from the basic underlying characteristics of the term structure of interest rates. More specifically, we typically observe an upward-sloping term structure, which implies that, on average, the annual cost of borrowing funds for a short period of time, say three months to two years, is less expensive than borrowing for a longer period of time, such as 10 to 30 years. Unfortunately, the situation is, for at least two reasons, somewhat more complex than it might appear. First, when one borrows funds for a short period of time, those funds must be refinanced in the relatively near future. Second, the short end of the term structure tends to be substantially more variable than the long end; in fact, there are occasions when the term structure will flatten, or even invert, for extended periods of time. While the first point is obvious, when combined with the second point it has important implications. Financing strategies that involve substantial amounts of short-term borrowing will tend to be less expensive relative to longer-term financing. But short-term funds require significantly more refinancing under potentially adverse conditions.

[6] Or perhaps it is not so obvious when considering the government as an infinitely lived organization. In this discussion, we assume that variable debt charges are negative from the government's perspective. For a more detailed economic argument as to why this might be a reasonable working assumption, see Missale (1994).

The consequence is greater cost uncertainty, or risk, associated with low-cost financing strategies. That is, we should not expect to find many financing strategies that simultaneously reduce both expected cost and risk. We do not fully understand the actual trade-off between cost and risk over a wide range of financing strategies, but this study aims to develop a better understanding of this cost-risk relationship. If we can do this, then we will be a step closer to an approach that involves finding the lowest expected cost for a given level of risk. Even better, we might find that certain financing strategies dominate others in terms of *both* cost and risk.

The randomness of the evolution of the term structure of interest rates is not the only random element in this analysis. As stated earlier, one important component of the financing requirement relates to the government's surplus or deficit position. A crucial question is, how is this quantity determined? The government's annual financial operating plan, the *budget*, depends importantly on the government's *receipts* and its *expenditures*. The interaction of these two elements, taking into consideration any non-cash accounting transactions, determines the deficit or surplus position for a given year. While the government makes detailed plans with respect to these receipts and expenditures, there remains a non-trivial amount of uncertainty which, again, increases as we move further out in time. Indeed, it is an incredibly complex process. We can, nevertheless, make several general statements. For instance, the government tax receipts depend substantially on the state of the economy. In recessionary periods, tax receipts will tend to fall, while they will typically rise during periods of strong economic growth. Government programs, which constitute the bulk of the government's expenditures, also exhibit a business-cyclical pattern. The bottom line is that the government's surplus position will depend, in some manner, on the general *macroeconomic* conditions prevailing in the economy during that period. Again, as with interest rates, we do not have a priori knowledge of future macroeconomic conditions. This introduces another source of uncertainty into our modelling exercise. Moreover, we also know that the term structure of interest rates is not independent of the macroeconomy. In particular, we empirically observe a steep term structure preceding periods of economic expansion and a flatter, or even inverted, term structure prior to recessionary periods.[7] Thus we should, in principle, observe a link between the surplus position of the government, the evolution of the term structure of interest rates, and the general macroeconomy. This matter is of particular importance given the long-term horizon of the debt strategy problem, which correspondingly unfolds over multiple business cycles. The consequence is that an approach is required to capture this additional element of uncertainty in our model. Bolder (2002) outlines a reduced-form model for the joint evolution of the term-structure of interest rates, the macroeconomic business cycle, and the government's financial position.

Using the models in Bolder (2001, 2002), the next step requires the computation of the average, or expected, cost of a given financing strategy and the riskiness of this estimate. A measure of expected cost

[7] In short, the term structure actually serves as a good *leading* indicator of economic activity.

and its associated risk is *not* necessarily sufficient to distinguish between various financing strategies, since there is an additional dimension to the problem. To understand this, we need to consider the intertemporal nature of our analysis. In fact, it is a careful consideration of the time dimension that suggests that measures of cost and risk alone may not adequately describe the government's exposure to some uncertain outcomes. In particular, a government may not wish to engage in a disproportionate amount of financing in any given period. Consider, for example, an admittedly simplistic, three-period financing strategy that requires a government to borrow $10 billion in the first period, $30 billion in the second period, and $20 billion in the final period. This might be considered inferior to a financing strategy of $20 billion each period. The reason for the first strategy's inferiority is the concentration of issuance in the second period. This lack of *time diversification* could expose our issuing sovereign to poor borrowing conditions (i.e., high interest rates) during this second period. Depending on the composition of the issuance in the second period, it may additionally expose the government to substantial refinancing risk in future periods. One possible measure to capture the time diversification of a given portfolio is the proportion of the total debt that must be financed in the upcoming year. This measure can be proxied by the ratio of floating debt stock to the total debt stock. Floating, in this context, means that part of the debt matures in the next twelve months. Traditionally, the Canadian government has looked at a simple transformation of the floating debt ratio termed the *fixed-debt ratio*. This measure provides some insight into the intertemporal debt issuance trade-off, or time diversification, made by a given financing strategy.[8]

To this point, in addressing our question, we have reviewed the sources of uncertainty and three potential measures for distinguishing between financial strategies: expected cost, risk, and time diversification. We have not as yet discussed the financing strategies themselves, because a financing strategy can be quite complicated. Complexity aside, financing strategies are of paramount importance; they represent the *only* element of our analysis that is under the control of the debt manager. As a consequence, they demand serious investigation. In its simplest form, a financing strategy indicates how much of the annual borrowing requirement to allocate to a given debt instrument. For example, with an annual borrowing requirement of $20 billion, one financing strategy might be 50 per cent in one-year treasury bills, 25 per cent in two-year coupon bonds, and a final 25 per cent in 10-year coupon bonds. Another financing strategy might prescribe the issuance of 33 per cent of the $20 billion, or $6.7 billion, in each of these three debt instruments. These types of financing strategies are called predetermined, or *deterministic*, because they are simple rules, known at the initial point in time, that do not require knowledge of future outcomes for their application. Of course, more complicated potential financing strategies could be followed. In particular, one might look at the term structure of interest rates in each period and use a specific value to create a decision rule. For instance, one might decide to meet the annual borrowing requirement entirely with 3-month treasury bills

[8] This is by no means a perfect measure of refinancing risk, or what we have been calling time diversification. It is, nevertheless, a reasonably intuitive concept, and thus serves as a good starting point.

if the one-year zero-coupon rate is below 4 per cent, but use exclusively two-year coupon bonds otherwise. This may seem a strange financing strategy but we have selected it to raise a point: we do not initially know what this rule will imply in terms of specific issuance. Actual issuance will depend upon the evolution of a random variable, the one-year zero-coupon rate. Thus, this is an example of a non-determinstic or *stochastic* financing rule. This is the most interesting type of financing strategy, but also the most difficult to handle in an analytical setting. It seems, therefore, that we have another entry for our checklist; in particular, we need the necessary analytical machinery to consider a wide array of financing strategies. The difference in difficulty between deterministic and stochastic financing strategies is such that we believe it makes sense to enter them separately on our list. That is, it makes sense to work towards understanding financing strategies of a deterministic nature first, and then progress towards the more complicated, stochastic strategies later if that proves necessary.

Not all financing strategies are *permissible*. A permissible strategy typically has a technical mathematical definition. For our purposes, it relates to a more general issue about the nature of financial markets. In particular, we have suggested that virtually any combination of debt instruments is acceptable to meet the government's annual borrowing requirements. This is not strictly true. The complication arises because there exists a large and active secondary market of government securities for a number of maturities across the term structure of interest rates. To maintain the liquidity of this marketplace—which is in the issuer's financial interest—the government must continue to issue at each of these key maturities.[9] This is not to say that there is no flexibility in the set of financing strategies that we may examine, but rather that our options are constrained by the desire to maintain liquid and well-functioning government security markets. The greater the liquidity of the secondary Government of Canada markets, the larger the demand for these securities and hence the lower the borrowing costs for the government. On a related note, the government fixed-income markets serve as a benchmark for a host of provincial and corporate issuers. In this way, the efficiency of the underlying government securities markets has implications for the entire financial market. Additionally, both the implementation and transmission of monetary policy occur through financial markets. Thus, to the extent that financial markets are liquid and well-functioning, monetary policy will tend to be more efficient. The larger point is that we must not lose sight of the realities of the underlying financial and economic variables that we are attempting to model.

There are two additional complicating factors. First, one cannot assume that the government's borrowing requirement is independent of the financing strategy. Government debt charges form a significant component of the government's annual expenditures. If the government engages in a relatively higher-cost,

[9] Liquidity, in this setting, is defined as the ability of a market participant to quickly buy or sell a security with minimal transactions costs and impact on the price of the security. This idea is very well expressed in Bennett, Garbade, and Kambhu (2000). In particular, they state in an analysis of the American situation that "maximizing liquidity in the Treasury market is coincident with minimizing the Treasury's long-term interest expense."

lower-variability financing strategy, then government expenditures will increase. That is, while the financial requirements may be less volatile on average, they will tend to increase government expenditures relative to revenues.[10] Conversely, a financing strategy consisting of entirely short-term debt may increase the variability of government expenditures and hence the variability of the government's financial requirements. This previous statement, however, is *not* true in a general sense. We will address this in a moment. Any reasonable simulation model will need to incorporate this feedback between the financing strategy and the government's financial requirements. The second complicating factor is also a form of feedback. Namely, alterations in the government's financing strategy will have implications for the cost of financing in each debt instrument. As stated in the previous paragraph, the cost of issuing a given debt instrument depends largely on the amount of that instrument currently outstanding (i.e., its liquidity). A financing strategy involving minimal issuance in a given debt instrument will be relatively costly. Conversely, very large issuance in a given sector will be difficult for the market to digest. There is, for example, only a limited appetite for 10-year Government of Canada coupon bonds. To place very large amounts of issuance in a given instrument will imply higher costs for the government. In short, there is an issuance range for each maturity sector that maximizes liquidity without creating a situation of oversupply. This issue is relevant to a sovereign borrower, but need not typically be considered by a corporate entity. Ultimately, a corporate entity issues in a sufficiently small size that it can be considered a price-taker in debt markets. The relatively large size of issuance by sovereign borrowers implies that moderate alterations in financing strategy can influence the relative rates at which they can borrow.[11] Again, to facilitate a fair comparison between various financing strategies, this relationship between financing strategy and the term structure of interest rates must be explicitly modelled.

As stated earlier, debt-cost variability does *not* immediately translate into budgetary volatility. To see this, we note that the government's financial requirements can be represented in the following form,

$$F_t = \underbrace{\eta_t - s_t}_{\substack{\text{Primary}\\ \text{balance}}} - c_t, \tag{1}$$

where,

$$\begin{aligned}
\eta_t &\triangleq \text{government tax revenues in period } t,\\
s_t &\triangleq \text{government spending in period } t,\\
c_t &\triangleq \text{debt-service charges in period } t.
\end{aligned}$$

[10] The result might be to put a government that is in a surplus position into a deficit, or to push a government already in a deficit even further into deficit.

[11] Not all sovereigns currently find themselves in this situation. Australia, for example, currently has a very small debt-to-GDP ratio, and may potentially exit debt markets altogether.

Equation (1) is relatively straightforward to interpret. If program spending and debt-service costs exceed tax revenues in a given period, then the government will be in a deficit position. This shortfall will have to be met with borrowing. If tax revenues are in excess of government expenditures, the government will be in a surplus position. The government is, quite understandably, concerned with keeping itself in a balanced-budget position.[12] That is not to say that deficits are undesirable, but rather that extended periods of deficit financing cannot be sustained. The key concern, for a government with a net-debt position, is that variability in debt-service costs will contribute to variability in the overall budget balance—excessive variance, for example, can contribute to a *vicious circle* in the budgetary cycle. We can explicitly write the financial position variance as,

$$\mathrm{var}(F_t) = \mathrm{var}(\eta_t - s_t) + \mathrm{var}(c_t) - 2\mathrm{cov}(\eta_t - s_t, c_t). \tag{2}$$

An inspection of equation (2) reveals that while the debt-service cost volatility ($\mathrm{var}(c_t)$) contributes to the variance of the government's financial position, the covariance of the debt-service costs and the primary balance is also quite important. In particular, if this covariance ($\mathrm{cov}(\eta_t - s_t, c_t)$) is sufficiently positive, then it will act to dampen debt-service cost volatility.

The key question, therefore, relates to the sign and magnitude of the covariance term. Is this covariance indeed positive? The answer is yes, in most circumstances, the covariance between the primary balance and debt-service charges is positive. Consider an expansionary period, where the economy is growing above capacity. In this case, tax revenues increase and program spending falls, contributing to a larger primary balance. At the same time, the monetary authority will increase short-term interest rates, thereby leading to higher debt-service costs.[13] The result is positive co-movement between these two macroeconomic variables. It also applies in the opposite direction. Easing of monetary conditions during recessionary periods will contribute to lower debt charges as the primary balance falls due to weaker tax revenue and increased program spending.

There are four main points to note about this relationship:

(i) Central banks typically tighten and ease monetary conditions in a gradual manner. Moreover, monetary policy operates with a lag of 12 to 18 months and, as such, a monetary authority will formulate and implement monetary policy on a forward-looking basis. The consequence is that while there may be a positive primary balance and debt-service cost covariance, its magnitude may

[12] The fiscal objectives of a particular government may vary, of course, depending on their debt-to-GDP ratio. A government that is a large net-debtor may, for example, attempt to maintain a surplus position over the medium term, to pay down their debt stock.

[13] This is something of a simplification, because tightening, or easing, of monetary conditions does not impact each sector of the term structure of interest rates in an equal fashion. Generally, particularly when inflationary expectations are anchored, the short end of the term structure is more sensitive to changes in monetary policy.

be relatively modest. Figure 1 demonstrates that, for sufficiently variable debt-service charges, a fairly substantial degree of positive covariance is required to dampen out the volatility in the government's financial position.

(ii) The size of the covariance depends on the composition of the debt portfolio. If, for example, the majority of the debt stock consists of long-term instruments then, its sensitivity to interest-rate changes will be dramatically reduced. A debt stock entirely denominated in treasury bills, conversely, will exhibit a greater degree of covariance.

(iii) The surplus-deficit position of the government also plays a role in this relationship. If the government is in a deficit position, all else being equal, the weight of debt-charge volatility will play a more important role in the volatility of the government's financial requirements.[14] A surplus position will have the opposite effect. Figure 1 illustrates this relationship for a balanced budget. From a risk-management perspective, a debt manager should be concerned about the potential for a deficit situation to erode the impact of positive covariance between debt charges and the primary balance.

(iv) The most important point is that there are economic examples where this relationship clearly does *not* always hold. In particular, if shocks to the economy arise from the supply side, then the previously described relationship between primary balance and debt-service charges can reverse. A supply-side shock—for example, a dramatic increase in oil prices—can be inflationary and lead to weaker economic conditions. In this situation, the primary balance will deteriorate, but interest rates will remain high.[15] The resulting *stagflation*—which was experienced in the 1970s—can have dramatic effects on the government's financial position. Moreover, the shorter the maturity structure of the debt portfolio, the worse the impact. Figure 1 reveals that negative covariance can contribute to significant increases in the volatility of the government's financial position. A supply-side shock is not a high-probability event, but it can have dramatic effects if it does occur.

The consequence is that the relationship between the variance in the government's financial position and the financing strategy is more complex than it might at first appear.[16] This analysis, however, does *not* consider supply shocks. Furthermore, our stochastic environment is structured such that positive covariance

[14]To see this, recall that, for computational purposes, we actually rewrite equation (2) as,

$$\mathrm{var}(F_t) = \omega_1^2 \mathrm{var}(\eta_t - s_t) + \omega_2^2 \mathrm{var}(c_t) - 2\omega_1\omega_2 \mathrm{cov}(\eta_t - s_t, c_t), \tag{3}$$

where ω_1 and ω_2 represent the weights of the primary balance and government debt charges, respectively. Mathematically, a government deficit implies that $\omega_1 < \omega_2$, while a government surplus implies the converse.

[15]Indeed, the monetary authority may even be forced to tighten monetary conditions to stave off inflationary pressures.

[16]I would like to thank Patrick Georges from the Department of Finance Canada for clarifying my thinking about this important issue.

between the primary balance and debt-service charges is evident. In future work, we wish to expand this to consider a broader range of outcomes.

Figure 1: **Financial Position Variance**: This graph illustrates the variance of the government's financial position for a primary balance standard deviation of 10 per cent and four possible debt-service cost standard-deviation values: 5 per cent, 15 per cent, 20 per cent, and 25 per cent. In each case, the primary balance and debt-service charge correlation values range from -1 to 1. For the purposes of this computation, debt charges and the primary balance are assumed to contribute equally (i.e., the budget is balanced) to the government's financial position standard deviation.

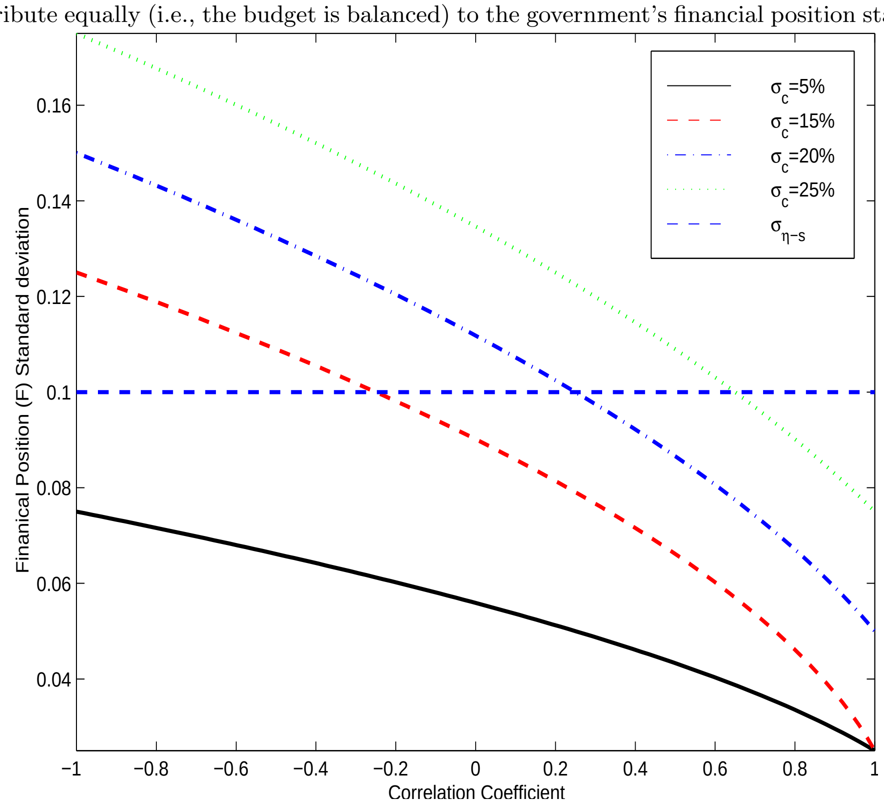

Therefore, while one might at first think that the government's debt strategy is merely a question of issuing debt at the lowest cost to cover its borrowing requirements, the situation is more complex. In particular, the existing debt stock, the evolution of the term structure of interest rates, and the way in which the macroeconomy affects these interest rates as well as government revenues and expenditures need to be considered in depth. We also must be able to use these random inputs to consider a variety of different financing strategies in terms of expected cost, risk, and time diversification. We must keep in mind the various real-world restrictions imposed on us by the need to maintain liquid and well-functioning government debt markets. It will be useful to keep these concepts in mind as we progress to the following sections of this discussion. Indeed, our goal in this paper is to create a road map that can aid us in finding a solution to our original question. In doing so, we want to define debt strategy in both qualitative and quantitative terms. To this point, we have discussed the problem in a qualitative sense; we will move to put the question on a more mathematical basis. While the nature of our question is such that we cannot analytically solve this problem, its formal mathematical definition provides substantial insights. We then abstract from our formal mathematical definition and consider the structure of our simulation framework.

2.2 An optimal-control problem

Having discussed debt strategy in broad qualitative terms and reviewed some of the requisite motivation for our analysis, we will attempt to put our problem on a mathematical footing. In its purest form, our question falls into the class of *stochastic optimal-control* problems. Appendix A sets out a more thorough definition of our problem, with the aim of developing some intuition for those who want more detail. The goal of this section is not technical. Instead, it is intended to highlight the sources of difficulty that we will confront in analyzing the problem.

Our problem, to paraphrase from section 2.1, is that we would like to minimize some measure of expected financing costs over some time interval while simultaneously keeping the risk, or variance, of these costs under control and maintaining a certain level of market liquidity. To formulate this in a mathematical sense, we need to develop some notation for the various quantities that we discussed in section 2.1. For technical reasons, we define a probability space, $(\Omega, \mathcal{F}, \mathbb{P})$, a filtration $\{\mathcal{F}_t, t \geq 0\}$, and a time interval, $t \in [0, T]$. The notation takes the form of the following definitions:

$$\begin{aligned}
\Xi &\stackrel{\triangle}{=} \text{ the stock of government debt at time 0,} \\
\theta \equiv \theta(t) &\stackrel{\triangle}{=} \text{ financing strategy at time } t, \\
\Theta &\stackrel{\triangle}{=} \{\text{the set of financing strategies that maintain market liquidity}\}, \\
\mathrm{S} \equiv S(t) &\stackrel{\triangle}{=} \text{ the state of the economy,} \\
\mathrm{r} \equiv r(t, y_1(t), ..., y_n(t)) &\stackrel{\triangle}{=} \text{ instantaneous interest rate,} \\
\{y_1(t), ..., y_n(t)\} &\stackrel{\triangle}{=} \text{ the sources of uncertainty driving the dynamics of } r, \\
\mathrm{P} \equiv P(t, T, S, r) &\stackrel{\triangle}{=} \text{ bond price function,} \\
\mathrm{F} \equiv F(t, S, P) &\stackrel{\triangle}{=} \text{ the government's financial requirements (i.e., surplus/deficit position),} \\
\mathrm{f} \equiv f(t, F, \theta, \Xi) &\stackrel{\triangle}{=} \text{ the government's borrowing requirement at a given instant in time,} \\
\mathrm{c} \equiv c(t, f, \Xi, P) &\stackrel{\triangle}{=} \text{ the cost of servicing the debt.}
\end{aligned}$$

While somewhat stylized, this set of definitions does capture the main points that must be considered in the analysis. Let us put all these symbols into words. Specifically, we are assuming that the government's borrowing requirements (F) are a random function that depends upon time and the short rate of interest (r). The bond price function (P) is critical; it tells us how much it will cost to borrow money in each debt instrument at each instant in time. The key assumption is that bond prices are a function of the instantaneous short rate, which is itself a random function—depending on n state variables $y_1(t), ..., y_n(t)$ and the state of the macroeconomy (S)—sufficient to describe the entire term structure of interest rates. We can see that the amount of financing at a given point in time (f) will depend upon the initial stock of debt (Ξ), the financial requirements of the government up to that point in time (F), the state of the economy (S), and the financing strategy (θ). Finally, the debt servicing costs (c) will depend upon the initial stock of debt, the financing that has occurred up to that point in time and, of course, the bond price function. In a general dynamic programming setting, what one would like to do is find the financing strategy or, more formally, the optimal control process from the set of permissible processes (Θ) that will minimize the cost of servicing the debt. Formally, we can represent this problem in a manner that looks very similar to equation (57) in Appendix A. Specifically,

$$\min_{\theta(t)} \mathbb{E}\left[\int_0^T c(t, f(\theta(t)), \Xi, P)\, dt\right], \tag{4}$$

subject to,

$$\mathrm{var}\left[\int_0^T c(t, f(\theta(t)), \Xi, P)\, dt\right] \leq \delta, \tag{5}$$

with $\delta > 0$, which is a *risk* contraint, and

$$\theta(t) \in \Theta, \tag{6}$$

which is, of course, a *liquidity* constraint. For a more detailed, though still largely heuristic, discussion of equation (4), see Appendix A. Stated another way, we are seeking a financing strategy (θ) that minimizes the government's expected debt charges (c) over some time interval ($[0, T]$), subject to restrictions on the amount of variance in these expected debt charges and maintenance of some minimum level of liquidity in the government securities market. This is a problem in stochastic optimal control. That is, given a system (our public debt) governed by random forces (interest rates, government financial requirements, and the macroeconomy), we seek to find a control (financing strategy) that optimizes our variable of interest (debt charges) over some time interval. Given the close correspondence to our qualitative definition of our debt strategy problem, this represents a clean and intuitive way to conceptualize it.

The news, however, is not all good. Despite the fact that we can define the debt strategy question in this straightforward manner, we cannot directly solve this problem in an analytical fashion. It does provide us with a framework for thinking about the problem. In particular, it helps us understand where (and why) we need to alter our analysis, from the theoretical case, to get at a potential solution. Although we cannot directly solve our problem, we may devise an approach for determining an approximate solution. Let us, therefore, consider the specific reasons why we cannot solve this problem in its theoretical form. There are at least three reasons. The first is that the definition of the problem assumes, by construction, that the surplus, financing, and debt-charge processes are time-continuous. Often, we use continuous functions or processes to approximate their discrete equivalents because of mathematical convenience. In this case, however, we will be forced to discretize the time dimension.

The second reason is that our function, c, is not known, because it depends on a number of stochastic processes, such as government surpluses and interest rates, that are themselves unknown. Worse, from a mathematical perspective, there are a number of feedbacks between these variables. In addition, even if the aforementioned processes were known and deterministic, the actual structure of the government's existing debt stock is enormously complicated. We can, however, compute c in a numerical fashion. That is, given an existing debt stock, a term-structure model, government financial requirements, and a financing strategy, we can determine the required financing and the associated debt servicing costs. This can be determined for any single realization of our random processes. In fact, we need to perform this computation for a very large number of realizations to calculate the expected cost and variance for the financing strategy in question. Our ability to perform this numerical calculation will form the basis of our approach to this problem. As we will see later, this presents a number of serious practical complications that need to be resolved.

The third reason is that our ability to numerically compute the expected cost and variance of a very large number of financing strategies does not necessarily bring us any closer to finding an *optimal* financing strategy. Indeed, optimality would imply that, from the entire set of permissible financing strategies, we have found the best financing strategy in terms of expected cost, variance, and time diversification. Mathematically, we cannot accomplish this task without some kind of optimization algorithm overlaid onto our simulation framework.[17] That is, we require clearly defined and well-behaved functions, along with detailed knowledge of the nature of the exact statistical nature of the randomness in the model, to find an optimal solution. None of this, unfortunately, is available in the present set-up, for a reason that is not entirely mathematical. While we know that the government is risk-averse, it is fundamentally difficult to quantify this risk aversion in terms of both variance of expected cost and time diversification. The result is that, instead of finding an optimal solution, we must seek to understand the nature of the potential trade-offs that can be made between expected cost, variance, and time diversification, and thereby better inform policy-makers about the nature of the debt strategy decision.

[17]This is clearly an interesting direction for analysis that we would like to explore in the future. In particular, the stochastic gradient methods outlined in Ermoliev and Wets (1988) and Benveniste, Métivier, and Priouret (1997) could prove useful in this respect. There is also a literature on stochastic programming that may very well be applicable to this problem.

3 The Model

We have discussed the issues of debt strategy analysis at length and have seen that the problem can be defined in a mathematically rigorous fashion. More importantly, we have demonstrated that a reasonable approach to characterizing the solution to the debt strategy problem involves stochastic simulation. Nevertheless, there are still a large number of things to consider in the construction of a simulation model, and thus some additional form of classification is necessary. We have correspondingly divided our simulation framework into three separate analytic components: the stochastic model, the control model, and distributional analysis. More specifically, the stochastic and control models feed into the *debt strategy engine*, which provides its outputs to the area of distributional analysis. This classification is outlined in the schematic in Figure 2.

Before delving into the specifics of each individual component in Figure 2, it is worth providing an overview of what we need to accomplish in our simulation model. In its most basic form, we are faced with the following sequence of tasks:

(i) First, we select a *time horizon* for our analysis (i.e., $[0, T]$). We have, for the purposes of this paper, elected to use a 10-year time horizon with quarterly time increments.

(ii) We must also determine the composition of the initial portfolio (i.e., Ξ). This is somewhat more involved than it might at first appear.

(iii) We need to select a *financing strategy* (i.e., θ). In our analysis, we treat this as a vector of weights in each debt instrument. It turns out this is not as immediately obvious as it sounds and we will discuss this issue in detail later.

(iv) We must generate a random sequence of future macroeconomic states, government financial requirements, and term structure of interest rates (i.e., $S(t), F(t), P(t,s)$ for $t \in [0,T]$ and $s \in \left[\frac{1}{12}, 30\right]$). In short, we need to create our *stochastic model.* This requires substantial computational expense.

(v) We then apply the financing strategy for each period over our time horizon. This requires that we take into account maturing debt stocks and new financial requirements. We also must ensure that we respect the government's reopening cycle and keep the portfolio composition in balance with the desired financing strategy.[18] The feedback between financing strategy and financial requirements, as well as the relationship between the financing strategy and issuance

[18] In Canada, as in a number of other countries, a bond may be issued in small increments over a period of time. A given 10-year Government of Canada coupon bond, for example, is typically auctioned to government securities dealers on four occasions over the course of a year, before a new maturity is introduced.

rates, must also be respected. This step describes the work of our *control model*. Again, this step is computationally expensive.

(vi) For each period, we compute the associated *debt charges* to the government associated with the selected financing strategy (i.e., $c(\theta, \cdots)$).

(vii) The preceding steps must be performed literally thousands of times to generate a distribution of government debt charges. This requires that we be cautious in the construction of our computer programs.[19]

(viii) Using this debt-charge distribution, we consider various *measures of cost and risk* to assess the relative desirability of a given financing strategy. We call this final step *distributional analysis*.

This sequence of steps falls into one of the four categories described in Figure 2. The remainder of this section will address the various issues involved in each of these categories and highlight our corresponding modelling choices.

3.1 Stochastic model

The first area of discussion is the stochastic model. This is perhaps the most important element of the analysis, because an ill-specified or incomplete stochastic environment will imply incomplete or flawed results. Bolder (2001, 2002) directly addresses the stochastic model that we will be using in our simulation framework for debt strategy. We stress that these works do *not* represent the final word on this area, but rather a reasonable first step towards approximating the complicated random element of this analysis. Moreover, we are *not* attempting to predict future outcomes. Instead, we characterize the range of possible future outcomes and use this information to approximate the corresponding distribution of future debt-charge outcomes.

This first step is to specify the dynamics of the term structure of interest rates. As described in Bolder (2001), we have identified three desirable characteristics in a term-structure model for use in Bank of Canada applications: an adequate temporal description of the dynamics of the Canadian term structure, the existence of an analytic representation for the relationship between the underlying state variables for speed of computation in our simulation setting, and a parameter set that is relatively easy to estimate and interpret.

These considerations point us towards a rich set of term-structure models, popularized by Duffie and Kan (1996), termed the class of *affine* term-structure models. This class encompasses the models of Vasicek

[19] It would appear that this would be a natural application for the idea of variance reduction. We did, in fact, write our code to permit the use of antithetic variables and control variates. Our concern was that, by dampening the variance of our estimates, we would actually be dampening the volatility of the debt-charge sample paths. The result would be to underestimate the risk associated with a given financing strategy. Appendix C provides a brief review of some of the principal variance-reduction techniques.

Figure 2: **A Model Schematic**: This figure illustrates the various components of our simulation framework for debt strategy.

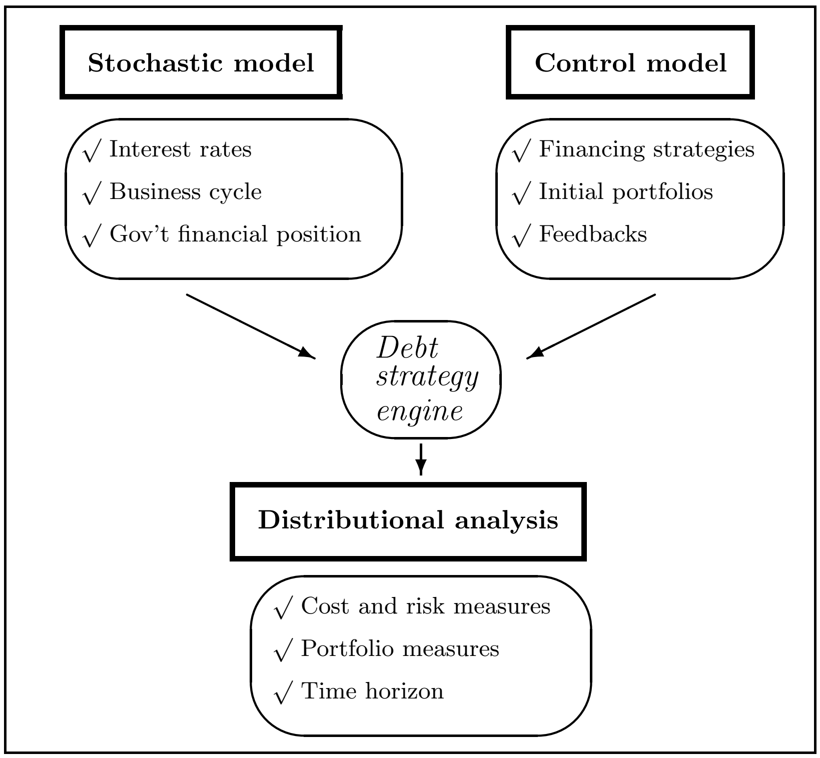

(1977), Cox, Ingersoll, and Ross (1985a,b), Longstaff and Schwartz (1992a,b), and a number of others. These models are formulated by assuming that future dynamics of the term structure of interest rates depend, in its simplest form, on the evolution of a single observed, or unobserved, factor. This factor, also termed a state variable, is a random process that is restricted by the assumption of an absence of arbitrage in the underlying financial market. The no-arbitrage restriction permits the derivation of a deterministic relationship between the term structure of interest rates and this state variable. Two special cases of this model, the Cox, Ingersoll, and Ross (CIR) and the Vasicek model, can readily be extended to incorporate multiple-state variables and permit analytic solutions for the bond price function; the two-factor CIR model is our selected model for term-structure evolution. Multiple-state variables are important, because substantial evidence suggests that the use of a single-state variable to explain the random future movement of the term structure is inadequate. This inadequacy stems from the fact that the dynamics of the term structure of interest rates are too complex to be summarized by a single source of uncertainty.

As noted, the term structure of interest rates is only one component of the stochastic environment.

The government's financial requirements also vary over time in a random fashion. Bolder (2002) constructs a reduced-form model that describes the joint evolution of the economic business cycle, the government's financial position, and the term structure of interest rates. To accomplish this goal, it models the dynamics of the business cycle with the hidden-Markov model suggested by Hamilton (1989). Thus, the stochastic framework is built on a conceptually straightforward and flexible foundation. Bolder (2002) then employs a transformation of the filtered probability of recession to capture the flat or inverted term-structure outcomes observed to occur prior to business cycle downturns. It captures these dynamics—in an admittedly simplistic manner—by constructing a time-varying market price of risk parameter through a convex combination involving the filtered probabilities. The government's financial position is specified as a modified Ornstein-Uhlenbeck process. The process is modified in the sense that the dynamics of the government's financial position depend importantly on the current state of the business cycle.

The previous paragraphs briefly discussed our stochastic model. The interested reader is referred to Bolder (2001, 2002) for an in-depth presentation of these models of randomness. A critical issue in any stochastic model is their parameterization. Not only do we use statistical estimation techniques to fit our models to historical data, but we also perform a number of diagnostic tests to ensure that the behaviour of our theoretical models is consistent with this data. These diagnostic tests allow us to accomplish three objectives. First, they provide assurance that our theoretical models are doing what we think they are doing. This is always critical when using complicated models. Second, because an estimated parameter set may not perform exactly as we might wish, diagnostics permit us to calibrate the parameters to obtain the desired behaviour of our models. Third, they help us to understand the key elements of our theoretical models and often provide useful suggestions for sensitivity analysis. Appendix B provides a more detailed description of the model estimation, calibration, and diagnostics.

3.2 Control model

The control model is exactly as the name suggests. It represents those elements of the debt strategy problem over which the government exercises a level of *control*. In particular, the government controls the choice of financing strategy. That is, the composition of debt issuance in any one period is determined by sovereign debt managers. As we learned in section 2, the stochastic environment is not entirely independent of the financing strategy. As such, our goal in this section is to describe the elements of debt issuance and the appropriate relationships with the stochastic environment in a defensible manner. Following the schematic in Figure 2, we address the key issues in the control model over the next three subsections.

3.2.1 Time horizon and initial portfolio

The first issue is the determination of an appropriate time horizon. Because the government has the option to issue very long-term debt instruments (i.e., 30 years), it would seem reasonable to argue that the minimum time horizon for our analysis would be the maturity of the longest-term debt instrument in the portfolio. On the other hand, a long-term horizon is both computationally expensive and difficult to model in a reasonable manner. How, for example, does one reliably model the evolution of key macroeconomic variables over a 30- to 40-year time horizon? There is also the confounding influence of the initial portfolio. If we analyze a financing strategy that is significantly different from the composition of the initial portfolio, then we are in essence examining the *transition* of the portfolio from its current composition to a new structure. The risk-cost characteristics of this financing strategy will, in fact, be the risk-cost characteristics of the *transition* from one government debt portfolio to another. This may be useful in some circumstances, but it is not helpful for an overall comparison of different financing strategies. Ultimately, we wish to compare financing strategies in equilibrium or in their steady state. What does this mean? A portfolio is in steady state if the proportion of debt instruments in the overall portfolio is identical to the portfolio weights in the financing strategy vector (i.e., Ξ is equivalent to θ).

The consideration of equilibrium portfolios is consistent with the work of Bergström and Holmlund (2000), and we believe it is crucially important that we use so-called *steady-state* initial portfolios for our analysis. To a large extent, this eliminates the problem of determining the time horizon in the analysis. In this way, our measures of cost and risk capture the dynamics of the portfolio in equilibrium and *not* a potentially noisy (and incomplete) transition to a new equilibrium level. Of course, the important question of optimal transition to the new state is not examined in this framework. Thus, our focus in this work is implicitly on the long-term strategic direction of the management of government debt. The current portfolio remains critically important for determining the short-term tactical decisions. Clearly, these are related issues, but we believe the strategic component of this decision best lends itself to quantitative analysis. Tactical analysis, by constrast, tends to be highly related to the institutional details of the government debt market. Long-term strategic analysis compares the risk-cost characteristics of various financing strategies. Once these are determined, the government can select the specific financing strategy that is most desirable, given their objectives. If this differs substantially from the current portfolio, then a transition strategy must be selected and short-term tactical analysis comes into play. This may not always be necessary. In particular, there need not be a one-to-one correspondence between a government financing strategy and its cost-risk exposure. Numerous European sovereigns use interest-rate swap contracts to partially decouple their issuance pattern and their risk exposure. That is, they might issue in several debt instruments and then use interest-rate swaps to create a portfolio with only one or two maturities. This theoretically appealing option does raise a number of thorny practical questions. Nevertheless, it illustrates a potential distinction between transition and equilibrium financing strategies.

Another way to consider this problem is from a control-theory perspective. In some sense, examination of equilibrium strategies is a form of characterizing the solution space of the control problem. It identifies those financing strategies in the overall set of permissible strategies that are desirable. One then proceeds to consider the current portfolio and works out how to get to this desired state. This is, roughly speaking, similar in spirit to the *backwards induction* argument used in the dynamic programming techniques employed to solve optimal-control problems.

3.2.2 The financing strategy

In a general optimal-control setting, the control—or, in our problem, what we have defined as the financing strategy—is itself a function of time. It is *not*, in fact, constant from one period to the next. In a stochastic optimal-control setting, therefore, this generalizes such that this function of time can depend on the stochastic environment. Thus, in full generality, one can consider the financing strategy to be itself a stochastic process, which means that the type of financing strategies the government could undertake are virtually unlimited. The government could entertain an issuance strategy, for example, that depends importantly on the evolution of the term structure of interest rates.

This could deeply complicate our analysis. At this point, therefore, it is useful to discuss the borrowing objectives of the Government of Canada to provide some guidance in our analysis. In a variety of public documents—Department of Finance Canada (2001) is a good example—the federal government clearly describes among its operational principles of debt strategy the ideas of transparency, liquidity, and regularity. Given the size of the federal government's debt and its importance as a benchmark for the Canadian fixed-income market, the Government of Canada does *not* engage in opportunistic borrowing. As a consequence, financing strategies that attempt to pick ideal market conditions for the issuance of various debt instruments are not considered. Instead, we need to consider financing strategies that reflect the government's long-term, strategic approach to financing its debt porfolio. We argue that a reasonable way to represent this approach is to consider the proportion of its overall portfolio the government wishes to hold in each of the available debt instruments. Does the government, for example, wish to have a debt portfolio with a large proportion of issuance in treasury bills and a small amount in two-, five-, 10-, and 30-year bonds? And, if so, what is the relative proportion of the debt portfolio in each of these bonds? Are they equally weighted or skewed towards the long end? It seems that a reasonable way to perform this analysis is to consider a large number of possible portfolio weights. In this analysis, therefore, we will define our financing strategies as fixed vectors of portfolio weights.

This is a useful idea, but how do we put it into operation? It turns out that applying a fixed vector of weights to total issuance in each period is not a reasonable approach; that is, a flow rule is difficult to put into place. Indeed, a flow-based financing strategy will not maintain the portfolio in a steady state. To illustrate this point, consider an extreme example. Assume that we have an initial one-time annual financial requirement of $100 billion. Moreover, we decide to follow a financing strategy half of which consists of 30-year bonds, with the other half consisting of 3-month treasury bills. In the first year, therefore, we issue $50 billion of each instrument. Observe that it will require some time before the first 30-year bond arrives at maturity. In the second quarter of the first year, therefore, there will be no 30-year bond maturity. This is contrasted with $50 billion of 3-month treasury bills that will mature in the second quarter. If we were to apply our flow-based issuance rule to these maturities, we would have $25 billion of new 30-year and 3-month treasury-bill issuance occurring in the second quarter. In the third quarter, because of the lack of 30-year bond maturities, we would be required to issue $12.5 billion of 30-year bonds and 3-month bills. Repeated application of this financing rule would, in short order, lead to a portfolio completely dominated by longer-term borrowing, in stark contrast to the initial intention. This is less extreme when starting with an existing portfolio that involves continuously maturity long-term debt, but the effect is quite similar. Ultimately, this approach fails, because it focuses exclusively on *flow* measures of the portfolio. The bottom line is that a flow-based approach to financing strategy construction is *not* workable.

What are the alternatives? We suggest a fairly reasonable solution. It requires consideration of both *flow* and *stock* measures of the portfolio and requires two steps.[20] First, we separate the quarterly financial requirement into refinancing of previously issued debt and new borrowing related to the government's budgetary position. Second, we proceed to reissue any maturing amounts at their original term to maturity. For example, if this maturing bond was initially issued as a five-year bond, then this will lead to a new five-year issue in that quarter equal to the notional value of that bond.[21] Any new borrowing will use the fixed vector of weights that represent our financing strategy.

As long as the government is always in a deficit or has a zero financial requirement, this will maintain the portfolio in its initial steady state. Problems in the model arise if the government moves into a surplus position, because our previously described financing strategy does not permit negative issuance. There are two potential solutions to this problem. First, we could allow any surplus position to be applied to the government's treasury-bill stock. This approach will not maintain the government's steady-state portfolio. Second, we could simply apply the fixed vector of weights in a negative fashion (i.e., assume that these are buybacks), to bring the portfolio equally back into line with the reduced financial requirements. This will maintain the steady-state portfolio over all periods. We adopt the second approach in this paper.

The success of this approach would seem to depend on an even issuance pattern for all bonds across

[20] This discussion of financing strategy modelling is, by necessity, fairly specialized to the Canadian situation. Other sovereign borrowers face similar circumstances and thus we suspect that, with some modification, this algorithm could be applied more generally.

[21] This will require that our computer program keep track of the initial issuance maturity of all instruments.

periods. For example, in our actual portfolio we actively use bond reopenings to ensure that we obtain large, liquid benchmarks. This implies, for example, that we typically have two reopenings for each two-year government coupon bond. This situation is magnified with approximately four to six reopenings for the five-, 10-, and 30-year bonds. Thus, if we applied the previously described solution to the problem, we would be forced to issue a huge amount of bonds in a single quarter to maintain the steady-state portfolio. With some care we can deal with this problem, but we will have to tolerate some short-term deviations from the steady-state portfolio. The idea is that when a large benchmark matures, we do not immediately refinance the entire bond, but instead follow a gradual reissuance pattern.[22] For example, if \$12 billion of originally issued 10-year bonds matures in a given quarter, it is not realistic to assume that we refinance the entire issue in that quarter. Instead, we reopen this bond over the next n quarters, where $n \in \{3, 4, 5\}$. This is accomplished by issuing a new bond in each quarter, at par value, but reducing the term to maturity of the bond in each subsequent quarter. With a two-year bond, for example, where $n = 3$, we would start with a maturity of $2\frac{1}{4}$ years, then 2 years, followed by $1\frac{3}{4}$ years.

Thus, 10-year issuance in each quarter would amount to $\$\frac{12}{n}$ plus any adjustments required from the government's deficit or surplus position. This would also have some repercussions for how we construct our steady-state portfolio. Instead of having equal issuance in each previous quarter, we could apply a more complicated rule. In particular, for our two-year bond example, we would have to go back to $2\frac{1}{4}$ years and reconstruct the past quarterly issuance following a cycle of three reopenings per bond. As long as $n = 3$ remained constant over the entire period of analysis, we would remain relatively close to the steady-state portfolio. The benefit of the increased complexity involved in this approach is that we could continue to examine the rollover risk of the portfolio. We chose this stylized solution because it is both reasonable and workable. Were we, for example, to have a new maturity for each bond in each quarter, then our redemption profile would be unrealistically smooth. This would, of course, simplify the implementation, but it would undermine completely our ability to say anything about refinancing risk.[23]

There is one additional complexity. If we decide to spread the maturity of a given, typically fairly large, benchmark over a number of quarters, we still have to meet the initial maturity repayment. The question is, where do we raise these funds? Our solution involves the creation of a *cash account*. The role of this cash account is to serve as a residual value in the overall balance equation. In general, maturing bonds plus new financial requirements of the government less any issuance in the period must be equal to zero. Consider a quarter wherein we have \$9 billion of two-year bonds maturing and we plan to refinance this issue over $n = 3$ quarters. Moreover, assume that the government's financing requirement in that quarter is zero. The

[22] In reality, some portion of a large benchmark may be repurchased by the government over the bond's lifetime. In this analysis, however, we abstract from the Canadian government's bond repurchase program.

[23] As a practical matter, we permit two reopenings per year for the two-year bond and four reopenings for the five-, 10-, and 30-year bonds.

consequence is that we have $9 billion of maturities, no financial requirement, and only $3 billion of new issuance. The required value, therefore, of the cash account in that period would have to be $6 billion. This could be financed at the 3-month treasury-bill rate and would, therefore, need to be refinanced in the next period. If the reopening cycle is restricted to a single fiscal year, then the portfolio should remain on a yearly basis in its steady state, although there would, of course, be some intra-year deviations.

3.2.3 Financial requirement feedback

In section 2, we discussed the role of debt charges in the government's financial position. The base model in Bolder (2002) aggregates the entire government financial requirement. In other words, it does *not* indicate what proportion stems from debt-service charges and what relates to program spending. In general, the impact of debt charges is countercyclical. That is, given generally lower interest rates during dips in the business cycle, the government typically experiences some debt-charge relief during these periods. Conversely, although revenue is stronger and expenditures smaller during strong economic conditions, interest rates are typically higher, implying higher debt-service costs. The strength of these relationships will, of course, depend on the government's financing strategy and the size of the existing debt stock. In particular, the greater the proportion of short-term debt, the more sensitive the portfolio to interest rate changes, and hence the greater the countercyclical effect. Clearly, this is an important factor in our analysis. What we require is an extension to our model to capture this interplay between the financing strategy and the government's financial position.

Perhaps the simplest approach would be to model the government's financial requirements on an ex-interest-charge basis. The long-term average government financial requirements—along with an initial portfolio—would allow us to determine some constant assumed value for the debt-service component. Let us say, for example, that this value is $40 billion. That is, we would set the long-term average financial requirement in the stochastic differential equation that describes the dynamics of the financial position process to $40 billion.[24] Then, we would dump the accumulated debt charges from that period into the government's financial position. In this manner, we could explicitly account for the dependence of the government's financial position on the financing strategy. This is problematic for two reasons. First, how does one find this *constant* value? It seems reasonable to expect that this value varies with both time and choice of financing strategy and, thus, is not constant at all. Second, we are ignoring the path-dependency in our problem. That is, the government will react to both positive and negative surprises in financial requirement—including the debt-service component—and adjust its budgetary position as required. This may occur with a lag, but it will nonetheless occur.

Another, somewhat more complete, approach is to constuct a prediction model for the government's

[24] In this case, the financial requirement is positive, indicating a surplus position.

debt-service costs. How might this work? As discussed in Bolder (2002), we model the financial requirement process, F_t, as a modified Ornstein-Uhlenbeck process. Let us denote the actual realized debt charges from period t as g_t. We suggest that we incorporate this form of feedback by adding the debt charges from the previous period to the financial requirement and simultaneously subtract the average realized debt charges from a number of previous periods, as follows:

$$F_t^{\theta} = F_t + g_{t-1} - \bar{g}, \tag{7}$$

where,

$$\bar{g} = \frac{1}{N} \sum_{i=2}^{N+2} g_{t-i}, \tag{8}$$

and F_t^{θ} represents the *total* financial requirements associated with financing strategy, θ. The idea is that the government predicts the debt-charge component for the overall financial requirement by constructing a simple average of the quarterly debt charges over the previous N periods.[25] If the financing strategy is quite variable (for example, consisting entirely of short-term issuance), then one would expect this prediction to have large variance. That is, $\bar{g}$, would be a fairly poor predictor for g_{t-1}. This variability would then, in turn, impact the variability of the financial requirement process and, in the true manner of feedback, lead to more variable debt-charge distributions. We considered a number of different financing strategies with various parameterizations of the financial requirement process, and concluded that $\bar{g}$ is perhaps too accurate a predictor of g_{t-1}. Part of the reason is that the government cannot, structurally speaking, update its forecast on a quarterly basis. Instead, this occurs on an annual basis. We recommend, therefore, the following revision to equations (7) and (8),

$$F_t^{\theta} = F_t + g_{t-1} - \tilde{g}, \tag{9}$$

where,

$$\tilde{g} = \frac{1}{N} \sum_{i=2}^{N+2} g_{\tau-i}, \tag{10}$$

where τ is the final quarter of the previous fiscal year. In this way, the government updates the debt-charge component of its financial requirements on an annual basis.[26]

We did also try an alternative specification, $\hat{g}$, such that,

$$\hat{g} \sim \mathcal{N}\left(\frac{1}{N} \sum_{i=2}^{N+2} g_{t-i}, \eta_{\theta}^2 \right), \tag{11}$$

[25] We need to work with a lag because the current financial requirements are necessary for the determination of the current debt charges. This could be solved using an iterative approach, but computationally this is not a feasible solution, using our already highly computationally intensive simulation approach.

[26] In the illustrative results described in section 4, we set $N = 8$.

where η_θ^2 is the forecasting error associated with financing strategy, θ. Put simply, this formulation implies that the government predicts the future debt-charge component of the financial requirement as an average of past behaviour plus a Gaussian random-error component. While this has some conceptual appeal, we rejected this approach in the end. We feel that the simple average approach is sufficiently naive that we could not justify adding an incremental forecasting error.

3.2.4 Issuance feedback

A final important relationship is the link between issuance and market yield. It is not reasonable to assume that market rates are constant while the amount of issuance varies substantially across various financing strategies. For example, the government cannot issue its entire financial requirement in a given sector of the term structure and not see fairly significant upward impact on market yields. This is, however, symmetric. The government cannot reduce issuance to very small levels and expect to issue at a neutral market rate. There is, therefore, an interval of issuance, quite likely depending upon the sector of the term structure, where debt can be issued at market rates. This is a second form of feedback, between the term structure of interest rates and the financing strategy. Again, we will try to consider a fairly ad hoc approach to modelling this relationship. Our approach is to add a penalty, for the ith debt instrument, to the par interest rate denoted p_i, such that,

$$p_i(x) = \begin{cases} p_i(x) > 0 : x \in [0, a_i) \\ p_i(x) = 0 : x \in [a_i, b_i] \\ p_i(x) > 0 : x \in (b_i, \infty) \end{cases}, \tag{12}$$

for $i = 1, ..., H$ where H denotes the number of available debt instruments. Thus, the interval $[a_i, b_i]$ would define the interval for the ith debt instrument where issuance does not impact market rates. Outside this interval, the cost of debt issuance would tend to increase. While this seems intuitively quite appealing, it does raise a number of difficult questions. First, how does one proceed to determine the size of the relative interval $[a_i, b_i]$, and what is the form of $p_i(x)$? Mathematically, there are a number of good choices for $p_i(x)$, such that equation (12) is satisfied.[27] Economically, however, the increase in par rates outside the interval is difficult to determine. This raises the second difficult question. Significant arbitrary changes in the par rates imply significant arbitrary changes in the zero-coupon term structure of interest rates. There is, therefore, the potential for our ad hoc alterations to par rates to introduce arbitrage into our interest-rate system. This is a concern, because our term-structure model was carefully constructed to avoid this problem. We do believe that, while this is clearly an issue, if the changes in rates are sufficiently small and we perform diagnostic tests to ensure no arbitrage opportunities in our interest rate system, there should not be any serious problems.

[27] A piecewise polynomial is one obvious choice.

These serious issues aside, we propose the following structure to model the penalty function for issuance outside of the interval $[a_i, b_i]$,

$$p_i(x) = \begin{cases} \alpha_i(a_i - x)^2 : x \in [0, a_i) \\ 0 : x \in [a_i, b_i] \\ \beta_i(x - b_i)^2 : x \in [0, 3b_i] \\ \beta_i(3b_i - b_i)^2 : x \in (3b_i, \infty) \end{cases}, \tag{13}$$

where, $i = 1, ..., N$. Essentially, this is a piecewise polynomial whereby issuance costs below a_i, and above b_i increase parabolically. The cost of issuance exceeding b_i flattens off at three times the upper bound. The choices of the parameters $\{\alpha_i, \beta_i, i = 1, ..., N\}$ are made to provide reasonable values for the penalty function; we chose to impose higher penalties to short-term issuance, given the higher absolute levels of issuance in these sectors. Figure 3 outlines the actual parameter selections for some debt instruments. The selected upper and lower bounds for non-penalized issuance are described in Table 1; these bounds were selected based on benchmark targets for government debt issuance described by the Department of Finance Canada (2001). Consider an example. Given our selected parameters for equation (13), borrowing \$4 billion at the 10-year maturity would occur at the best possible cost: the par interest rate generated by our model. The government, however, would pay a five-basis-point penalty if it issued \$15 billion during a given quarter.

Table 1: **Quarterly Issuance Intervals**: This table outlines the upper and lower bounds for quarterly debt issuance by instrument. These values are used in the computation of the penalty function described in equation (13).

Debt instrument	**Lower bound (\$ billions)**	**Upper bound (\$ billions)**	**Maximum penalty (basis points)**
3-months	18.00	40.00	43
6-months	9.00	20.00	12
1-year	9.00	20.00	12
2-year	2.50	5.00	9
5-year	2.25	3.75	9
10-year	2.25	3.75	5
30-year	0.80	1.50	3

Careful modelling of the interplay between financing strategy and debt issue costs is, in fact, a form of issuance constraint. It captures the idea that market liquidity is a function of the amount issued in a given sector. Moreover, it avoids a dangerous ceteris paribus assumption that one can alter the government's financing strategy and *not* alter the cost of debt issuance. Analysis performed using this ceteris paribus assumption would tend to yield corner solutions. It might suggest, for example, that the government should fund itself entirely with 3-month treasury bills. Figure 3 illustrates the penalty functions, described in equation (13), for each of the debt instruments used by the Government of Canada for a somewhat arbitrary,

Figure 3: **Piecewise Polynomial Issuance Penalty Function**: Given our selected parameters, these graphs illustrate the penalty function, described in equation (13), for 3-, 6-, and 12-month treasury bills and for 2-, 5-, and 10-year Government of Canada coupon bonds.

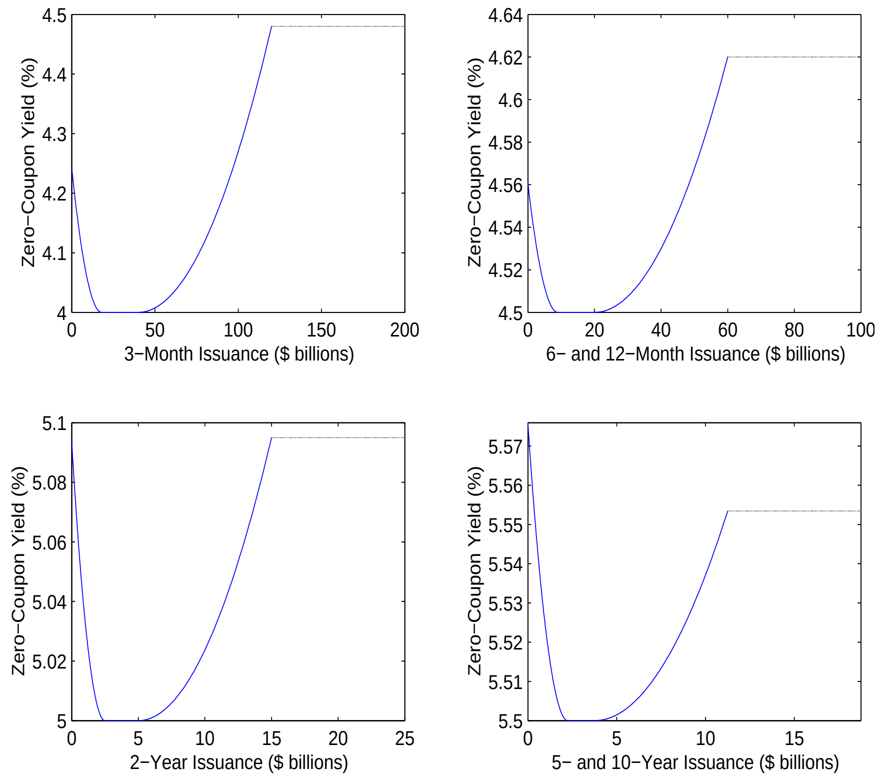

but hopefully reasonable, choice of parameters.[28] Table 1 outlines the upper and lower bounds associated with this parameterization. Despite the arbitrary nature of these parameters, we believe that it is important to include this feature in the model. Close inspection of the penalty functions shows quite clearly that these are small values. Because it is not obvious how to establish these quantities, it did not seem prudent to use overly large penalty values.

3.3 Distributional analysis

To this point, a substantial amount of work has gone into constructing a debt-cost distribution for each individual financing strategy. Constructing this distribution itself is not our ultimate objective. Instead, we wish to look at various aspects of this distribution to assess the relative desirability of each financing

[28]It was not obvious as to how we might estimate these parameters.

strategy. In it simplest form, we would merely look at the cumulative debt charges over the entire time-horizon of our analysis and select that financing strategy providing the lowest expected cumulative debt charges. Although this quantity is appealing from a conceptual perspective, it is difficult to interpret. What does it mean, for example, to distinguish between financing strategies based on expected cumulative debt charges for a 10-year time horizon? Moreover, this approach fails to capture the intemporal variation that is a key concern of a sovereign debt manager. Ultimately, the focus needs to be on an annual frequency to aid with the annual decision-making process. We will see in the following discussion that how this should be done is not immediately obvious.

Our simulation framework provides us with a tremendous amount of data. Clearly, this is a job for statistical analysis. Note that we have been using the term *distribution* quite loosely. In fact, the government's debt charges at a given time are described as,

$$c_t \equiv c(t, f, \Xi, P, S, \theta). \tag{14}$$

In words, this means that government debt charges depend upon a wide range of factors, including time, the government's financial requirements, the initial portfolio, the term structure of interest rates, and the state of the economy. In our simulation model, therefore, when we generate N sequences of values for $\left\{c_t, t = \frac{1}{4}, \frac{1}{2}, ..., 10\right\}$, we have actually constructed N sample paths for the stochastic process, c_t.

There are two primary ways we can think about the collection of c_t sample paths generated by our simulation model. First, we can freeze the time axis at any point and examine the distribution of debt charges. Figure 4 demonstrates this point graphically by examining the debt-charge distribution after five years. Given this structure, it is natural to consider various moments of this distribution as well as percentile-based tail measures. Notationally, we will denote this distribution as,

$$f_1(c_T), \tag{15}$$

for $T = \frac{1}{4}, \frac{1}{2}, ..., 10$. Because this is the standard approach employed in this type of problem, we will begin our discussion with a consideration of measures of distributions of the form f_1.

Some measures of these distributions are straightforward. Given a choice of T, the mean is perhaps the most obvious choice. As no assumptions of normality are made in the underlying interest-rate processes, it is also reasonable to consider a measure of central tendency not influenced by extreme observations such as the median. Correspondingly, one would also recommend that measures of disperson be examined, such as the standard deviation and the interquartile range.[29] These common measures are well-known and conceptually straightforward, and hence require no explanation. They do, however, form an important component of our analysis.

[29] The interquartile range is defined as the difference between the 75th and 25th percentiles.

Figure 4: **A Debt-Charge Distribution at a Given Point in Time**: In the first graph of this figure, we observe a collection of 200 debt-charge sample paths. In the second graph, we construct the distribution of debt-charge values at time T equal to five years.

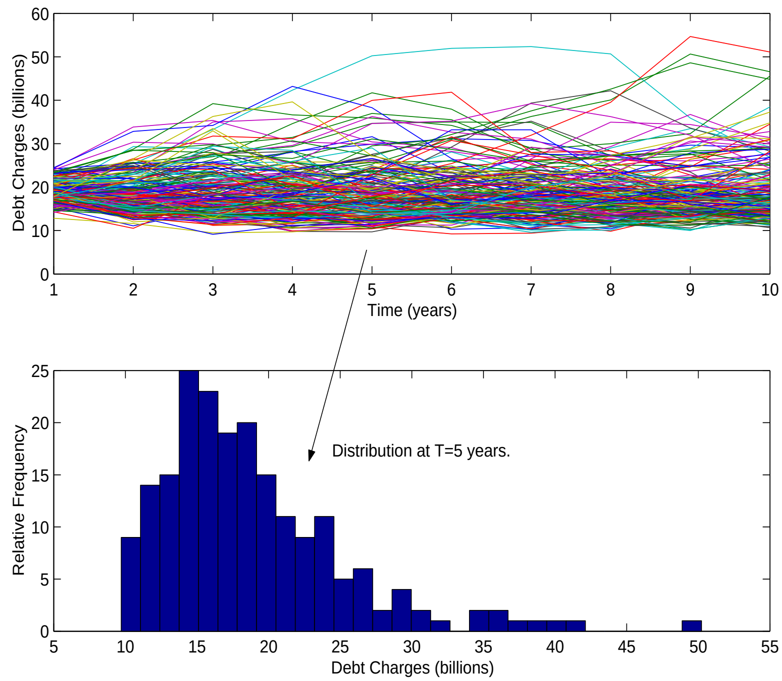

Often, in examining distributions, one also considers higher moments such as skewness and kurtosis as specific numerical quantities. In this setting, these are not terribly easy to interpret, because we are quite interested in tail events. That is, we ask: what is the worst outcome that we might normally observe based on the assumptions of our simulation model? Skewness and kurtosis measures might indicate that there is a greater chance of movements in one direction, or more probability mass in the tails of our distribution, but they do not provide a magnitude or an easy basis for comparison among various financing strategies. The most common measure used in this area was pioneered by Danish Nationalbank (1998), which borrowed a page from the well-known *value-at-risk* methodology. This measure, termed *cost-at-risk* (CaR), is defined as,

$$\mathrm{CaR}(x, p) = \sup\{z : \mathbb{P}(x \leq z) \leq 1 - p\}, \tag{16}$$

where x represents the random variable in question—in our case, government debt charges—and p is the critical percentile cut-off value. In words, it is the largest amount of government debt charges, over a given time horizon, that is not exceeded with probability $1 - p$. Under assumptions of normality, the CaR measure is merely a multiple of the standard deviation of x. In our more general setting, the CaR measure is a percentile measure of the distribution. If, as is the case in our analysis, we set $p = 0.95$, then the CaR measure provides us with the largest debt-charge realization, such that it is exceeded by 5 per cent of the debt-charge observations (i.e., the 95th percentile).

In our setting, it is often more interesting to consider the CaR measure relative to its mean (or median) value, because, in general, financing strategies that are dominated by longer-term borrowing are more expensive to the government. They nevertheless exhibit substantially less variability over time, because of lower variability in long-term rates and less refinancing risk. As the CaR measure is attempting to measure deviations from normal conditions, it makes sense to consider the distance between the CaR measure and the mean observation. This measure is termed *relative Cost-at-Risk* and is defined as,

$$\text{RCaR}(x,p) = \underbrace{\sup\{z : \mathbb{P}(x \leq z) \leq 1-p\}}_{\text{equation (16)}} - \mathbb{E}(x). \tag{17}$$

In short, the relative CaR measure provides us with a sense of a kind of worst-case deviation from the mean.

Relative CaR is nevertheless somewhat arbitrary, given that we select the critical cut-off value, p. Moreover, it does not provide any real sense of how bad things could get once we move further out into the tail of the debt-charge distribution. An alternative measure of the tail of the distribution is termed the *conditional tail CaR*. It is defined as,

$$\text{TCaR}(x,p) = \mathbb{E}(x : x \geq \underbrace{\text{CaR}(x,p)}_{\text{equation (16)}}). \tag{18}$$

The conditional tail CaR is the expected debt charges for a given period conditional on being in the tail of the distribution. Furthermore, the tail of the distribution is defined as being all observations that occur beyond the CaR. If we are examining 10,000 different debt-charge outcomes at $T = 5$ years, then the conditional tail CaR is the average of the largest 500 outcomes.[30] In simple terms, the conditional tail CaR is telling us that if things go badly and the government finds itself in the tail of the distribution, then, under these circumstances, it provides the expected debt charges. We can, of course, also define the relative conditional tail CaR in a manner identical to equation (17),

$$\text{RTCaR}(x,p) = \underbrace{\mathbb{E}(x : x \geq \text{CaR}(x,p))}_{\text{equation (18)}} - \mathbb{E}(x). \tag{19}$$

Overall, we believe the conditional tail CaR measure is more robust to non-Gaussian distributions with large positive skewness. Given that the state variables in our term-structure model have non-central χ^2 transition densities, this is precisely the situation in which we find ourselves.

[30] Similarly, to compute the CaR measure one simply orders the outcomes from smallest to largest and then reads off the 9,500th element.

As a side note, there is an entire literature devoted to the study of *coherent* risk measures. This involves specifying a number of properties that would be desirable in a measure of risk.[31] The seminal paper in this area is by Artzner et al. (1999). Without getting embroiled in the mathematical details of this work and getting too far afield, this work provides theoretical reasons that support the use of conditional tail CaR in our analysis. Specifically, they demonstrate that the VaR measure is not coherent, because it fails to satisfy the subadditivity property. Thus, if x_1 and x_2 are two random variables, then,

$$\text{CaR}(x_1 + x_2, p) \nleq \text{CaR}(x_1, p) + \text{CaR}(x_2, p). \tag{20}$$

The idea is that the combination of two entities, or portfolios, should not create incremental risk. The conditional tail CaR measure, however, does satisfy this property. As such, in this analysis, we will consider both the relative CaR and relative conditional tail CaR measures.

A challenge associated with these annual-based measures of cost and risk pertains to the time dimension. In a one-period problem, the interpretation of cost and risk is relatively straightforward. For example, consider a situation where we generate a debt-charge distribution for a one-year time horizon. To summarize the cost of this portfolio, we might look at the expectation of this distribution or perhaps some median value. Risk could be summarized by variance, relative CaR, or relative conditional tail CaR. If we extend our time horizon to five or 10 years, how do we alter our measures of risk and cost? It is not obvious as to how one should interpret the annual expected debt charges for a period 10 years in the future. How, therefore, do we *integrate out* the time dimension for our analysis?

A natural approach would be to consider some set of average risk and cost measures over the interval. While this is appealing in many respects, it could potentially hide trends in the data over time. Swedish debt managers, in particular, have given this issue substantial consideration. Holmlund and Lindberg (2002) average their measure of debt costs across sample paths and examine this distribution of averages. They then proceed to disaggregate risk into two dimensions: scenario risk and time-series risk. Scenario risk is defined as the dispersion of the distribution over time and is calculated as the relative distance between the 95th and 50th percentiles of the distribution of debt-charge averages. The second element, time-series risk, is defined as the deviation, or variation, around the trend in the individual sample paths. To compute this quantity, Holmlund and Lindberg fit a trend using OLS to each individual sample path. They then calculate the average absolute deviation from this trend and examine the resulting distribution of averages of absolute deviations from the linear trend.

This novel approach is a good first step at examining the problem. Our concern with this methodology relates to the fitting of the trend. The entire sample path is required to fit the trend, but in reality the entire sample path is not known to the government until the end of the period. This led us to consider how one

[31] Conceptually, this is similar to the definition of a metric in mathematical analysis and the concept of a measure in a measure-theoretic setting.

Figure 5: **A Conditional Debt-Charge Distribution Over Time**: In this figure, we observe 25 debt-charge sample paths. At $T = 5$, assume that c_5 is \$25 billion. The question then becomes, what is the distribution of possible values of c_6? The conditional distribution provides an answer to this question.

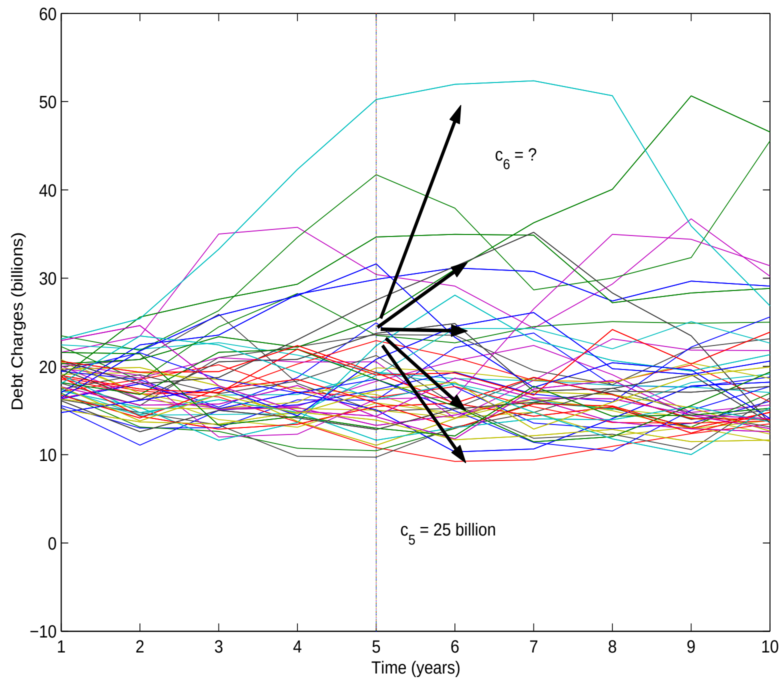

might better describe the evolution of the sample paths over time. One way of thinking about this collection of sample paths is to focus on the conditional debt-charge distribution.[32] Figure 5 illustrates a conditional distribution. Essentially, this is a pathwise measure. That is, if we are told the value c_t, we ask ourselves what the distribution of possible values is for c_{t+1}. Notationally, we denote this distribution as,

$$f_2(c_{t+1} \mid \sigma\{c_t\}), \tag{21}$$

for $t \in T = \left\{\frac{1}{4}, \frac{1}{2}, ..., 10\right\}$ and where $\sigma\{c_t\}$ is the σ-algebra generated by the debt-charge process, $\{c_t, t \in T\}$. A bit of reflection will reveal that this is an interesting distribution for a government debt manager. It provides a sense of the evolution of this process over time. The variance of the conditional distribution is of particular interest. It provides a measure of the uncertainty facing the government at each step in time as it attempts to forecast debt charges for the subsequent period. Ultimately, this is another alternative to capturing the time-series risk defined in Holmlund and Lindberg (2002).

How, then, might we characterize this conditional distribution? We suggest using a very simple approach

[32] This is occasionally termed the transition density.

by assuming that its distribution is Gaussian, and using a first-order autoregressive model to determine its first two moments. For example, given the debt-charge process c_t, we assume that,

$$c_t = \phi_0 + \phi_1 c_{t-1} + \epsilon_t, \tag{22}$$

where,

$$\epsilon_t \sim \mathcal{N}(0, \xi^2). \tag{23}$$

It is well known, and quite easy to show using a recursive argument, that the unconditional mean of c_t associated with this model is,

$$\mathbb{E}[c_t] = \frac{\phi_0}{1 - \phi_1}, \tag{24}$$

which precludes ϕ_1 from taking the value of unity.[33] By a similar argument, the unconditional variance is,

$$\text{var}[c_t] = \frac{\xi^2}{1 - \phi_1^2}. \tag{25}$$

These will be useful to determine whether our estimates are giving sensical results. The conditional distribution associated with equation (22) is given as

$$f_2(c_t \mid \sigma\{c_{t-1}\}) \sim \mathcal{N}(\phi_0 + \phi_1 c_{t-1}, \xi^2). \tag{26}$$

Estimating this model directly poses some practical issues. We have N sample paths, but in our model construction each individual sample path has only 10 annual debt charge outcomes. To overcome this issue, we estimate the model N times using OLS and report the average values for our parameters, ϕ_0, ϕ_1, and ξ.[34] The conditional variance ξ^2 is of the most interest. It is an admittedly imperfect measure of the average prediction error facing a government that attempts to forecast debt charges for the subsequent period in its budgetary process, because we are making the strong assumption of Gaussianity in deriving these estimates. Nevertheless, we believe this assumption is reasonable given its simplicity and ease of interpretability. In section 4, we explore the validity of this assumption by examing the empirical distribution of debt-charge first-differences.

The conditional volatility is a risk measure that captures the path dependency inherent in the debt strategy problem. One way to think about the AR(1) model is as a sequence of forecasts. One selects the value of ϕ_0 and ϕ_1 such that, at the beginning of each fiscal year, one produces the best linear forecast that one can construct (i.e., $\phi_0 + \phi_1 c_{t-1}$) based on the past year's debt charges (c_{t-1}). This forecast is subject to error, of course, and this is captured by the conditional volatility, ξ. This notion of *risk* is closely aligned with the manner in which a debt manager conceptualizes risk. Furthermore, it provides a long-term description of the annual, conditional uncertainty associated with a given financing strategy.

[33] Indeed, this would imply that c_t was a non-stationary distribution or, equivalently, that c_t was *white noise*.

[34] We found maximum likelihood to be relatively non-robust to the small number of data points and, more importantly, very slow, because it requires the use of a non-linear optimization routine.

4 Illustrative Results

In this section, we amalgamate the understanding of the problem developed in section 2 with the modelling details outlined in section 3 and demonstrate our simulation framework in action. We stress that these results represent a portfolio that is structurally similar to the Government of Canada's domestic debt portfolio, but this analysis neither represents an optimal policy for the government nor attempts to arrive at any conclusions about future debt strategy. Instead, our goal is to learn about the nature of the risk and cost trade-offs associated with different financing strategies. There are a series of items that we would like to answer:

(i) How does the introduction of a more complicated modelling environment influence the results of our analysis relative to a simpler approach focusing only on interest-rate variability?

(ii) How sensitive are the results to assumptions about the stochastic environment?

(iii) The results represent the risk-cost characteristics of a set of financing strategies under normal financial-market conditions. How do they perform under situations of market stress that involve extreme macroeconomic outcomes?

(iv) How sensitive are measures of cost and risk to changes in the financing strategy? That is, if we make a small change in the financing strategy, can we expect to observe large changes in cost and risk outcomes? Moreover, what is the nature of this relationship? Is it relatively simple or more complex?

(v) How do various financing strategies compare when we adjust for their risk? Are certain financing strategies on this risk-adjusted basis, for example, more desirable than others?

(vi) Understanding these four previous points, what is a reasonable set of cost and risk measures to use in analyzing debt strategy questions?

These six questions will be addressed in turn in the bulk of this section. We begin in section 4.1 by reviewing the key assumptions employed in this study. Section 4.2 examines the differences between measures of cost and risk for a simple model that includes only interest-rate variability and for an extended model of the stochastic environment. This examination is performed by considering five distinct financing strategies. Section 4.3 plays with the parameterization of our stochastic models to help us better understand the sensitivity of our results to the modelling assumptions. In section 4.4 we describe and demonstrate a novel approach to stress testing that we believe will prove quite useful. This stress-testing idea is based on an exploitation of the hidden Markov model for the evolution of the business cycle. In section 4.5, we consider 225 distinct financing strategies and examine the relevant trade-offs in terms of cost, risk, and time diversification. We also use regression analysis to help us describe the relationship between risks, costs, and the financing strategy. Also in section 4.5, we attempt to provide an answer to the last two questions listed above by extracting the lessons learned from the analysis of the base portfolio, different sensitivities, extreme scenarios, and a wide range of financing strategies. Our objective will be to outline a framework for analyzing the risk and cost trade-offs associated with different financing strategies and to develop a *risk-adjusted* measure of debt cost.

4.1 Key assumptions

Ideally, to answer the above-noted questions, we need to consider a wide range of different approaches to the problem. Ultimately, we would like to examine these different approaches for a large number of financing strategies, but the exposition of our results in this case would prove cumbersome. Our solution is to consider a smaller number of financing strategies that will, we hope, allow us to reach qualitatively similar conclusions without a ridiculous amount of computer effort or a bewildering myriad of results. Specifically, we consider the first three questions in our list by examining five quite substantially different financing strategies.

How did we select these five sample financing strategies? To best replicate the current Canadian debt strategy environment, we used seven possible debt instruments, including three-, six-, and 12-month treasury bills and two-, five-, 10-, and 30-year bonds.[35] Then, in an attempt to create very different financing strategies, we discretized the amount of treasury bills used in each strategy. That is, we selected *five* strategies with 0 per cent, 25 per cent, 50 per cent, 75 per cent, and 100 per cent of the portfolio weight in treasury bills; the residual weight is financed with nominal coupon bonds. Given this discretization, we split the remaining weight in each individual debt instrument equally. The actual weights for the five financing strategies used in the following analysis are described in Table 2.

Table 2: **Five Sample Financing Strategies**: To demonstrate the various features of our debt strategy model, we selected these five significantly different financing strategies. We consider the debt-charge distributions of each of these different strategies under various model assumptions.

Financing strategy	**3-month**	**6-month**	**1-year**	**2-year**	**5-year**	**10-year**	**30-year**
100% bills	1/3	1/3	1/3	0	0	0	0
75% bills	1/4	1/4	1/4	1/16	1/16	1/16	1/16
50% bills	1/6	1/6	1/6	1/8	1/8	1/8	1/8
25% bills	1/12	1/12	1/12	3/16	3/16	3/16	3/16
100% bonds	0	0	0	1/4	1/4	1/4	1/4

Having selected our financing strategies, let us now describe the methodology used to compute our results.

[35] For the purposes of this analysis, we have excluded Canada's index-linked debt, which is termed the Canadian *Real Return bond*. In recent years, Real Return bonds comprised approximately 6 per cent of the Canadian government's outstanding market-debt stock.

First, we perform our analysis in the context of an initial *steady-state* portfolio. Specifically, the portfolio at the starting point of each simulation is in equilibrium. Each financing strategy, as discussed in section 3.2.2, is constructed to maintain the portfolio in this steady state over the 10-year time horizon of our analysis. We can therefore look at summary measures for these portfolios, including the fixed-debt ratio, the average term to maturity (ATM), and the MacCauley duration, because, since the portfolios are in steady state, these do not change over across either time or stochastic scenarios.[36] These summary measures, for each of our five financing strategies, are described in Table 3. The ATM and MacCauley duration measures are, by construction, very similar. The key difference is that the ATM measure does not consider intermediate cash flows and thus is a generally larger measure than the MacCauley duration. The fixed-debt ratio represents the proportion of debt, at a given point in time, that need *not* be refinanced in the next year. This will not map directly into the proportion of treasury bills in the portfolio, because bonds with less than one year remaining to maturity must also be considered as floating-rate debt.

Table 3: **Summary Portfolio Measures**: As each of these portfolios is in its steady state, each will maintain its portfolio composition over the 10-year analysis horizon. This table summarizes three popular portfolio measures—the fixed-debt ratio, average term to maturity, and duration—for each of our five sample financing strategies.

Financing strategy	**Fixed-debt ratio**	**ATM**	**Duration**
100% bills	0.00	0.42	0.42
75% bills	0.20	1.75	1.25
50% bills	0.40	3.09	2.12
25% bills	0.59	4.43	3.14
100% bonds	0.79	5.77	4.37

For this analysis, we have selected an initial portfolio size of $400 billion. This is rebalanced—following the relevant financing strategy—on a quarterly basis over a 10-year time horizon. For each financing strategy, we need to decide the number of separate realizations of our stochastic environment in the construction of our debt cost distributions. This decision is important, because we are attempting to achieve a reasonable degree of accuracy in our solution. The results of a simulation are, of course, subject to approximation error. It turns out that we can characterize this error in our Monte Carlo approximations quite accurately by making use of two well-known statistical results. The first result, termed the *weak law of large numbers*, holds that, for a sequence of independent, identically distributed random variables, $X_1, X_2, ...$, where $\mu = \mathbb{E}X_i < \infty$,

$$\lim_{n \to \infty} \mathbb{P}\left[\frac{X_1 + \cdots + X_n}{n} - \mu > \epsilon\right] = 0, \tag{27}$$

[36] Clearly, the MacCauley duration is a function of the term structure of interest rates and thus will vary across scenarios. We have nevertheless elected to use an average set of zero-coupon rates to keep this measure constant over time.

for all $\epsilon > 0$. In words, this means that the average of independent trials will converge to its mean in probability—or rather, the Monte Carlo approximation error becomes arbitrarily small for very large values of n. Even better, if $\sigma^2 = \text{var}(X_i) < \infty$, we have that,

$$\lim_{n\to\infty} \mathbb{P}\left[\frac{\frac{X_1+\cdots+X_n}{n} - \mu}{\sigma\sqrt{n}} < \alpha\right] = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\alpha} e^{-\frac{z^2}{2}}\,dz, \tag{28}$$

for all $\alpha \in \mathbb{R}$. This is one version of the *central-limit theorem* and it gives us a prescription for determining the error in our Monte Carlo estimate for sufficiently large values of n. In particular, we construct a confidence interval for our Monte Carlo estimate as,

$$\frac{X_1+\cdots+X_n}{n} \pm \Phi^{-1}\left(1-\frac{\alpha}{2}\right)\frac{\hat{\sigma}_n}{\sqrt{n}}, \tag{29}$$

where, Φ^{-1} is the inverse of the standard normal cumulative distribution function and $\hat{\sigma}_n$ is the sample standard deviation.[37] Inspection of equation (29) reveals that the error of our simulation estimate decreases at the rate of $O(\sqrt{n})$.[38] Figure 6 illustrates the debt-charge values from the first year of an arbitrarily selected simulation. Observe their variation around the mean value of \$22.66 billion. Using equation (29), we can proceed to characterize the error in this estimate for a range of scenarios.

Table 4: **Simulation Error Estimates**: This table provides an example of the simulation error for various values of n in the first-year debt costs illustrated in Figure 6. We use a mean debt charge of \$22.6629 and sample standard deviation of \$1.4636 billion.

Scenarios (n)	**Lower bound**	**Upper bound**	**Confidence interval size**
100	22.3760	22.9497	0.5737
1,000	22.5722	22.7536	0.1814
2,500	22.6055	22.7203	0.1147
5,000	22.6223	22.7035	0.0811
10,000	22.6342	22.6916	0.0574
100,000	22.6538	22.6720	0.0181

The results of a small experiment with different selections of n are summarized in Table 4. Note that for 1,000 scenarios we observe a 95 per cent confidence interval of more than \$180 million, whereas increasing the number of scenarios to 10,000 reduces this to just under \$6 million. The incremental effort of performing

[37]Fishman (1995) indicates that, while this is an asymptotically valid confidence interval, there are, at least, three potential sources of error: non-uniform convergence of the sample statistic to the normal distribution, the error associated with using the sample standard deviation instead of the true variance, and the potential for positive correlation between the sample mean and the sample standard deviation. For the purposes of this study, we will assume that n is sufficiently large for these errors to be relatively minimal.

[38]$O(\sqrt{n})$ means that the speed at which the error declines is proportional to the speed at which $\frac{1}{\sqrt{n}}$ goes to zero.

100,000 simulations to reduce the size of the confidence interval to around $2 million does not seem worthwhile. Based on this analysis, therefore, and the computational capacity available to us, we have elected to consider 10,000 scenarios for each of the financing strategies analyzed in this section.[39]

Figure 6: **Simulated Debt Charges**: This figure outlines a set of 2,500 annual debt charges from the first year of an arbitrarily selected simulation.

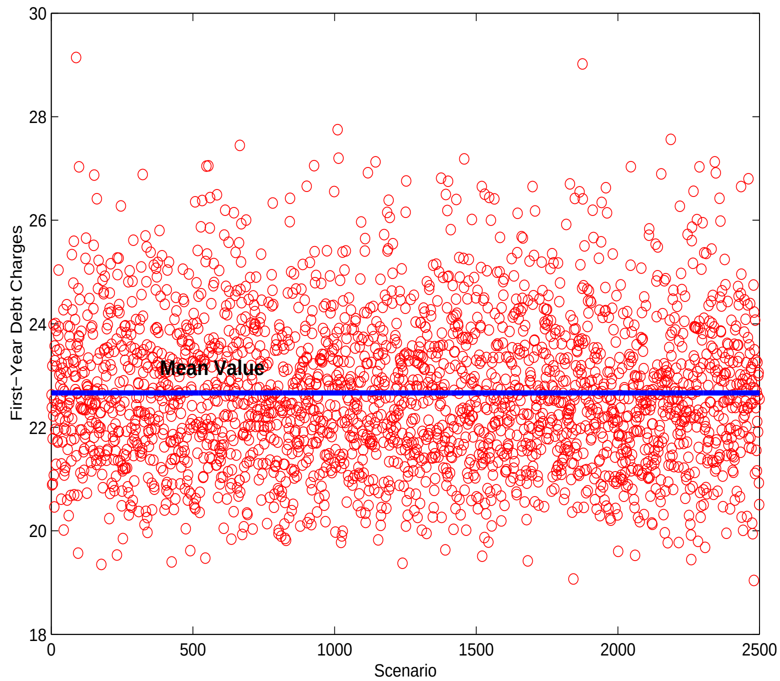

4.2 The basic model

In this subsection, we examine the cost-risk trade-offs associated with our five previously defined financing strategies in the context of two stochastic environments. The first setting is a *simple* stochastic environment. More specifically, the *only* source of uncertainty relates to the evolution of the term structure of interest rates. This means that the financial requirements of the government are assumed to be constant at zero, there is no business cycle, the government's issuance costs are independent of the financing strategy, and the government's financial requirements—given that they are constant—are naturally independent of the financing strategy. In the second stochastic environment, we use our full model. That is, we permit variation

[39] On average, the computation of 10,000 stochastic simulations for each financing strategy required between 45 minutes to one hour of computing time, depending on the options selected in the program. Our model was constructed using the mathematical software, MATLAB, and was run on a Sun Microsystems Blade Workstation running the Solaris 2.8 operating system.

of the government's financial requirements, incorporate a business cycle model, and allow for interplay between the financing strategy and both the government's financial requirements and the cost of issuing these debt instruments. The reason for this two-step approach is to construct a *base case*—involving only interest-rate variation—to compare the effects of the relaxation of these assumptions in the full stochastic model. On balance, owing to our careful model calibration, we should not expect drastic differences in the model results. Instead, we expect to observe a trend towards greater volatility in the general model, given its heightened realism and additional sources of uncertainty and interplay between key variables.

Our first step is to investigate the nature of our five financing strategies. It is instructive to look at the financing strategy weights as described in Table 2, but this does not provide a very clear sense of the amount of periodic debt issuance this implies for the government. In Table 5, we outline the average quarterly issuance associated with our five financing strategies. Clearly, the financing strategies dominated by treasury bills require significant quarterly issuance in these instruments. Also note that the 100 per cent bonds financing strategy actually requires a relatively modest amount of 3-month treasury bills each quarter, because to accommodate the reopening cycle of our four benchmark bonds, we need to use a cash account that takes the form of a 3-month treasury bill.

Table 5: **Average Quarterly Issue by Instrument**: This table summarizes the average quarterly issuance by debt instrument implied by our five different financing strategies. Observe that, for cash-management purposes, it is necessary to use the 3-month treasury bill in a limited manner for the 100 per cent bond financing strategy.

Financing strategy	3-month	6-month	1-year	2-year	5-year	10-year	30-year
Simple and full stochastic environments							
100% bills	133.69	66.85	33.43	0.00	0.00	0.00	0.00
75% bills	103.71	50.14	25.07	3.13	1.25	0.63	0.21
50% bills	73.73	33.42	16.71	6.27	2.51	1.26	0.42
25% bills	43.75	16.71	8.36	9.40	3.76	1.88	0.63
100% bonds	13.77	0.00	0.00	12.54	5.02	2.51	0.85

Note that the average quarterly issuance is virtually identical under the simple and full stochastic environments. This does not imply that the issuance is identical under our two specifications. The key difference stems from the fact that, because the simple stochastic model does not permit any variation in the government's financial requirements, the government's budgetary balance is assumed to be constant at zero. This implies that, in the simple setting, the average issuance values are constant across all scenarios. This is *not* the case in the full stochastic environment, which permits variability in government financial requirements. The average issuance levels are highly similar due to the calibration of the parameters in the full stochastic model. In expectation, this model will generate a zero level for the government's financial requirements. It will, of course, vary around zero for any given realization. The model parameters and calibration procedure are discussed in Appendix B.

Table 6: **Standard Deviation of Quarterly Issue by Instrument**: This table summarizes the standard deviation of quarterly issuance by debt instrument implied by our five different financing strategies. As the government's financing requirements vary over time, there will be a corresponding variation in quarterly issuance. This is reflected in these statistics.

Financing strategy	3-month	6-month	1-year	2-year	5-year	10-year	30-year
Full stochastic environment							
100% bills	2.35	1.20	0.63	0.00	0.00	0.00	0.00
75% bills	1.76	0.88	0.46	0.06	0.03	0.02	0.02
50% bills	1.18	0.58	0.30	0.12	0.06	0.04	0.04
25% bills	0.64	0.28	0.15	0.18	0.09	0.06	0.06
100% bonds	0.11	0.00	0.00	0.24	0.12	0.08	0.08

Table 6 summarizes the standard deviation of average quarterly issuance by debt instrument for the full stochastic environment. Financing strategies dominated by treasury-bill issuance quite clearly exhibit greater variability than financing strategies with larger proportions of bond issuance. This is intuitively reasonable given that more frequent refinancing exposes the government to additional variability in its issuance pattern. Another reason for the incremental variation in treasury-bill-dominated financing strategies is their sheer amount of average quarterly issuance. Higher absolute issuance should imply higher volatility.

There is also an additional element driving issuance volatility: the interplay between the financing strategy and the government's financial requirement. The debt charges associated with short-term borrowing are more volatile, thereby contributing, on average, to less accurate financial requirement forecasts for the government. This, in turn, adds to the overall volatility in average quarterly debt issuance. A separate examination of average quarterly debt issuance—not shown here due to space limitations—nevertheless reveals that the variability in the financial requirements is the more important effect for two reasons. First, the forecast rule we assume the government uses for projecting future debt charges, described in equation (50), is quite stable. Second, the assumed variability in government financial requirements is relatively small. We address this second point in section 4.3 when we consider the sensitivity of our model to two alternative parameterizations.

How do these differences in issuance impact measures of cost and risk? Table 7 describes the mean debt charges and their associated volatility for each of the five financing strategies across the 10,000 stochastically generated outcomes. These results are reported for the first, fifth, and tenth years of the 10-year simulation horizon. In both the simple and full stochastic environments, we observe a clear inverse trend between average cost and volatility. More specifically, as we increase the proportion of long-term debt in the financing strategy, we increase the expected cost to the government, but simultaneously reduce the volatility of these debt charges. In the first year, in the simple stochastic setting, for example, we observe expected debt charges of about \$18 billion for a financing strategy composed entirely of treasury bills, as compared with approximately $25 billion for a bonds-only financing strategy. The volatility of the treasury-bill-dominated financing strategy stands at almost $3 billion in the first year versus only $500 million for the 100 per cent bonds financing strategy.

Table 7: **Expected Debt Charges and Volatility by Financing Strategy**: This table outlines the expected debt charges and the associated volatility (standard deviation) of these debt charges for the first, fifth, and tenth year of the analysis horizon.

Financing strategy	Year 1		Year 5		Year 10	
	Mean	Volatility	Mean	Volatility	Mean	Volatility
Simple stochastic environment						
100% bills	18.33	2.74	17.96	5.90	17.96	6.76
75% bills	20.08	2.16	19.45	5.06	19.31	5.99
50% bills	21.82	1.59	20.94	4.23	20.66	5.28
25% bills	23.57	1.03	22.44	3.45	22.01	4.64
100% bonds	25.31	0.46	23.93	2.75	23.37	4.11
Full stochastic environment						
100% bills	19.28	2.79	18.93	6.22	19.02	7.26
75% bills	20.92	2.21	20.33	5.35	20.22	6.48
50% bills	22.26	1.63	21.45	4.48	21.22	5.71
25% bills	23.76	1.05	22.71	3.66	22.37	5.04
100% bonds	25.53	0.48	24.15	2.93	23.69	4.48

Inspection of Table 7 also reveals that the full stochastic setting generates *both* higher cost and volatility estimates across all periods relative to the simple stochastic environment. The expected cost, for example, in the full stochastic setting is higher by approximately $1 billion for 100 per cent treasury bill financing strategy, but decreases to only about $200 million for the 100 per cent bond financing strategy. The volatility of these expected costs increases fairly uniformly across all observations. This is fairly clear evidence of the incremental cost and risk associated with frequent refinancing of the government's portfolio. It also demonstrates that failure to model the random nature of the government's financial position will tend to bias downwards one's cost and risk estimates.

The next point is that, while the average debt charges remain stable across the 10-year horizon of our analysis—indeed, they even tend to fall slightly—their volatility rises dramatically across time.[40] Furthermore, although the shorter-term debt-based financing strategies continue to remain more volatile relative to the longer-term-debt strategies, the spread between them decreases with time. This is natural, because the uncertainty in the future realizations of our stochastic environment increases with time. It is clear, for example, that the range of possible term-structure outcomes is greater over a five-year period than a one-year period.

[40] The slight decline in debt charges is due to the slightly higher interest rates used to generate the cost of the initial portfolio.

Figure 7: **Annual Cost Distributions for 100 per cent Treasury Bills**: This figure is a sequence of annual histograms for the simple and full stochastic environments. The first graph relates to the simple setting and the second describes the full setting.

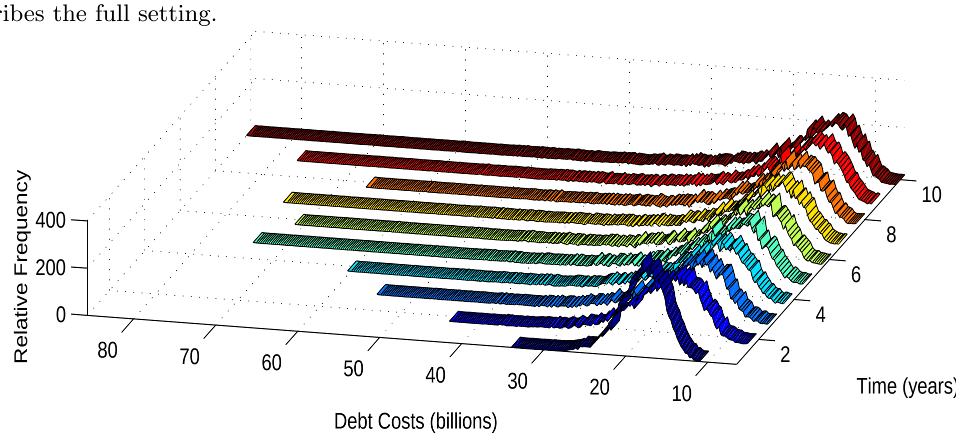

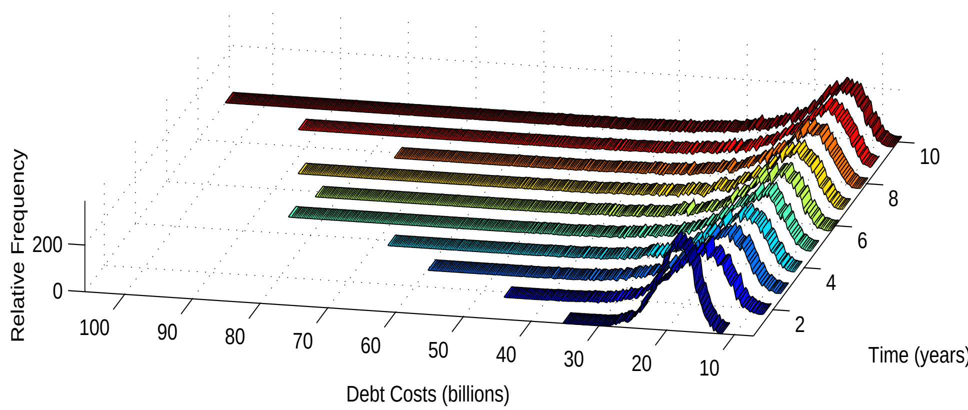

Although Table 7 summarizes the first two moments of our debt-cost distribution, it tells us relatively little about the tails of this distribution. Tail events are important because they provide information about the *worst-case* outcomes that we can expect to observe under our stochastic environment. Figures 7 and 8 graphically describe the annual debt cost distributions associated with the 100 per cent treasury bill and the 100 per cent bond strategies for both the simple and full stochastic environments; the first set of histograms, in each figure, relates to the simple setting, and the second set describes the full stochastic setting. Note in Figures 7 and 8 the positive skew in the annual debt-charge distributions. This stems from the form of the conditional distributions for the state variables used in our two-factor CIR interest-rate model. These state variables have non-central χ^2 distributions that preclude interest rates from becoming negative. This implies that there is a negative bound on debt charges. With relatively small probability—primarily because this behaviour is dampened by the mean-reversion inherent in these models—interest rates can take large positive values. This feature of our interest-rate model, which incidentally is a reasonable approximation of reality, contributes to the positive skew in our debt-charge distributions. Again, this positive skewness increases over time as the range of possible interest-rate realizations becomes more dispersed over time.

Also note in Figures 7 and 8 that, on balance, there is greater dispersion in debt-cost outcomes under the full stochastic environment. This relates to both the greater volatility in government financial requirements and the fact that the full stochastic model includes negative shocks to government financial requirements under recessionary conditions—this is described in section 3.1 and Bolder (2002). This generates a moderate negative tendency to financial requirements when the economy is in recession. On average, this generates somewhat larger debt charges and the deficit position acts to offset the positive covariance between the primary balance and debt charges, thereby generating greater financial requirement volatility and contributing to a lengthier tail for the annual debt-cost distribution.

A final observation about Figures 7 and 8 is the significantly lower volatility in the 100 per cent bond strategy relative to the 100 per cent treasury-bill strategy. Furthermore, the probability mass in Figure 8 is centred around \$25 billion, rather than the \$18 billion central value in Figure 7. Although these observations should come as no surprise, given our analysis of Table 7, it nonetheless lends some credibility to these figures. That is, an extreme short-term debt-based financing strategy leads to low expected cost but high volatility, while the longest-term based financing strategy (i.e., 100 per cent bond) generates high expected costs with low volatility.

While these graphs are helpful, we would also like to describe the tail of the annual debt-charge distributions in quantitative terms. In section 3.3, we defined two different measures that accomplish exactly this task. In particular, we defined the relative cost-at-risk (RCaR) in equation (17) and the relative conditional tail cost-at-risk in equation (19). Both of these quantities are outlined in Table 8 for the first, fifth, and tenth year of our analytical time horizon; the table also considers both the simple and full stochastic environments. To paraphrase the discussion of section 3.3, the best way to think about these measures is as the *worst-case* deviation from the mean that we would expect to observe under normal market conditions.

We would expect, based on our review of Figures 7 and 8, that these two measures for the full setting should dominate those for the simple setting. This is exactly the case. In the first year, the difference, ranging from \$40 to \$100 million, is not very extreme. Over time, however, the spread between the full and simple setting increases more substantially. This is a quantification of the greater volatility inherent in the full stochastic environment. If we focus only on interest-rate dynamics, we ignore other potential aspects that could contribute to extreme debt-charge outcomes.

Within each stochastic environment, we continue to observe a greater amount of *tail* risk associated with the treasury-bill-dominated financing strategies relative to the longer-term-based borrowing strategies. In the 100 per cent treasury-bill strategy under the simple stochastic environment, for example, we see that

Figure 8: **Annual Cost Distributions for 100 per cent Bonds**: This figure includes a sequence of annual histograms for the simple and full stochastic environments. The first graph relates to the simple setting and the second describes the full setting.

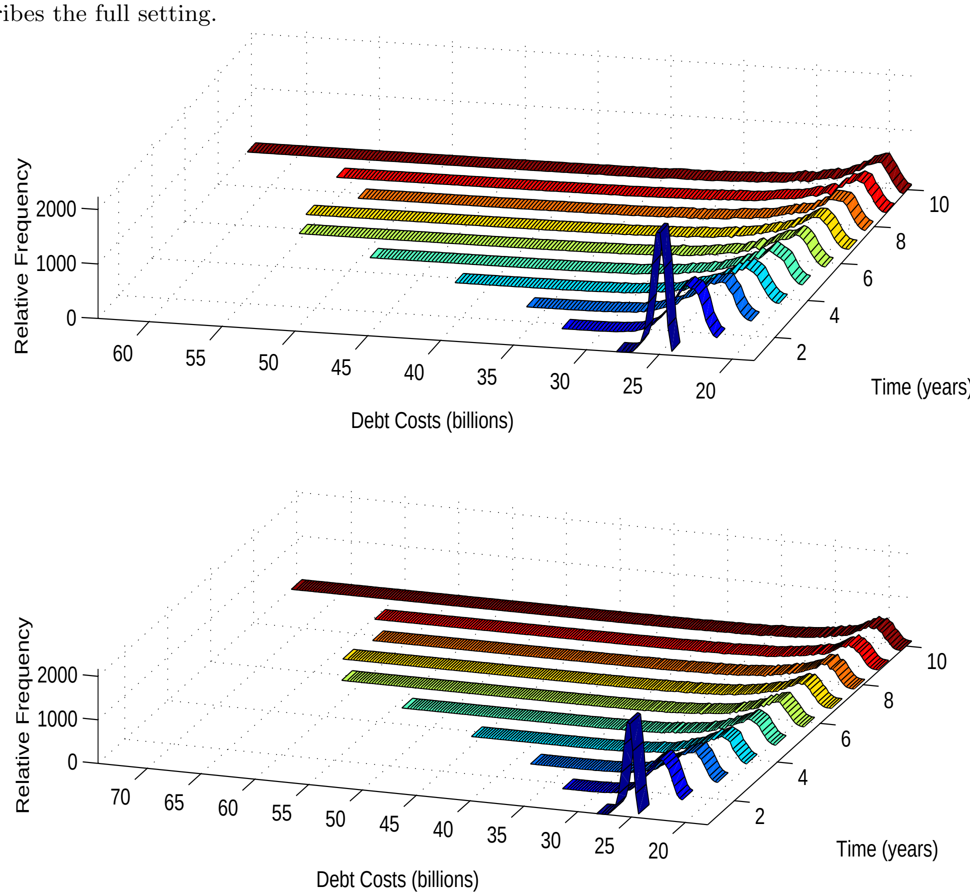

the worst-case deviation from the mean ranges from \$4.8 to \$6.3 billion in the first year under our two measures. In contrast, the 100 per cent bonds strategy exhibits a worst-case deviation of only \$800 million to \$1.1 billion. We also observe that, in a manner similar to the cost volatility measure, both the RCaR and RTCaR measures increase over time and the distinction between the various financing strategies falls in relative terms. The first-year RCaR of the 100 per cent treasury-bill financing strategy, for example, is almost six times greater than the RCaR of the 100 per cent bond strategy; in the tenth year, however, it is only 1.6 times larger.

This raises another question. How do we interpret annual debt-charge-based measures of risk, whether they be RCaR, RTCaR, or cost volatility? While the reason for the increased volatility of these measures over time is clear, it does not help us with their application. Indeed, in the tenth year of the analysis, the two

Table 8: **Percentile Measures of Risk**: This table describes the relative CaR and expected tail CaR for the first, fifth, and tenth year of the analysis horizon. These values are computed based on the 95th percentile of each annual debt cost distribution using equations (17) and (19).

	Year 1		**Year 5**		**Year 10**	
Financing strategy	**Relative CaR**	**Relative tail CaR**	**Relative CaR**	**Relative tail CaR**	**Relative CaR**	**Relative tail CaR**
Simple stochastic environment						
100% bills	4.87	6.37	11.73	16.50	13.33	19.97
75% bills	3.87	5.05	10.08	14.18	11.91	17.83
50% bills	2.85	3.73	8.43	11.91	10.54	15.84
25% bills	1.83	2.41	6.85	9.74	9.33	14.03
100% bonds	0.83	1.11	5.43	7.74	8.36	12.50
Full stochastic environment						
100% bills	4.98	6.51	12.39	17.60	14.16	21.74
75% bills	3.95	5.17	10.63	15.16	12.75	19.49
50% bills	2.91	3.82	8.86	12.73	11.22	17.29
25% bills	1.87	2.47	7.18	10.41	10.16	15.36
100% bonds	0.87	1.15	5.65	8.29	9.11	13.72

CaR-based measures provide estimates that are of similar magnitude to the expected cost of the portfolio. This would suggest that, as we move further along the time dimension, annual debt-charge-based measures of risk become substantially less useful as a policy tool.

One of the reasons for the tremendous variance in annual debt-charge based measures is that we are examining the *unconditional* distribution of debt costs at each point in time. The only exception, of course, is the first year of the analysis, because all of the stochastic realizations begin from the same initial starting point. When we examine the debt cost distribution for the fifth year, however, we are not conditioning on the value of the debt charges for the fourth year in the simulation. This information is nevertheless critical, because the entire simulation is strongly *path-dependent*. That is, current debt charges will depend in an important manner on the debt charges—and thus the stochastic environment—prevailing in the previous periods. We need to explicitly consider this path-dependency in any useful measure of risk to be applied over a longer-term time horizon.

It is exactly due to these concerns that we developed the conditional cost volatility measure, in section 3.3, that deals with the entire time dimension. Indeed, this is quite a simple approach. If our concerns stem from not conditioning on previous outcomes in each sample path, then why not explicitly condition on these values? In other words, our risk measure represents the residual uncertainty in forecasting the next year's debt charges conditioning on the current value. Moreover, given the nature of our stochastic model, which involves strong mean reversion of our stochastic processes, this distribution is stationary. As such, we can safely interpret our measure of risk in the natural way.

Table 9 demonstrates the results of the pathwise estimation of an AR(1) model by financing strategy and stochastic environment. The first columns of Table 9, ϕ_0, ϕ_1, and ξ, are the parameters that we use to describe the conditional debt cost distribution. ϕ_0 is an intercept term and can be considered as the base debt charges that are independent of the current state of the stochastic environment. The second parameter, ϕ_1, essentially describes the stationarity of conditional distribution. Consider this value as representing how much weight is being placed on the current value when forecasting the amount of debt charges for the next year. Caution is required, however, because when the model places too much weight on the previous observation we run into stability problems. A value of $\phi_1 = 1$, for example, would imply that,

$$c_t - c_{t-1} = \phi_0 + \epsilon_t. \tag{30}$$

In other words, this would imply that the process, $\{c_t, t \geq 0\}$, is entirely random and has non-finite variance.[41] This would be highly problematic. Fortunately, all of the estimated values suggest stationary conditional distributions.

The third parameter, ξ, representing the conditional volatility is the most useful for our purposes. As stated earlier, this single measure describes the uncertainty associated with forecasting debt charges for the subsequent period using this AR(1) model. We note the same trends in this measure as we detected in the previous analysis. In particular, the values are higher for the full stochastic environment than for the simple setting. Moreover, the short-term debt dominated financing strategies exhibit significantly more—on the order of five times for the two extreme financing strategies—conditional volatility relative to longer-term-debt focused financing strategies. Thus, this measure is consistent with our previous analysis.

The final two columns exploit the structure of our AR(1) model to compute the first two moments of the unconditional debt cost distribution. The unconditional mean values are similar in both level and trend to the previously considered expected cost values. Also observe that the unconditional volatility is significantly larger—on the order of \$300 to \$400 million—than the conditional volatility. This represents the increase in the efficiency of the forecast associated with being able to condition one's forecasts of future debt costs on the current level.

One concern with this approach is that these parameters are computed based on an implicit assumption of a Gaussian conditional distribution. Unfortunately, based on the nature of our stochastic environment, this is not strictly true. To assess this more carefully, Figure 9 includes four separate histograms of the first differences ($\Delta c_t = c_t - c_{t-1}$ for $t = 2, 3, 6$, and 10) of the cost distribution for the 50 per cent treasury-bill financing strategy. These histograms, which each have an overlaid normal density, indicate that they are *not* Gaussian.[42] In particular, each histogram exhibits positive skewness, an excess of observations in the tails, and too much mass in the centre of the distribution. Nevertheless, as a first-order approximation, the

[41] In the econometric literature, of course, this is referred to as a *unit root*.

[42] The results are highly similar for the other four financing strategies.

Table 9: **Conditional Cost Distribution**: This table describes the parameters of the conditional cost density. These are estimated assuming that this density is reasonably described by a first-order autoregressive process.

Financing strategy	**Conditional parameters**			**Unconditional parameters**	
	Intercept (ϕ_0)	**Slope** (ϕ_1)	**Conditional volatility** (ξ)	**Mean** $\left(\frac{\phi_0}{1-\phi_1}\right)$	**Volatility** $\left(\sqrt{\frac{\xi^2}{1-\phi_1^2}}\right)$
Simple stochastic environment					
100% bills	8.17	0.52	2.42	16.89	2.83
75% bills	8.39	0.55	1.91	18.59	2.29
50% bills	8.02	0.60	1.41	20.27	1.76
25% bills	6.60	0.70	0.93	21.69	1.30
100% bonds	4.45	0.80	0.50	22.12	0.83
Full stochastic environment					
100% bills	8.49	0.52	2.47	17.88	2.89
75% bills	8.59	0.56	1.97	19.50	2.38
50% bills	8.00	0.61	1.46	20.77	1.84
25% bills	6.57	0.70	0.96	21.80	1.34
100% bonds	4.53	0.80	0.53	22.16	0.88

assumption of Gaussianity does not seem unreasonable.

Clearly, more work is required in considering more elaborate techniques for estimating these conditional densities. We could, for example, attempt to find a better approximation for these first-difference distributions and explicitly account for this using maximum-likelihood estimation. One could also use non-parametric kernel-based regressions to avoid making any assumptions about the form of the conditional density. Nevertheless, we consider the current approach as a first-order approximation that demonstrates a general approach to measuring the longer-term risk associated with a given financing strategy.

With a good understanding of the conditional distribution, we could easily extend these ideas by computing a worst-case conditional volatility. Assuming that the assumption of Gaussianity is reasonable, for example, we could calculate something like a *time-conditional cost-at-risk* (TCCaR) as,

$$\text{TCCaR} = \Phi^{-1}\left(1 - \frac{\alpha}{2}\right)\xi, \tag{31}$$

where α represents the critical level for the tail of the distribution. This could, of course, be generalized to other distributional assumptions or computed non-parametrically.

4.3 Sensitivity analysis

The results of any mathematical model depend on the nature of the assumptions made in its construction. It is therefore essential that we include a facility in our analysis to test these assumptions. In this subsection, we consider some alternative parameterizations for our full stochastic model. Obviously, although there is

Figure 9: **First-Difference Distributions**: This figure summarizes four separate histograms for the first differences of the 50 per cent treasury-bill financing strategy. Overlaid on each histogram is a normal density to assess the hypothesis of a Gaussian conditional density for the government's debt costs.

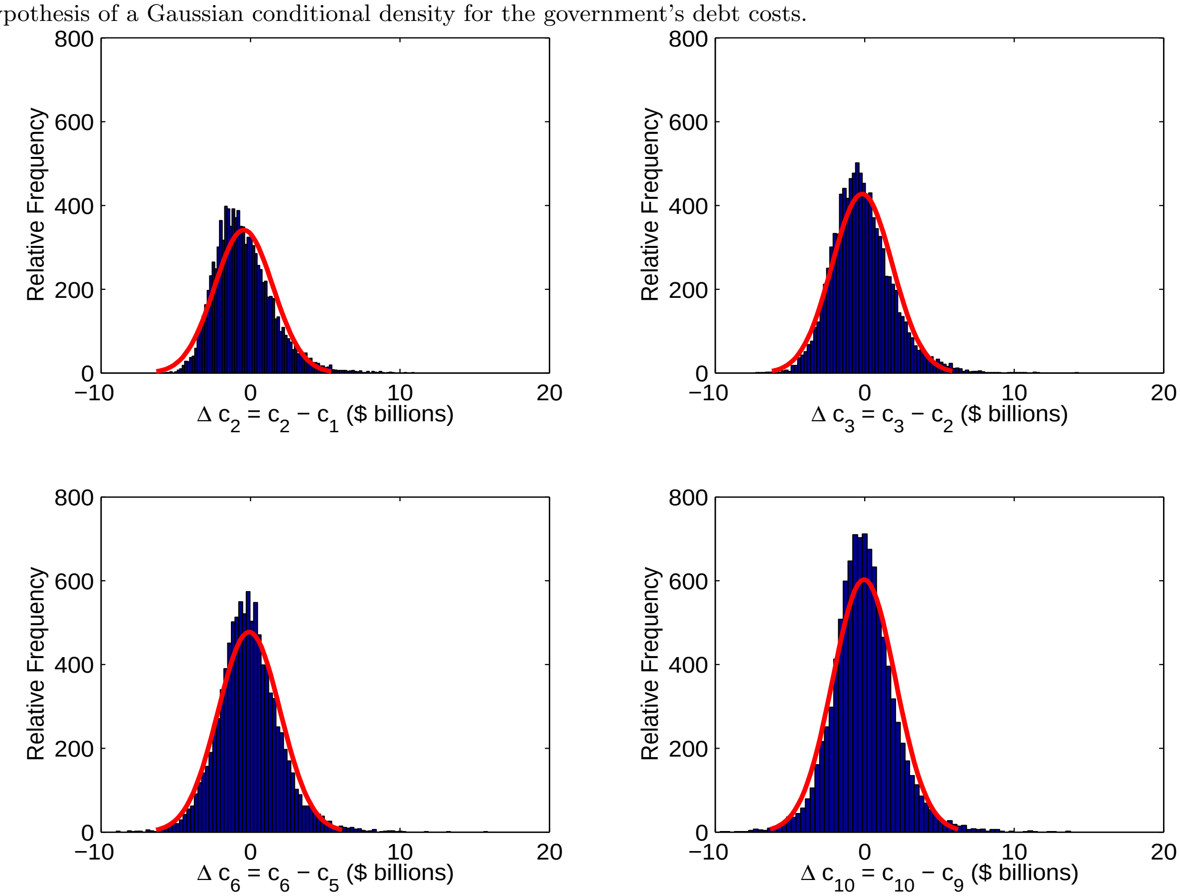

scope for a much higher level of sensitivity analysis, we are constrained in what can be effectively presented. We have selected, therefore, two alternative specifications for the volatility of our stochastic system that we believe are quite interesting. First, we hold all other parameters constant and increase the volatility of the government's financial requirement process by a factor of two and a half. We selected this specification for our sensitivity analysis because the government's financial requirements are critical in determining *how much must be borrowed* in each given period. Note, however, that the impact of a recession on the government's financial requirements remains unchanged. The altered parameter describes the financial requirement volatility that is orthogonal to the general state of the macroeconomy. In other words, we are considering a situation where the government's budgetary position is subject to a greater degree of noise.

The second parametric change to our model involves our interest-rate model. Again, we hold all other parameters constant and increase the volatility of the dynamics of the term structure of interest rates. Specifically, we increase the volatility of the first state variable—associated with the general slope of the term structure—by a factor of 1.3 and augment the volatility of the second state variable—which describes parallel movements in the term structure—by a quarter. This second candidate for sensitivity analysis was selected given that interest rates are essential in determing *how much must be paid* to finance new borrowing in each period.[43]

As in the previous section, we consider the five financing strategies described in Table 2 for 10,000 scenarios with an initial debt stock of $400 billion. Table 10 illustrates the mean and associated volatility of the debt charges across all financing strategies for each sensitivity analysis. A natural point of comparison is Table 7 on page 41. Interestingly, there is a relatively small increase in the mean and volatility of debt charges when we increase the variability of the government's financial requirement process. The increase in cost volatility for the first sensitivity analysis, in fact, seems to appear primarily for the longer-term-debt dominated strategies.

Table 10: **Expected Debt Charges and Volatility by Financing Strategy**: This table outlines the expected debt charges and the associated volatility (standard deviation) of these debt charges for the first, fifth, and tenth year of the analysis horizon.

Financing strategy	Year 1		Year 5		Year 10	
	Mean	**Volatility**	**Mean**	**Volatility**	**Mean**	**Volatility**
Increased financial requirement volatility						
100% bills	19.22	2.76	18.74	6.19	18.87	7.50
75% bills	20.87	2.19	20.16	5.40	20.10	6.83
50% bills	22.22	1.62	21.29	4.59	21.07	6.15
25% bills	23.73	1.05	22.56	3.87	22.22	5.61
100% bonds	25.51	0.51	24.02	3.28	23.53	5.23
Increased term-structure volatility						
100% bills	19.24	5.10	18.61	8.66	18.98	9.27
75% bills	20.88	4.01	19.99	7.18	20.06	7.80
50% bills	22.22	2.91	21.11	5.71	20.94	6.37
25% bills	23.73	1.82	22.36	4.32	21.97	5.07
100% bonds	25.50	0.74	23.81	3.09	23.16	4.01

Conversely, we observe a significant jump in the volatility of debt costs associated with augmented term-structure variance. Specifically, the 100 per cent treasury-bill financing strategy demonstrates a cost volatility of more than $2 billion more than the base parameterization outlined in Table 7. This dampens out substantially as we move towards financing strategies with a greater emphasis on long-term borrowing. The

[43]The exact specifics of our parameter selection and associated diagnostics for our stochastic model are described in Appendix 5.

mean level of debt costs, however, remains almost identical to that evidenced by the base parameterization. This suggests that, on average, the term structure has the same general form, albeit with a greater level of volatility.

Table 11, which is directly comparable with Table 8 on page 45, provides an overview of our two *tail* risk measures for the two sensitivity analyses. The same trend is evident for the increased financial requirement volatility. Specifically, the short-term-debt based financing strategies are virtually identical to the base parameterization, while the longer-term-debt based strategies exhibit a somewhat larger tail risk. One possible reason for this effect is that the total amount of short-term borrowing required under the treasury-bill-dominated strategies are relatively insensitive to small changes in the volatility of the government's financial requirements. The absolute level of borrowing among the bond-dominated financing strategies, however, is quite small, and is perhaps more sensitive to increased volatility of government financial requirements. That is, the error in the financial requirement model that is *independent* of the business cycle appears to be relatively unimportant to the overall results. Other aspects of the business cycle model also need to be examined to determine their importance to the simulation results.

Table 11: **Percentile Measures of Risk**: This table describes the relative CaR and expected tail CaR for the first, fifth, and tenth year of the analysis horizon.

Financing strategy	**Year 1**		**Year 5**		**Year 10**	
	Relative CaR	**Relative tail CaR**	**Relative CaR**	**Relative tail CaR**	**Relative CaR**	**Relative tail CaR**
Increased financial requirement volatility						
100% bills	4.87	6.40	11.88	17.32	14.38	22.60
75% bills	3.86	5.08	10.33	15.09	12.94	20.53
50% bills	2.85	3.76	8.79	12.83	11.51	18.45
25% bills	1.85	2.45	7.37	10.75	10.35	16.68
100% bonds	0.91	1.20	6.14	8.96	9.80	15.28
Increased term-structure volatility						
100% bills	9.24	12.74	16.44	24.39	17.55	26.66
75% bills	7.27	10.00	13.64	20.25	15.01	22.60
50% bills	5.30	7.26	10.74	16.15	12.00	18.67
25% bills	3.31	4.53	8.13	12.31	9.61	15.14
100% bonds	1.36	1.84	5.84	8.91	7.83	12.22

The increased interest-rate volatility sensitivity analysis, not surprisingly, generates large increases in the estimated tail risk. The relative CaR for the 100 per cent bills strategy, for example, is close to twice the value associated with the base parameterization. Increased interest-rate volatility has a greater impact on financing strategies with larger portions of short-term debt that must be frequently refinanced. This is consistent with both the behaviour of our term-structure model and generalized stylized facts about the term structure of interest rates. In particular, the short end of the term structure is substantially more variable than the long end.

Table 12: **Conditional Cost Distribution**: This table describes the parameters of the conditional cost density. These are estimated assuming that this density is reasonably described by a first-order autoregressive process.

Financing strategy	Conditional parameters			
	Intercept (ϕ_0)	**Slope (ϕ_1)**	**Conditional volatility (ξ)**	**Unconditional mean $\left(\frac{\phi_0}{1-\phi_1}\right)$**
Increased financial requirement volatility				
100% bills	8.25	0.54	2.47	17.76
75% bills	8.20	0.58	1.97	19.37
50% bills	7.58	0.63	1.47	20.44
25% bills	6.19	0.71	1.00	21.49
100% bonds	4.41	0.80	0.61	21.92
Increased term-structure volatility				
100% bills	9.72	0.45	4.55	17.54
75% bills	10.17	0.47	3.59	19.23
50% bills	9.98	0.51	2.62	20.52
25% bills	8.83	0.59	1.68	21.75
100% bonds	6.03	0.73	0.82	22.56

The final set of measures involves the conditional cost volatility outlined in Table 12. There are relatively few surprises in this table compared with the conclusions we have drawn from Tables 10 and 11. That is, there is relatively little difference—compared with the base parameterization—for the financial requirement sensitivity analysis, but a substantial increase in the conditional cost volatility associated with increased interest-rate variability.

The strong concordance of the various results in this section would suggest that interest rates are more important to the final results than the government's financial requirements. While this is not a terribly surprising result, the scope of the difference is larger than one might expect. One possible reason for the relative unimportance of the government's financial requirements relates to the assumed level of accuracy of the government's debt forecasts. In equation (8), we postulate our model for the government's determination of the magnitude of debt charges for the subsequent period. If this assumption overestimates the government's ability to predict and react to shocks to the level of financial requirements, then we will systematically underestimate the importance of financial requirement volatility in our analysis.

4.4 Stress testing

The stochastic simulation framework proposed in this paper attempts to model the impact of key macroeconomic variables under normal market conditions on various financing strategies. It is nevertheless reasonable to ask how our financing strategies might fare under *abnormal* market conditions. The examination of abnormal outcomes is called *stress testing* and it represents an important element in our debt strategy analysis. The usual approach to this problem is to construct a negative set of outcomes and examine how one's portfolio performs under these situations. This standard stress-testing methodology, however, does not make any statement about the probability of the occurrence of abnormal market conditions. This feature of the conventional stress-testing approach is not ideal because, in the absence of any sense of their likelihood, it makes stress-test results difficult to interpret. For example, imagine that a particular financing strategy performs quite badly under one possible stress scenario, but is otherwise quite attractive under normal market conditions. If debt managers believe that the stress scenario has a relatively high probability of occurring, then this would likely preclude consideration of this strategy. If, conversely, the expected probability of this stress scenario is relatively low, then this may not be as problematic for the financing strategy in question. Lacking an assessment of the probability of the extreme event, therefore, the situation is intractable.

Given this drawback of standard stress-testing approaches, we suggest an alternative methodology. We wish to specify a state of the world that involves extreme outcomes for our stochastic model. While this may seem vague, it is easily put into operation in the context of our Markov-chain based model of the evolution of the business cycle. As discussed in Bolder (2002), the macroeconomic business cycle is described by a two-state hidden Markov model with the estimated transition probabilities,

$$\begin{aligned}\mathbb{P}[S_t = 0 \mid S_{t-1} = 0] &= q = 0.53, \\ \mathbb{P}[S_t = 1 \mid S_{t-1} = 1] &= p = 0.96,\end{aligned} \tag{32}$$

yielding the following transition matrix,

$$P = \begin{bmatrix} q & 1-p \\ 1-q & p \end{bmatrix} = \begin{bmatrix} 0.53 & 0.04 \\ 0.47 & 0.96 \end{bmatrix}. \tag{33}$$

Using this structure, we propose the addition of a third state to our Markov chain that corresponds to a set of extreme outcomes. This could be a large parallel shift in the term structure of interest rates, substantially increased interest rate volatility, or term-structure inversion. It could also involve negative outcomes for the government's financial position. This is similar in spirit to the *peso problem*. That is, the large, black-market depreciation of the Mexican peso during the late 1970s was attributed to the small probability of a hidden regime with huge consequences.[44] Clearly, we have no way to estimate the transition probabilities of such a state, so, instead, we assign some arbitrary but small probability of occurrence to this third state.

To alter our transition matrix, we define $S_t = 2$ as the economy residing in the extreme state in period

[44] In this case, the event was the devaluation of the peso by Mexican authorities. It eventually occurred in 1982. For more details, see DeGrauwe (1989, page 129).

t. Furthermore, we assume that the transition probabilities have the following form,

$$\begin{aligned} \mathbb{P}[S_t = 2 \mid S_{t-1} = 0] &= \varepsilon, \\ \mathbb{P}[S_t = 2 \mid S_{t-1} = 1] &= \varepsilon, \end{aligned} \tag{34}$$

where ε is a small, arbitrary, positive value. Our goal is to introduce a third extreme state with the minimum possible perturbation to our macroeconomic model. We are assuming, therefore, that the probability of transitioning into the extreme state is equal, and small, whether one currently resides in a recession or an expansion. This assumption is based on the idea that these extreme states are related to large economic shocks—such as war or some other form of geopolitical turmoil—that are independent of the current state of the economy. We construct the revised transition matrix as,

$$\begin{aligned} \tilde{P} &= \begin{bmatrix} \mathbb{P}[S_t = 0 \mid S_{t-1} = 0] & \mathbb{P}[S_t = 0 \mid S_{t-1} = 1] & \mathbb{P}[S_t = 0 \mid S_{t-1} = 2] \\ \mathbb{P}[S_t = 1 \mid S_{t-1} = 0] & \mathbb{P}[S_t = 1 \mid S_{t-1} = 1] & \mathbb{P}[S_t = 1 \mid S_{t-1} = 2] \\ \mathbb{P}[S_t = 2 \mid S_{t-1} = 0] & \mathbb{P}[S_t = 2 \mid S_{t-1} = 1] & \mathbb{P}[S_t = 2 \mid S_{t-1} = 2] \end{bmatrix}, \\ &= \begin{bmatrix} q & 1-p-\varepsilon & 1-m \\ 1-q-\varepsilon & p & 0 \\ \varepsilon & \varepsilon & m \end{bmatrix}, \\ &= \begin{bmatrix} 0.53 & 0.04-\varepsilon & 0.70 \\ 0.47-\varepsilon & 0.96 & 0 \\ \varepsilon & \varepsilon & 0.30 \end{bmatrix}. \end{aligned} \tag{35}$$

Inspection of this transition matrix reveals our remaining assumptions. First, we have assumed that the probability of remaining in an extreme state is relatively low at 30 per cent. Second, we have assumed a zero probability for the economy's transition from an extreme state to an expansionary state. That is, if the economy currently resides in the extreme state, it can either remain in that state of the world (with probability $m = 30$ per cent) or transition into a recessionary state (with probability $1 - m = 70$ per cent).

In this section, we will examine the risk and cost characteristics—using our usual battery of measures—for $\varepsilon = 0.05$ per cent and 1 per cent. These are arbitrarily selected values intended to demonstrate how the approach works in general. It is nevertheless natural to ask how we interpret a given selected value for ε. One of the pleasant analytical properties of Markov chains is that we can—under certain assumptions that are satisfied in our analysis—compute the unconditional probabilities implied by our choice of transition matrix.[45] The unconditional probability of recession in the two-state model is 8.06 per cent, implying that

[45] For an N-state Markov chain, the $N \times 1$ vector of ergodic (or steady-state) probabilities, denoted π, satisfies the following,

$$P\pi = \pi. \tag{36}$$

the corresponding unconditional probability of expansion is 91.94 per cent. When we set

$$\mathbb{P}[S_t = 2 \mid S_{t-1} = 0] = \mathbb{P}[S_t = 2 \mid S_{t-1} = 1] = \varepsilon, \tag{38}$$

then we slightly alter these probabilities. If $\varepsilon = 0.5$ per cent, then the unconditional (or ergodic) probability of being in the extreme state is 0.7 per cent. If, however, we double ε to 1 per cent, then the ergodic probability also doubles to 1.4 per cent. Another way to think about this is that, under this assumption, the econonomy should on average find itself in an extreme state 1.4 quarters (just over four months) out of every 100 quarters (25 years). These assumptions for ε lead to slight alterations in the ergodic probabilities of recession and expansion of 8.09 per cent and 91.20 per cent for $\varepsilon = 0.5$ per cent. These probabilities are correspondingly 8.11 per cent and 90.48 per cent when we set $\varepsilon = 1$ per cent.

Table 13: **Expected Debt Charges and Volatility by Financing Strategy**: This table outlines the expected debt charges and the associated volatility (standard deviation) of these debt charges for the first, fifth, and tenth year of the analysis horizon.

Financing strategy	**Year 1**		**Year 5**		**Year 10**	
	Mean	**Volatility**	**Mean**	**Volatility**	**Mean**	**Volatility**
First stress test: $\varepsilon = 0.5\%$						
100% bills	19.49	3.25	19.45	6.82	19.50	7.84
75% bills	21.09	2.59	20.77	5.87	20.66	6.99
50% bills	22.39	1.91	21.83	4.92	21.61	6.16
25% bills	23.84	1.25	23.01	4.03	22.71	5.44
100% bonds	25.57	0.59	24.39	3.24	23.99	4.86
Second stress test: $\varepsilon = 1\%$						
100% bills	19.72	3.80	19.91	7.35	20.01	8.32
75% bills	21.27	3.03	21.17	6.33	21.12	7.41
50% bills	22.52	2.24	22.17	5.31	22.01	6.52
25% bills	23.93	1.47	23.29	4.37	23.06	5.76
100% bonds	25.60	0.70	24.60	3.54	24.27	5.14

The next question is, what is the nature of the extreme scenarios? We decided to alter our term-structure model by increasing the long-term instantaneous interest rate by 600 basis points from 4 per cent to 10 per cent, raise the volatility of the first state variable by 130 per cent and the second state variable by 70 per cent, and finally increase the average slope of the term structure by approximately 75 basis points.[46] Clearly,

It is the vector, π, that describes the steady-state probabilities of a Markov chain. Moreover, we have the following result,

$$\lim_{m\to\infty} P^m = \pi\vec{1}, \tag{37}$$

where $\vec{1}$ is a $1 \times N$ row vector of ones. See Bolder (2002, Appendix A) for more details.

[46]One could also, in this framework, alter the parameters of the term-structure model such that the term structure of interest rates is primarily inverted.

this is both a negative and extreme state of affairs for a sovereign borrower. It is, nevertheless, arbitrary and was selected entirely for illustrative purposes. One could also simultaneously generate extreme outcomes for the government's financial requirements.

Table 14: **Percentile Measures of Risk**: This table describes the relative CaR and expected tail CaR for the first, fifth, and tenth year of the analysis horizon.

	Year 1		Year 5		Year 10	
Financing strategy	**Relative CaR**	**Relative tail CaR**	**Relative CaR**	**Relative tail CaR**	**Relative CaR**	**Relative tail CaR**
First stress test: $\varepsilon = 0.5\%$						
100% bills	4.48	8.62	12.85	19.73	15.08	23.86
75% bills	3.53	6.86	11.17	17.00	13.80	21.33
50% bills	2.60	5.08	9.32	14.28	12.10	18.87
25% bills	1.67	3.31	7.64	11.74	10.78	16.76
100% bonds	0.76	1.56	6.10	9.46	9.72	15.03
Second stress test: $\varepsilon = 1\%$						
100% bills	4.93	11.17	13.73	21.58	15.94	25.40
75% bills	3.91	8.90	11.78	18.61	14.39	22.64
50% bills	2.86	6.60	9.84	15.65	12.72	19.98
25% bills	1.82	4.32	8.09	12.92	11.31	17.71
100% bonds	0.82	2.05	6.50	10.53	10.21	15.94

Table 13 outlines the mean and cost volatility for our five financing strategies with $\varepsilon = 0.05$ per cent and 1 per cent. Again, we should be comparing these results with the *normal* conditions summarized in Table 7. Note that, even with these very small probabilities of extreme events, we observe higher expected costs and associated volatility across all financing strategies. The impact, in terms of both cost and volatility, of the presence of an extreme state is greater for those financing strategies with larger proportions of short-term debt instruments. This is to be expected, given the much larger refinancing burden associated with these portfolios, and thus the enhanced exposure to negative market conditions. In other words, the increased expected cost of a longer-term-debt dominated financing strategy can be considered as something of an *insurance premium* given its relative insulation against extreme events. Ultimately, we see this style of stress testing as a tool for determining the probability of an extreme state (i.e., ε) that is required to make this insurance premium worth paying.

Stress testing, in general, is an exercise in examining tail events. Therefore, it is critical to examine tail-based measures such as CaR and conditional tail CaR. We have already established that these measures are of questionable use over longer-term time horizons. Even worse, as these are events that occur with very low probability, we need to examine long-term time horizons to get a reasonable estimate of their impact on the government's debt cost distribution. A one-year relative cost-at risk measure, for example, is not sufficiently long for us to assess the impact of low-probability extreme events. This is clear from a comparison of Table 14 with Table 8 on page 45. The one-year RCaR and RTCaR measures are virtually identical to the results under normal conditions. The longer-term RCaR and RTCaR measures, however, do dominate the corresponding values for the normal stochastic environment. Clearly, a one-year time horizon is not sufficiently long to adequately perform stress testing. This would suggest that the measures in Table 14 are of little practical use.

Table 15: **Conditional Cost Distribution**: This table describes the parameters of the conditional cost density. These are estimated assuming that this density is reasonably described by a first-order autoregressive process.

Financing strategy	**Conditional parameters**			
	Intercept (ϕ_0)	**Slope** (ϕ_1)	**Conditional volatility** (ξ)	**Unconditional mean** $\left(\frac{\phi_0}{1-\phi_1}\right)$
First stress test: $\varepsilon = 0.5\%$				
100% bills	8.95	0.52	2.59	18.48
75% bills	9.02	0.55	2.06	19.99
50% bills	8.40	0.60	1.53	21.09
25% bills	6.89	0.69	1.01	22.20
100% bonds	4.60	0.80	0.56	22.51
Second stress test: $\varepsilon = 1\%$				
100% bills	9.24	0.51	2.74	19.01
75% bills	9.26	0.55	2.17	20.41
50% bills	8.63	0.60	1.61	21.45
25% bills	7.12	0.68	1.07	22.57
100% bonds	4.75	0.79	0.60	22.75

Our conditional cost volatility measure avoids this problem. It considers, by construction, the entire time interval used in the analysis and synthesizes this information into a single risk measure. While not a panacea, we do believe it is a step in the right direction. Table 15 summarizes the parameters for our model of the conditional cost distribution under extreme conditions for our five financing strategies. These are directly comparable with the results under normal conditions outlined in Table 9. Observe that the conditional volatility, unconditional mean, and unconditional volatility for both assumptions on the value of ε dominate those values in Table 9. The conditional cost volatility for $\varepsilon = 1$ per cent, for example, is approximately \$250 million greater than the associated value computed under the normal stochastic environment. It is nevertheless reassuring to see that the differences are relatively small, given the small probability of occurrence associated with these extreme outcomes.

Overall, we consider that this stress-testing methodology is useful for a sovereign debt manager who is determining the probability of an extreme state (i.e., ε) required for the government to alter its risk preferences and, consequently, its financing strategy.

4.5 Examining multiple financing strategies

In analyzing the five different financing strategies, we have learned a few lessons about the nature of debt strategy. In addition, we have explored a number of features of our stochastic framework. This analysis has nonetheless failed to describe the relationship between cost and risk for a large number of financing strategies. In this subsection, we will address this issue in some detail. The first challenge is how to organize and analyze the overwhelming number of results associated with a substantial number of financing strategies—indeed, we found it cumbersome in the context of five financing strategies. Our suggested approach is to examine the relationship between the portfolio weights—in other words, the financing strategy—and various measures of the cost and risk of the resulting debt-charge distributions for each of these financing strategies.

To make this clearer, let us define the number of financing strategies as K and the number of available debt instruments as H. We define the proportion of debt in instrument h for the kth financing strategy as ω_{kh}. In vector form, we have

$$\omega_k \triangleq \begin{bmatrix} \omega_{k1} \\ \omega_{k2} \\ \cdots \\ \omega_{kH} \end{bmatrix}, \tag{39}$$

for $k = 1, ..., K$. In aggregate, we may represent the portfolio weights in matrix form as,

$$\Omega \triangleq \begin{bmatrix} \omega_1 & \omega_2 & \cdots & \omega_K \end{bmatrix}. \tag{40}$$

Thus the matrix Ω, with dimensions $H \times K$, is a compact representation for the proportion of debt issued in the available instruments for each financing strategy. Let us further define a financing strategy that puts all of its weight on a single debt instrument as a *bullet* portfolio.

We need a tool to facilitate our analysis. If y is an arbitrary measure of debt cost or risk, then we can examine the relationship between portfolio composition and this measure as follows,

$$y = \Omega\beta + \epsilon, \tag{41}$$

where β is a vector of K parameters and ϵ represents the error in this linear representation.[47] One approach, of course, to determine the parameter vector is ordinary least squares. In other words, we plan to examine the relationship between Ω and various measures of risk and cost in a highly straightforward univariate setting. Note, however, that in this setting we cannot arbitrarily select Ω. The well-known solution to the

[47] A potential complicating factor in this approach is that y, our measure of cost or risk, is computed through simulation and thus is observed with error. This is discussed in Judge et al. (1985), among others, and it can give rise to certain statistical problems. We will, however, abstract from these potential complications during this analysis.

least-squares problem arising from equation (41) is,

$$\beta = (\Omega^T\Omega)^{-1}\Omega^T y. \tag{42}$$

If $\Omega^T\Omega$ is singular, we cannot solve this problem. One certain way to achieve this is to require that the rows of Ω, or rather each ω_k, $k = 1, ..., K$, be linearly independent. Moreover, we require the additional two conditions,

$$\omega_{kh} \geq 0, \tag{43}$$

for $h = 1, ..., H$ and $k = 1, ..., K$, and

$$\sum_{i=1}^{H} \omega_{kh} = 1, \tag{44}$$

for $k = 1, ..., K$. Even with these constraints, there is a very large, indeed an infinite, number of possibly linearly independent weights. For relatively large N, however, it is not terribly difficult to generate a singular $\Omega^T\Omega$. The resulting approach used in this analysis is to carefully construct 225 linearly independent sets of portfolio weights. This is *not* a basis for the space of linearly independent weighting vectors with the conditions in equations (43) and (44).[48]

Table 16: **OLS Regression of Mean and Volatility**: This table lists the coefficients and R^2 for the linear regression describing the relationship between the financing strategy and both debt-charge mean and volatility.

Regression statistics	**Year 1**		**Year 5**		**Year 10**	
	Mean	**Volatility**	**Mean**	**Volatility**	**Mean**	**Volatility**
3 months (β_1)	19.78	5.12	20.58	7.89	20.57	8.51
6 months (β_2)	18.70	3.57	18.67	6.87	18.75	7.82
1 year (β_3)	20.08	1.88	18.89	5.80	19.03	7.00
2 years (β_4)	22.19	1.04	19.82	5.05	19.97	6.47
5 years (β_5)	24.84	0.50	22.75	4.19	22.76	6.25
10 years (β_6)	26.22	0.49	25.55	2.54	24.47	5.00
30 years (β_7)	30.29	1.09	30.59	1.96	30.03	2.42
R^2 coefficient	0.97	0.80	0.94	0.78	0.94	0.81

To permit this analysis, we ran 225 financing strategies using the full stochastic environment with the same parameterization used in section 4.2. We considered 2,500 scenarios for each financing strategy with an initial portfolio of \$400 billion.[49] Table 16 illustrates the results of the estimation of equation (41) for

[48] The idea is that we consider K bullet portfolios consisting entirely of issuance in the kth debt instrument. We then construct linearly independent portfolios of instrument k and $k+1$, then $k, k+1, k+2$, and repeat this until $K-1$. This leaves one debt instrument absent. We then move forward to instrument $k+1$ and repeat. The final addition is the equally weighted portfolio.

[49] We were forced to reduce the number of scenarios, relative to the previous analysis, given computational constraints.

the first two moments of the debt cost distribution for the first, fifth, and tenth year of the analysis. Given the nature of the data, it is quite straightforward to interpret the regression coefficients. The value of β_1 for the first-year mean debt cost is the predicted first-year cost of a financing strategy consisting entirely of 3-month treasury bills. In other words, each regression coefficient describes the cost or risk for a bullet portfolio. An equally weighted financing strategy, therefore, would involve a proportionate share in each of the coefficients. Finally, a review of the coefficient values indicates that the estimated values are quite reasonable. Bullet portfolios of longer-term debt, for instance, are more costly and less volatile than bullet portfolios of shorter-term debt.

The most striking feature of Table 16 is the size of the R^2 coefficients in the final row. This regression statistic summarizes the amount of overall sample variance that is described by our linear model; the measure of both mean and volatility of cost are very substantial. We observe, for example, that 97 per cent of the variation in average debt costs is summarized by this linear model in the financing strategy weights; the associated value is 80 per cent for the first-year cost volatility. A very similar pattern is evident in Table 17, which outlines the regression results for our two measures of tail risk. While the RTCaR measure is somewhat less well described by a linear model than the RCaR measure, they both exhibit relatively high R^2 values.

Table 17: **OLS Regression of Percentile Measures of Risk**: This table lists the coefficients and R^2 for the linear regression describing the relationship between the financing strategy and both the RCaR and RTCaR measures.

	Year 1		**Year 5**		**Year 10**	
Regression statistics	**Relative CaR**	**Relative tail CaR**	**Relative CaR**	**Relative tail CaR**	**Relative CaR**	**Relative tail CaR**
3 months (β_1)	7.99	15.84	14.80	23.23	16.22	25.75
6 months (β_2)	5.69	9.98	13.13	20.03	14.99	23.54
1 year (β_3)	3.22	4.46	11.30	16.76	13.41	21.14
2 years (β_4)	1.98	2.27	10.07	14.30	12.44	19.87
5 years (β_5)	0.80	1.37	7.81	12.34	12.68	19.71
10 years (β_6)	0.60	1.68	4.59	7.73	9.77	15.61
30 years (β_7)	1.03	4.95	3.27	6.29	4.59	7.60
R^2 coefficient	0.93	0.63	0.85	0.71	0.80	0.77

It would seem reasonable to expect the cost of the portfolio to be linear in the portfolio weights, but one would generally predict the relationship between the financing strategy and risk to be inherently non-linear. Note that this does imply that the risk measures are *not* entirely linear in the financing strategy, but they are nevertheless relatively well described by a linear model. This is worth considering in more detail. It essentially suggests that the benefits of diversification are relatively modest. What does this mean? In the extreme, it means that if we have two debt instruments, a and b, then the standard deviation of their annual debt charges for a portfolio, p, constructed from these two instruments is given by,

$$\sigma_p = \alpha\sigma_a + \beta\sigma_b. \tag{45}$$

In general, however, the portfolio standard deviation is summarized by,

$$\sigma_p = \sqrt{\omega_a^2\sigma_a^2 + \sigma_b^2\sigma_b^2 + 2\omega_a\omega_b\sigma_a\sigma_b\rho_{ab}}. \tag{46}$$

When are equations (45) and (46) equivalent? If we set them equal, expand, and equate the coefficients, we find that this holds only if the correlation between debt instruments a and b is equal to unity. In the perfect correlation situation, therefore, diversification is not possible. Clearly, this is *not* entirely the case in our situation, but our analysis does suggest that the benefits of diversification are lower than we might have initially anticipated.

Table 18: **Estimated Par-Rate Correlation**: This table summarizes the estimated correlation of the par-interest rates arising from our two-factor CIR model estimated using Canadian term-structure data from 1994 to 2001. The results are estimated from 2,500 randomly generated, 40-period sample paths.

Par interest-rate pair	Estimated correlation	Standard deviation
3 mos. - 30 yrs.	0.65	0.21
3 mos. - 10 yrs.	0.73	0.17
3 mos. - 5 yrs.	0.84	0.11
3 mos. - 2 yrs.	0.96	0.03
3 mos. - 1 yr.	0.97	0.01
3 mos. - 6 mos.	0.98	0.00

What is driving this high degree of correlation? One possible reason is that our interest-rate model is parameterized such that all interest rates are perfectly correlated. This is, in fact, exactly the case in a single-factor model of the term structure of interest rates. In our two-factor setting, we have one factor that describes parallel movements in the term structure, the perfect correlation case, and another that describes changes in the steepness of the term structure over time. The introduction of the second factor eliminates the perfect correlation between zero-coupon interest rates. Table 18 highlights the pairwise correlation between the 3-month interest rate and the other key nodes on the term structure. Note that this analysis was performed with par interest rates, rather than with zero-coupon interest rates, because these are the rates at which the debt is actually issued; as par rates are essentially coupon-weighted averages of zero-coupon rates, they will exhibit a higher relative correlation. Nonetheless, actual interest rates exhibit quite a substantial correlation between adjacent values, but less when we consider non-adjacent values on the term structure (i.e., 3 months versus 10 years). This highlights the fact that we should be exploring other possible term-structure models that permit less correlation between nodes on the term structure.[50]

[50] There is a class of *jump-diffusion* models, for example, that may be worth considering.

Another possible reason for the modest diversification benefits is the importance of the initial portfolio. When the initial portfolio is large, the importance of newly issued debt in the overall debt charges is dampened. In fact, the debt-charge distribution should be overwhelmingly driven by the large amount of previously issued debt that has a fixed and known cost. This is particularly evident in the early years of the simulation. Some rough analysis indicated that the initial portfolio is not terribly important in the results. The main reason is the influence of treasury-bill borrowing, which by definition is refinanced within one year. Any financing strategy with a non-trivial amount of short-term debt will have the corresponding amount of refinancing.

The final, and most plausible, reason for the relatively high degree of linearity between risk and financing strategy relates to the nature of debt issuance in any given financing strategy. As previously described, our analysis employs a stock rule for the application of the vector of portfolio weights associated with each financing strategy. The idea is to issue the appropriate amount of debt instruments to maintain the overall portfolio in balance, as described by the set of portfolio weights. This implies that we need to be careful about the interpretation of the portfolio weights. Consider the simple example of a portfolio of two assets composed equally of 3-month treasury bills and 30-year bonds. This does *not* imply that each quarter the government will issue 50 per cent of its financing requirement in 3-month treasury bills and the other 50 per cent of its financing requirement in 30-year bonds[51]: the treasury bills need to be completely refinanced each quarter, whereas the 30-year bonds are quite infrequently refinanced. In any given period, then, there is much more issuance at the short end of the term structure than at the long end. Table 19 explains this phenomona by summarizing the average quarterly issuance for a $400 billion initial portfolio with four reopenings per bond and equal weights in each of the usual seven debt instruments.[52]

Table 19: **Average Quarterly Issuance**: This table outlines the average quarterly issuance, in $ billions, of a $400 billion initial portfolio with equal portfolio weights.

3 mos.	6 mos.	1 yr.	2 yrs.	5 yrs.	10 yrs.	30 yrs.
72.8	32.0	16.0	8.0	3.2	1.6	0.8

Observe that the quarterly issuance in the 3-month sector is more than 90 times larger than the average quarterly issuance in the 30-year sector, and more than 45 times greater than the average quarterly 10-year issuance. The majority of the issuance occurs in the short end of the term structure, where the correlation between instruments approaches unity. The standard formula for standard deviation, as outlined in equation (46), is thus not terribly useful in a dynamic setting, because it does not take into account

[51] In fact, for each stock rule there exists an equivalent flow rule that achieves the same result. This would involve altering the proportions of issuance to take into account their maturity.

[52] For the two-year bond, there are only two reopenings.

the concept of issuance. The benefits of diversification are, therefore, less important than might have been expected. This is clearly a strong conclusion that suggests a sovereign's debt strategy decision is driven more by relative risk preferences and concern about extreme events than attempts at diversification.

We have argued that the benefits of diversification are less substantial than one might have initially suspected. Let us now look more closely at these trade-offs in graphical form to see whether this is a reasonable conclusion. Figure 10 illustrates the two-dimensional cost-risk trade-off for four different measures of risk: one-year RCaR, one-year RTCaR, one-year cost volatility, and conditional cost volatility (ξ). We characterize the first three risk measures as short term, given their one-year focus. Conversely, the final measure, conditional cost volatility, is a longer-term measure of the risk associated with a given financing strategy.

There appear to be two separate clouds of points on each graph. This *discontinuity* stems from our model of the relationship between financing strategy and issue cost.[53] Financing strategies that involve highly concentrated or very limited issuance in a few debt instruments will tend to be more expensive. Thus, these *non-diversified* strategies will be inefficient for a given level of risk. An important diversification effect, therefore, in our analysis stems from the interplay between the financing strategy and issuance costs.

What is meant by this diversification effect? Long-term dominated financing strategies demonstrate high cost but low risk; these are the financing strategies represented in the bottom right-hand corner of each graph in Figure 10. Those financing strategies in the top left-hand corner, or the shorter-term based financing strategies, involve lower costs at the expense of higher risk. The interior points are financing strategies that involve combinations of long- and short-term debt instruments. Inspection of Figure 10 reveals evidence of a diversification effect. That is, financing strategies using short-, long-, and medium-term debt instruments appear *both* less risky and costly than a simple linear combination of short-term and long-term financing strategies would imply.

The fact that this relationship is reasonably linear between all four measures of risk and cost is also quite evident in Figure 10. This underscores the relatively limited nature of the diversification benefits. The diversification benefit associated with the feedback between the financing strategy and debt issuance costs appears to be more important than the diversification effect derived from the correlation structure of the term structure of interest rates. Moreover, once a government decides on its relative preferences for short- and long-term risk, a relatively small number of financing strategies should be considered.

At this point, while considering Figure 10, a natural question is what is the *best* financing strategy (i.e., ω_k) for the government. One might be tempted, given its widespread popularity, to use the usual Markowitz

[53] It is discontinuous because we use a discontinuous function to mathematically model this relationship.

Figure 10: **Dispersion Risk and Cost Tradeoffs**: This figure includes three short-term risk-cost trade-offs; one-year relative cost-at-risk, relative tail cost-at-risk, and one-year cost volatility. It also includes a long-term measure of risk in the form of conditional cost volatility. Cost is measured by the first-year mean cost for the first three measures and unconditional mean debt cost.

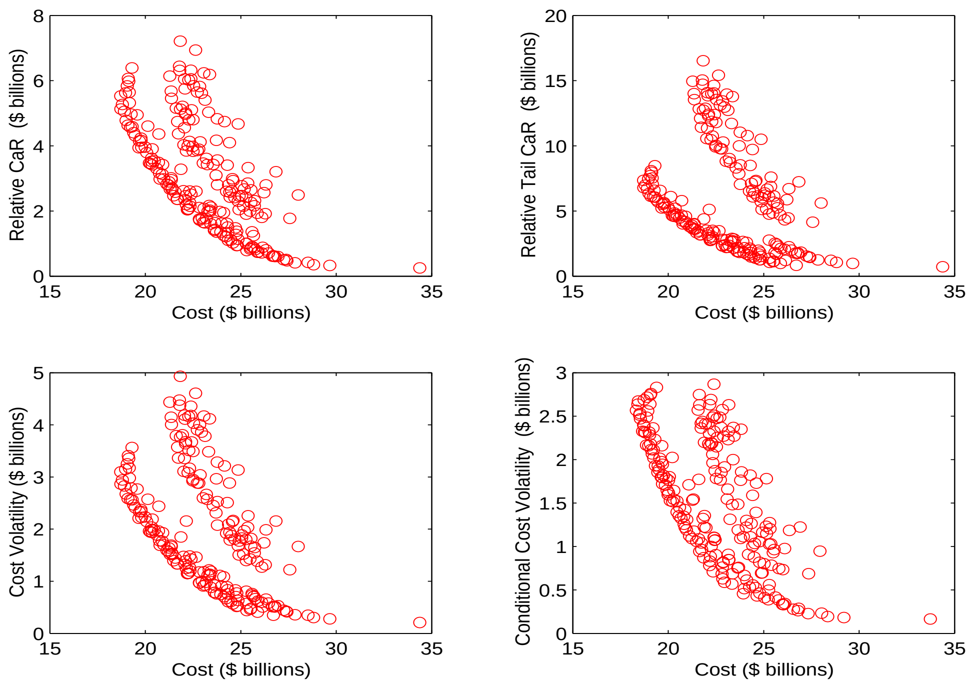

efficient frontier approach that is described as the following straightforward optimization problem,

$$\begin{aligned} \min_{\omega_k} \quad & \frac{\omega_k^T \Sigma \omega_k}{2}, \\ \text{subject to:} \quad & \omega_k^T \vec{1} = 1, \\ & \omega_k^T \vec{\mu} = \mathbb{E}(\mu), \\ & \omega_{kh} \in [0,1] \text{ for all } h \in \{1, ..., H\}, \end{aligned} \tag{47}$$

where the costs of the vector of seven debt instruments are $\mathcal{N}(\vec{\mu}, \Sigma)$, $\vec{1}$ is a vector of ones, and $\mathbb{E}(\mu)$ is the desired expected cost of the portfolio. In other words, the usual way to deal with this type of setting is to minimize the variance of cost subject to a given level of expected cost (i.e., $\mathbb{E}(\mu)$). There are, at least, three reasons why this is not applicable in our setting. First, we make no assumption of normality of our debt costs. This is a key assumption that permits one to examine only the first two moments of the debt-charge distribution. Second, this is a single-period model and we are faced with a multiperiod problem. Third, and perhaps most importantly, a generalization of this approach does not permit us to consider the inherent path dependency in the debt strategy problem. That is, the set of outcomes in period n depend importantly on the outcomes in period $n-1$.

We can, of course, solve equation (47) for the first period and examine the results. The consequent *efficient frontier* is outlined in Figure 11 along first-year expected debt costs and associated first-year debt cost volatility for the 225 financing strategies under consideration. Observe that it does appear to capture the most desirable locus of risk and cost points. This approach is of limited usefulness, however, because it cannot be generalized into a multiperiod setting and assumes a Gaussian debt cost distribution. We are not, to be more precise, particularly interested in the best portfolio based on one-year debt charge volatility, but rather we are interested in determining the best financing strategy based on conditional cost measures over a long-term time horizon. Clearly, this is more complex. As we saw in section 2.2, this problem requires a substantially more complicated structure than that described in equation (47); in fact, one needs to use techniques from dynamic programming to solve this issue. As stated earlier, we do not attempt in this paper to address the concept of optimality. Instead, we are trying to understand the risk and cost characteristics of a number of different financing strategies.

Figure 11: **An Efficient Frontier**: This figure outlines the first-year expected debt costs and associated first-year debt cost volatility. Solving equation (47), we then determine an efficient frontier.

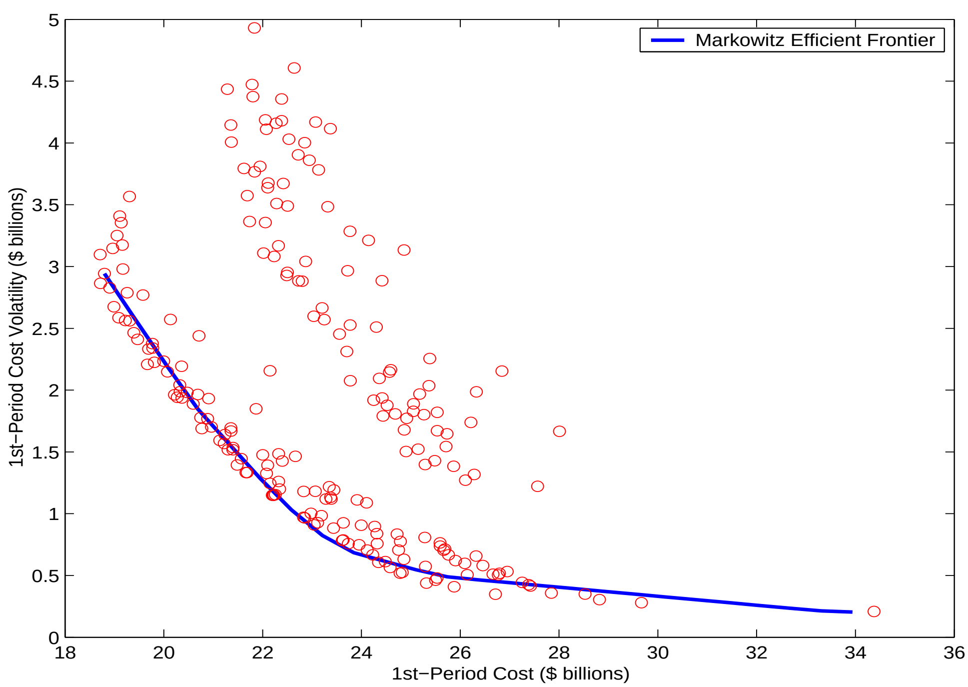

The four measures in Figure 10 address the risk associated with dispersed realizations of our stochastic environment. As we discussed in section 2, however, there is another element of risk in the debt strategy problem; we termed this *time diversification*, or refinancing, risk. In the previous analysis of this section, given the relatively small number of financing strategies, it was not appropriate to consider these measures. In this setting, we need a general approach to quantitatively describe this type of risk. Figure 12 outlines scatterplots for two different potential measures of time diversification risk: the fixed-debt ratio and the average quarterly debt redemption (AQR).

Figure 12: **Time-Diversification Risk and Cost Trade-offs**: This figure includes four scatterplots that attempt to describe time diversification (or refinancing) risk. These include the floating debt ratio and average quarterly debt redemptions (AQR) as compared with both the unconditional cost of the portfolio and the conditional volatility.

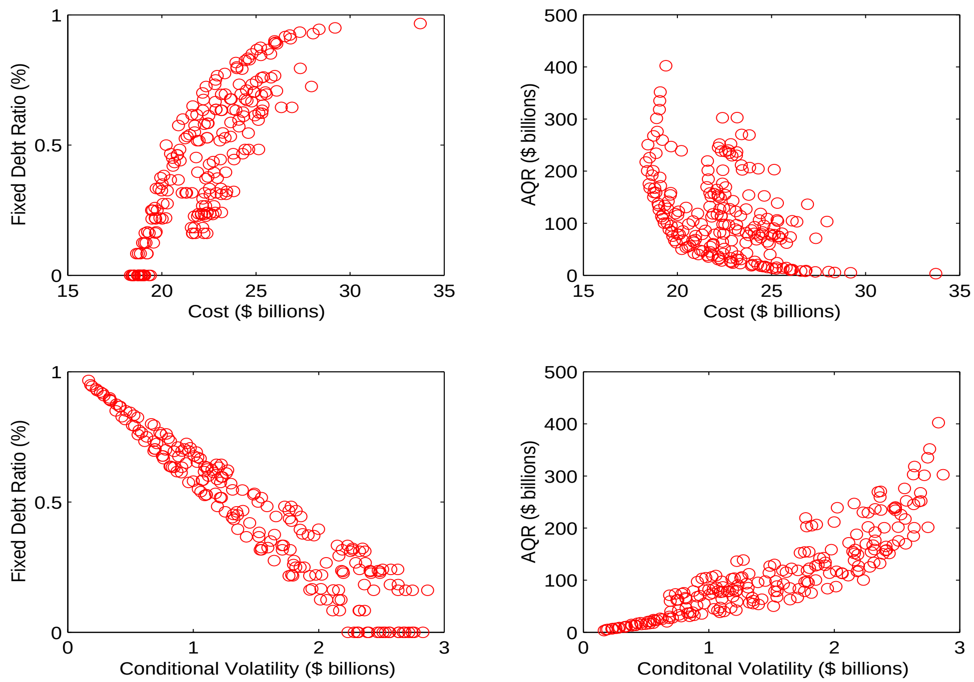

The first two quadrants of Figure 12 outline the relationship between our two time-diversification measures and the unconditional cost of the portfolio. We need to be somewhat cautious in interpreting these measures because the amount of refinancing risk *decreases* as the fixed-debt ratio increases, but *increases* as the AQR measure increases. We can make two broad conclusions from Figure 12. First, as refinancing exposure increases, the expected cost of the portfolio decreases; that is, there is a negative relationship between refinancing risk and debt cost. The second observation is a positive relationship between refinancing risk and dispersion risk. More specifically, as the refinancing risk of the portfolio increases, there is a corresponding increase in dispersion risk.

It appears that the fixed-debt ratio is somewhat more useful than the AQR measure, because a large number of different financing strategies appear to have very similar AQR values. Consider, for example, the fact that more than one half of the portfolios have average quarterly redemptions of less than $100 billion. With the fixed-debt ratio, the values are approximately equally distributed from zero to 100 per cent. We would suggest that the fixed-debt ratio is a more efficient measure of refinancing risk relative to average quarterly debt redemptions.

These relationships are identical to those uncovered in a general review of the financing strategies. In fact, it appears that refinancing risk is actually a proxy for the basic composition of the portfolio associated with a given financing strategy. A large amount of refinancing risk, for example, essentially describes a financing strategy dominated by short-term debt. Indeed, one might argue that, given the financing strategy, a measure of refinancing risk is redundant. We would nevertheless argue that the ease of interpretation of fixed-debt ratio—relative to a financing strategy—makes these measures useful for a sovereign debt manager. It effectively acts as a constraint to the problem. A government, for example, could place limits on the amount of dispersion risk and refinancing risk it would undertake. This would provide a prescriptive direction for the sovereign debt manager, in that certain financing strategies would no longer be permissible.

Figure 13: **A Long-Term Cost, Risk, Time-Diversification Surface**: This figure illustrates the three-dimensional interplay between unconditional average debt costs, conditional volatility, and the fixed-debt ratio.

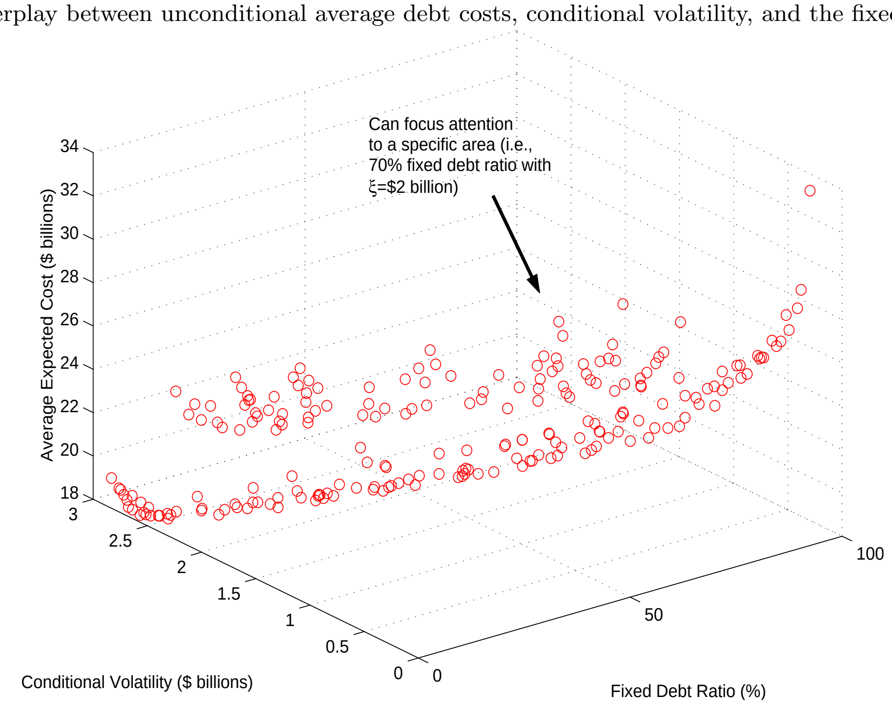

Figure 13 describes the cost, dispersion risk, and time-diversification risk in three-dimensional space.

This surface—and some decisions on risk preferences—essentially allows us to narrow down a large number of financing strategies into a smaller preferred set. We could, for example, consider a target of 70 per cent fixed-debt ratio and a conditional volatility of $2 billion. Inside the cloud of points that satisfy these two criteria, of course, there are only a few associated financing strategies that are the most efficient in terms of cost.

While Figure 13 narrows the financing strategy decision from a long-term perspective, one could repeat the analysis for a shorter-term perspective using the first-year mean debt cost, the first-year relative tail cost-at-risk, and the fixed-debt ratio. This type of analysis would permit the sovereign debt manager to reconcile the long- and short-term perspective when selecting a financing strategy. Clearly, in practice, this would be more difficult. It would be quite complicated, for example, if the long-term strategy—in steady state—differed substantially from the current portfolio. One would then have to explicitly consider the transition strategy. Moreover, the analysis in this paper does not help in determining a transition strategy.

Figure 14: **A Short-Term Cost, Risk, Time-Diversification Surface**: This figure illustrates the identical three-dimensional trade-off, as in Figure 13, from a short-term perspective. We consider the first-year mean debt cost, the first-year RTCaR measure, and the fixed-debt ratio.

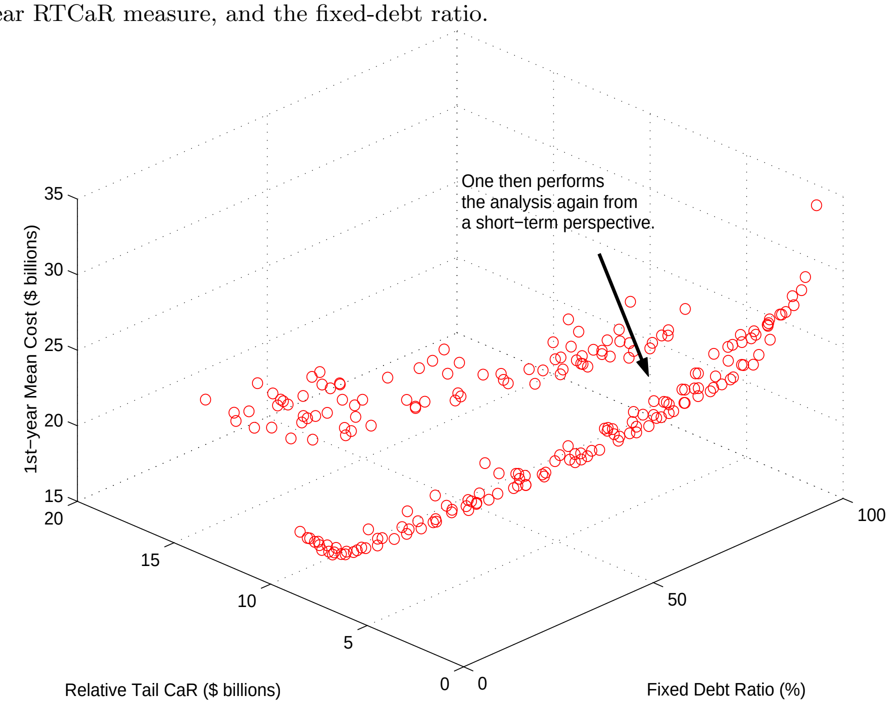

There is an additional issue to be considered. We have demonstrated in Figures 13 and 14 that there is no unique financing strategy that will meet a given set of government time diversification and dispersion risk preferences. A large part of the reason for this is the fact that many relatively distinct financing strategies share the same fixed-debt ratio. In theory, a strategy of 50 per cent treasury bills and 50 per cent 30-year bonds will have a very similar fixed-debt strategy to an alternative strategy of 50 per cent treasury bills and 50 per cent 2-year bonds. While the differences in these portfolios become quite evident in terms of both cost and risk, the fact does remain that there is additional room for another approach to describing the financing strategy.

Ideally, one could simply use the actual financing strategy. This is, however, somewhat cumbersome. Another alternative, used by a number of different sovereigns, is to use a summary portfolio measure to describe the specific financing strategy in more detail. What we are suggesting here is that any single measure of a given financing strategy is insufficient to describe its risk and cost characteristics. We suggest a multifaceted approach with long-term cost and dispersion risk measures, a measure of refinancing risk, and an additional higher-level portfolio measure. Two possible portfolio measures include the average term to maturity (ATM) and the MacCauley duration of the portfolio. The relationship between each of these measures, the unconditional debt cost, and the conditional volatility is summarized in Figure 15.

Figure 15: **Possible Portfolio Measures**: This figure illustrates two potential portfolio measures: average term to maturity and MacCauley duration. The relationship between these two measures and both the unconditional debt cost and the conditional volatility are summarized.

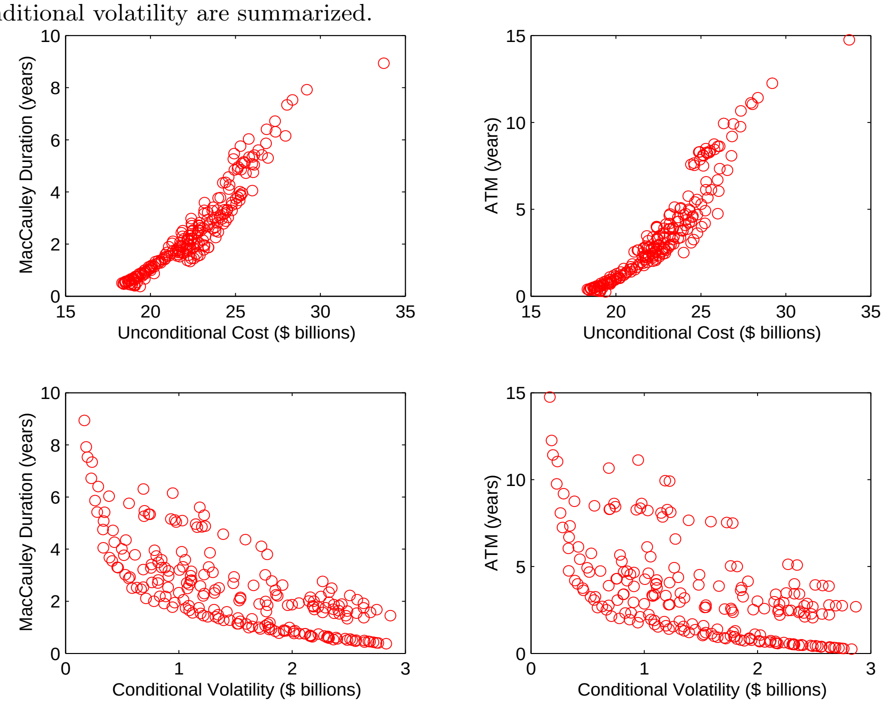

There does not appear to be much difference between the ATM and MacCauley duration measures. This should hardly be surprising, given the fact that the correlation between these two measures approaches unity.

This would suggest that, to decide between these two portfolio measures, we need to consider additional factors. The key distinction between these two approaches is that the MacCauley duration attempts to consider the impact of coupon payments on the overall term to maturity of a given fixed-income security. The ATM, however, is a simple weighted average of the various terms to maturity of the debt instruments in the portfolio. The MacCauley duration appears to be a more realistic measure, but it is also more difficult to interpret. In particular, assumptions about the term structure of interest rates are necessary for a computation of this measure; as a consequence, it can vary over time. We avoid this problem by using an *average* term structure for the computation of the MacCauley duration; this may not be a reasonable assumption. For the purposes of this study, we use the ATM as a high-level summary of the government's financing strategy.[54] This high-level measure is complemented by a long-term dispersion risk (conditional volatility, ξ), short-term dispersion risk (relative tail CaR, RTCaR), and a refinancing risk measure (fixed-debt ratio). We suggest that the use of the surfaces outlined in Figures 13 and 14 can help isolate the most efficient financing strategy, in terms of cost, for a given set of pre-specified risk preferences. The methodology is described graphically in Figure 16. To summarize, we suggest that the debt manager examine dispersion risk, or the volatility of debt-service charges, from both a short-term and long-term perspective. In doing so, the debt manager can also explicitly consider time-diversification, or refinancing, risk. The key idea is that the debt manager should focus on multiple measures to describe the risk and cost characteristics of a given financing strategy, because the complicated nature of the problem requires examination from multiple perspectives.

The question still remains as to how one determines a target level for these various elements of our risk-cost framework. We have seen, for example, that short-term-debt dominated financing strategies are less costly but more risky than longer-term-debt focused financing strategies. One possible way to address this issue is to adjust the cost of each individual financing strategy for the general level of risk. In other words, we could normalize the cost of each financing strategy for its underlying level of uncertainty. In portfolio theory, it is common to use the *Sharpe ratio*, defined as,

$$\frac{\mu - r}{\sigma}, \tag{48}$$

where a given return of an asset, $a \sim \mathcal{N}\left(\mu, \sigma^2\right)$, and r represents a risk-free interest rate. This will not be particularly useful for a debt manager. First of all, we do not have an obvious candidate for the risk-free rate, because the risk-free rate is itself typically a government rate. More importantly, in an asset-based approach, high expected returns are associated with high variance. Normalization by the standard deviation reduces these returns somewhat. In our liability setting, high expected costs are associated with *low* variance. Normalization by standard deviation, or some other measure of risk, would have the opposite effect.

[54]There are some practical limitations with the ATM measure as well. In particular, it is quite sensitive to the amount of long-term issuance. This type of sensitivity may not be a desirable quality in a high-level summary measure.

Figure 16: **Risk-Cost Framework**: This figure outlines a recommended approach to using our stochastic simulation framework to select a specific financing strategy.

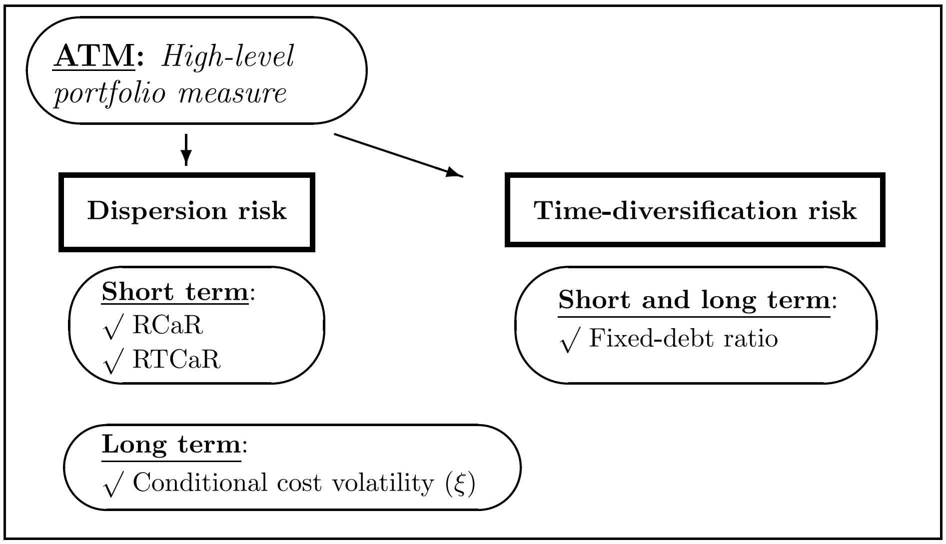

We suggest instead that the risk of the portfolio be added to its expected cost. This measure of adjusted cost (AC) is given by,

$$\text{AC}_\theta = \mathbb{E}_\theta(c_t) + \nu_{\theta,t}, \tag{49}$$

where $\nu_{\theta,t}$ represents a measure of risk for financing strategy, θ, at time t. The idea behind this is quite simple; if the expected cost of a portfolio is high, it will exhibit low risk and the adjusted cost will be very similar to the unadjusted cost. If, conversely, the expected cost of a portfolio is quite high, it generally has high risk and thus there will be a large upwards adjustment in the cost for that financing strategy. This will place all the portfolios on an equal footing.

Before we actually examine these adjusted costs, let us alter equation (49), but putting these costs into a percentage of the overall portfolio. That is, we define the relative adjusted cost (RAC) as,

$$\text{RAC}_\theta = \frac{\mathbb{E}_\theta(c_t) + \nu_{\theta,t}}{\text{card}(\Xi)}, \tag{50}$$

where card(Ξ) represents the size of the steady-state portfolio. Figure 17 summarizes the risk-adjusted expected debt costs for our 225 financing strategies relative to their ATM.[55] We have adjusted costs based

[55] In the case of RCaR and RTCaR, the definition of equation (50) is equivalent to the CaR and TCaR measures defined in equations (16) and (18), respectively.

Figure 17: **Risk-Adjusted Expected Cost**: This figure summarizes the risk-adjusted expected debt costs for our 225 financing strategies relative to their ATM. We have adjusted costs based on four different measures of risk: conditional volatility (ξ), one-year RCaR, one-year RTCaR, and one-year cost volatility.

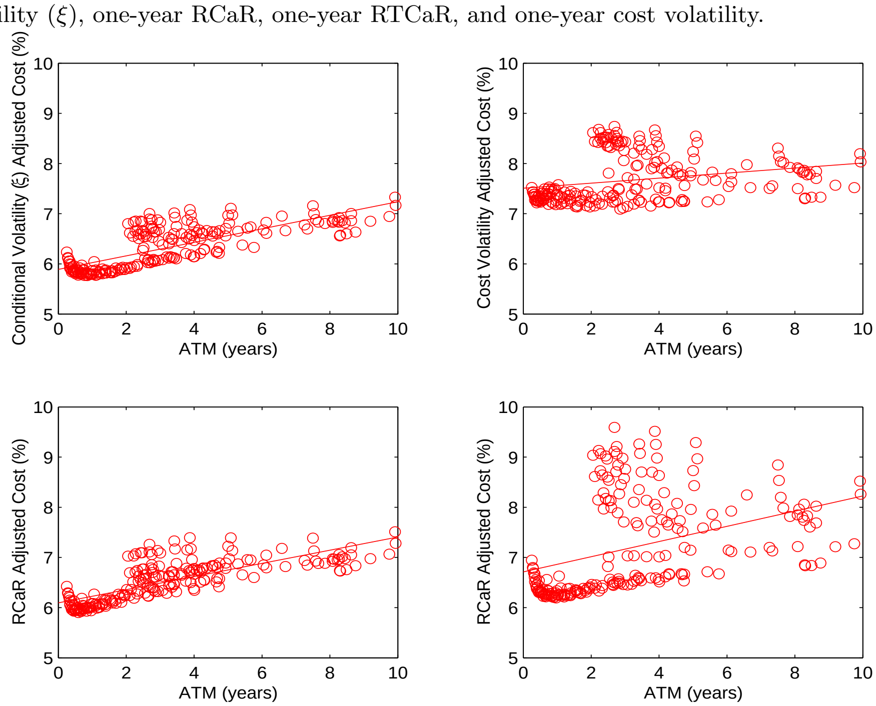

on four different measures of risk: conditional volatility (ξ), one-year RCaR, one-year RTCaR, and one-year cost volatility.

One might be tempted to conclude, from examining Figure 17, that across all four risk-adjusted cost measures the shorter ATM financing strategies are somehow *preferable* in risk-adjusted cost terms as compared with longer-term financing strategies. That is, one could argue that there is a bias towards those financing strategies that reduce the ATM of the overall debt portfolio. Caution, however, should be exercised, for at least three reasons. First, we observe that risk-adjusted cost appears to spike upwards at very short-term ATM levels. This underscores the riskiness of these strategies. This leads to the second point, that these results were performed under normal market condition assumptions for our stochastic environment. In section 4.4, we saw that the shorter-term-debt dominated financing strategies were more seriously impacted in both cost and risk terms by the introduction of a low-probability extreme state. The effect, therefore, of potentially extreme outcomes would be to flatten, or even invert, this relationship. This is underscored by the fact that when we adjust for tail risk, the trend is somewhat weaker than when we consider the variance-based measures of risk. Finally, a decision rule based entirely on this risk-adjusted cost analysis has nothing to say about the risk preferences of the government. The government could be either risk-averse or risk-seeking, and the risk-adjusted cost would continue to provide the same result. Therefore, we would suggest that the concept of risk-adjusted cost is a useful tool that, when considered in conjunction with the other measures described in this text, can assist in the decision-making process.

5 Conclusion

We have covered a tremendous amount of ground. In section 2, we defined the debt strategy problem in both qualitative and quantitative terms. We then worked through the conceptual and practical details of our stochastic simulation framework in section 3. Combining the details of this discussion, we turned to examine a set of illustrative results in section 4. This included an investigation of the new features of our stochastic framework, a sensitivity analysis, a novel approach to stress testing, and the examination of a relatively large number of financing strategies. This final analysis, in section 4.5, helped us develop a general approach to using the stochastic simulation framework to select a given financing strategy predicated on the consideration of a range of different measures of cost and risk. What have we learned from this analysis? We summarize the key conclusions as follows:

- The use of the full stochastic environment, including a model of the macroeconomy and the government's financial position, appears to be an improvement on the sole consideration of interest-rate dynamics. The reliance on the simple stochastic environment tends to underestimate *both* the risk and cost associated with a given financing strategy.

- The sensitivity analysis in section 4.3 revealed that assumptions regarding the random evolution of the term structure of interest rates are a critical component in the overall analysis. Indeed, the term-structure model appears to be significantly more important than the government's financial requirement process.

- Stress testing is an essential part of any complete debt strategy analysis. Even minute probabilities of extreme future interest-rate realizations contribute to non-trivial increases in both the expected cost and risk of a given financing strategy.

- The use of conventional risk measures, such as annual cost volatility and relative CaR, for longer-term analysis is not always helpful, because they fail to condition on the state of the world in the previous period. This is important because of the path-dependent nature of the debt strategy problem. We recommend the use of a measure that explicitly conditions on previous outcomes: the conditional cost volatility. This measure, which complements shorter-term measures such as RCaR and RTCaR, provides the flexibility for the consideration of cost and risk from *both* a short- and long-term perspective.

- The diversification benefits across different debt instruments are relatively modest. This is a result of two factors. First, our models produce a high level of correlation across all nodes of the term structure, with particularly high correlation at the short end—this is consistent with empirical term-structure behaviour. Second, there is a preponderance of short-term issuance across all financing strategies, which implies that most actual *issuance* occurs at highly correlated short-term rates. The consequence is a closer to linear relationship than expected between measures of risk and financing strategy. Ultimately, this implies that the most important elements of the government's choice of financing strategy are its risk preferences and the protection it affords against extreme events.

- Diversification appears through the explicit consideration of the interplay between issuance cost and financing strategy. Limited or concentrated issuance in a given debt instrument will precipitate increased cost to the government. This implies that diversification among various debt instruments will contribute to lower debt charges for a given level of risk.

- We have identified the need for additional measures to describe in more detail the nature of a given financing strategy. In particular, we have shown the fixed-debt ratio to be a reasonable measure of time-diversification risk and highlighted the usefulness of the average term to maturity as a high-level portfolio summary measure.

- We have constructed a risk-adjusted measure of cost to permit us to determine a potential target for the overall portfolio ATM. Our analysis suggests that a slight bias towards a shorter ATM is evident, although this is somewhat tempered by a number of factors, including the potential for extreme events to impact short-term strategies in a disproportionate manner.

This work is *not* the last word on the debt strategy problem. More work is required in a variety of directions. We suggest *four* possible avenues for future analysis; there are surely a number of others. First, given the sensitivity of the results to our term-structure model, greater effort is required to consider more flexible models to describe the stochastic interest-rate environment. Second, we saw that the conditional cost distributions are not normally distributed. It would be useful to consider more flexible parametric or non-parametric approaches towards characterizing conditional debt-cost volatility. It is also obvious that our overall stochastic environment is highly stylized. It does not, for example, consider inflation, thereby precluding consideration of index-linked debt in our analysis. Moreover, ignoring inflation does not permit us to examine the impact of *supply* relative to *demand* shocks on the macroeconomy. This issue is critical to understanding the issues surrounding the covariance between the state of the economy, monetary policy rules, and the primary balance. Future work is required to enhance the richness of our stochastic environment while simultaneously keeping computational effort under control. Third, we suggest a more detailed investigation of the relative advantages and disadvantages of the MacCauley duration and ATM as high-level portfolio measures. Fourth, it would be extremely helpful to consider numerical techniques that might help us determine—given some assumptions about government risk preferences—an optimal financing strategy. There are a number of dynamic programming techniques that could, with some creativity and a few simplifying assumptions, perhaps be profitably be applied to this problem.

Bibliography

Abken, P. A. (2000): "An Empirical Evaluation of Value at Risk by Scenario Simulation," *Journal of Derivatives*, 7, 12–18.

Aoki, M. (1989): *Optimization of Stochastic Systems*. Academic Press, 1250 Sixth Avenue, San Diego, CA, second edn.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath (1999): "Coherent Measures of Risk," *Mathematical Finance*, 9(3), 203–228.

Aström, K. J. (1970): *Introduction to Stochastic Control Theory*. Academic Press, 111 Fifth Avenue, New York, New York, first edn.

Auditor General of Canada (2000): "Managing Canada's Debt: Facing New Challenges," Discussion paper.

Barro, R. J. (1995): "Optimal Debt Management," National Bureau of Economic Research: Working Paper No. 5327.

Batley, R. (2002): "A Long Run View to Optimal Debt Issuance Strategy," Unclassified Internal UK Debt Management Office Document.

Bennett, P., K. Garbade, and J. Kambhu (2000): "Enhancing the Liquidity of U.S. Treasury Securities in an Era of Surpluses," *FRBNY Economic Policy Review*, 3, 89–105.

Benveniste, A., M. Métivier, and P. Priouret (1997): *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, New York, New York, first edn.

Bergström, P., and A. Holmlund (2000): "A Simulation Model Framework for Government Debt Analysis," Riksgälds Kontoret: The Swedish National Debt Office.

Billingsley, P. (1995): *Probability and Measure*. Wiley, Third Avenue, New York, New York, third edn.

Black, R., and C. Telmer (1999): "Liability Management Using Dynamic Portfolio Strategies," *Algo Research Quarterly*, 2, 9–19.

Bolder, D. J. (2001): "Affine Term-Structure Models: Theory and Implementation," Bank of Canada: Working Paper No. 2001-15.

——— (2002): "Towards a More Complete Debt Strategy Simulation Framework," Bank of Canada: Working Paper No. 2002-13.

Bolder, D. J., and S. Gusba (2002): "Exponentials, Polynomials, and Fourier Series: More Yield Curve Modelling at the Bank of Canada," Bank of Canada: Working Paper No. 2002-29.

Committee on the Global Financial System (2000): "How Should we Design Deep and Liquid Markets: The Case of Government Securities," Discussion paper, Bank for International Settlements.

Cox, J. C., J. E. Ingersoll, and S. A. Ross (1985a): "An Intertemporal General Equilibrium Model of Asset Prices," *Econometrica*, 53, 363–384.

——— (1985b): "A Theory of Term Structure of Interest Rates," *Econometrica*, 53, 385–407.

Cumberworth, M., A. Hitchcox, W. McConnell, and A. Smith (1999): "Corporate Decisions in General Insurance: Beyond the Frontier," Institute of Actuaries.

Danish Nationalbank (1998): "Management of Interest Rate and Refinancing Risk: Cost-at-Risk," Discussion paper, Working Party on Debt Management: Organization for Economic Co-operation and Development.

Danmarks Nationalbank (2000): "Danish Government Borrowing and Debt," Discussion paper.

Danmarks Nationalbank (2001): "Danish Government Borrowing and Debt," Discussion paper.

DeGrauwe, P. (1989): *International Money: Post-War Trends and Theories*. Oxford University Press, Walton Street, Oxford.

Department of Finance Canada (2001): "Debt Management Strategy: 2001-2002," Discussion paper.

Duffie, D., and R. Kan (1996): "A Yield-Factor Model of Interest Rates," *Mathematical Finance*, 6, 379–406.

Durrett, R. (1996): *Probability: Theory and Examples*. Duxbury Press, Belmont, California, second edn.

Dynamic Financial Analysis Committee of the Casual Actuary Society (1999): "Overview of Dynamic Financial Analysis," Discussion paper.

Ermoliev, Y., and R. Wets (1988): *Numerical Techniques for Stochastic Approximation*. Springer-Verlag, New York, New York, first edn.

Fishman, G. S. (1995): *Monte Carlo: Concepts, Algorithms, and Applications*. Springer-Verlag, 175 Fifth Avenue, New York, NY, Springer Series in Operations Research.

Freedman, C. (1990): "Why Should the Government Aim at Reducing the Variance of Interest Payments," Bank of Canada.

Gibson, M. S., and M. Pritsker (2000): "Improving Grid-based Methods for Estimating Value at Risk of Fixed Income Portfolios," Division of Research and Statistics, Federal Reserve Board.

Gilbert, P. (2001): "Debt Management Strategy: 2001-2002," Discussion paper.

Graveline, J. J. (2000): "A Survey of Topics in Optimal Consumption and Investment," University of Waterloo.

Gray, S. (1996): "The Management of Government Debt," Bank of England.

Hamilton, J. D. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–384.

Hermanutz, D., and C. Matier (2001): "Modelling Federal Finances under Uncertainty," Department of Finance Canada.

HM Treasury (UK) (2002): "Debt and Reserves Management Report 2002-03," Discussion paper.

Holmlund, A. (1999): "The Debt Office's Model for Analyzing Duration Choice for the Swedish Kronor Debt," Riksgälds Kontoret: The Swedish National Debt Office.

Holmlund, A., and S. Lindberg (2002): "The SNDO's Simulation Model for Government Debt Analysis (Preliminary Draft)," Riksgälds Kontoret: The Swedish National Debt Office.

Hörngren, L. (1999): "Methods for Analyzing the Structure of the Central Government Debt," Riksgälds Kontoret: The Swedish National Debt Office.

Hostland, D. (2001a): "Monetary Policy and Medium-Term Fiscal Planning: Preliminary Version," Department of Finance Canada.

——— (2001b): "Specification of a Stochastic Simulation Model for the Analysis of Monetary and Fiscal Policy," Department of Finance Canada.

Hostland, D., and C. Matier (2001): "An Examination of Alternative Strategies for Reducing Public Debt in the Presence of Uncertainty," Department of Finance Canada.

IMF and the World Bank (2001): "Guidelines for Public Debt Management," Discussion paper, Prepared by the Staffs of the International Monetary Fund and the World Bank.

Judge, G. G., W. Griffiths, R. C. Hill, H. Lütkepohl, and T.-C. Lee (1985): *The Theory and Practice of Econometrics*. John Wiley and Sons, New York, New York, second edn.

Karatzas, I., and S. E. Shreve (1991): *Brownian Motion and Stochastic Calculus*. Springer-Verlag, Berlin, second edn.

——— (1998): *Methods of Mathematical Finance*. Springer-Verlag, Berlin.

Kaufmann, R., A. Gadmer, and R. Klett (2001): "Introduction to Dynamic Financial Analysis," Risklab Switzerland Research Paper Series.

Lèonard, D., and N. V. Long (1992): *Optimal Control Theory and Static Optimization in Economics*. Cambridge University Press, 40 West 20th Street, New York, New York, first edn.

Longstaff, F. A., and E. S. Schwartz (1992a): "Interest Rate Volatility and the Term Structure: A Two-Factor General Equilibrium Model," *The Journal of Finance*, XLVII, 1259–1282.

——— (1992b): "A Two-Factor Interest Rate Model and Contingent Claims Evaluation," *The Journal of Fixed Income*, pp. 16–23.

McLeish, D. (2002): *Monte Carlo Methods in Finance*. Forthcoming.

Missale, A. (1994): "Public Debt Management," PhD Dissertation: Massachusetts Institute of Technology.

Robbins, H., and S. Monro (1952): "A Stochastic Approximation Model," *Annals of Mathematical Statistics*, 22, 400–407.

Swedish Ministry of Finance (Regeringskansliet) (2001): "Guidelines for Central Government Debt Management: Decision taken at the Cabinet Meeting 8 November 2001," Discussion paper.

Swedish National Debt Office (Riksgälds Kontoret) (2001): "Central Government Debt Management - Proposed Guidelines," Discussion paper.

Swedish National Debt Office (Riksgälds Kontoret) (2002): "Central Government Borrowing: Forecast and Analysis," Discussion paper.

United Kingdom Debt Management Office (2001): "Index Linked Gilt Re-design: Consultation Document," Discussion paper.

United Kingdom Debt Management Office (2002a): "Index Linked Gilt Re-design: Response to Consultation," Discussion paper.

United Kingdom Debt Management Office (2002b): "Exchequer Cash Management in the United Kingdom: A DMO Handbook," Discussion paper.

Vasicek, O. (1977): "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177–188.

Wang, T. (1999): "A Class of Dynamic Risk Measures," University of British Columbia.

Youngdahl, J., B. Stone, and H. Boesky (2000): "Implications of a Disappearing Treasury Debt Market," Goldman Sachs: Global Economics Paper No. 48.

Appendix A: The Control Problem

As discussed in section 2.2, our question, in its purest form, falls into the class of *stochastic optimal-control* problems. Because this area of mathematics may not be immediately familiar, the brief development that follows is intended to build some intuition for the non-technical reader. The easiest way to understand this problem, we believe, is to compare stochastic optimal control with a simpler set of problems and work upwards to our more detailed question. We can, for example, start by considering classical optimization. That is, in a classical optimization setting we are usually given some kind of real-valued, deterministic function of some number of variables, say for simplicity $h(x)$ where $x \in \mathbb{R}$. We are then asked to find the value of x that minimizes (or, conversely, maximizes) the value of h.[56] More formally, we wish to solve the following,

$$\min_{\mathrm{x}} \; h(x), \tag{51}$$

subject to,

$$x \leq \mu,$$

where $\mu \in \mathbb{R}$. Our success in this venture will, of course, depend on the characteristics of our function, h, and any restrictions that we might wish to impose on x or h. Nevertheless, if h is not too poorly behaved, this is a relatively simple calculus exercise. Indeed, this should be a very familiar problem. We can, however, add an additional level of complexity. In particular, we might wish to ask what would happen if the value of h depended instead upon time in some simple way. Say, for example, that time could take one of two possible values $t \in \{t_0, t_1\}$ and that our function, h, was time-separable. We could restate the problem as follows,

$$\min_{\mathrm{x}_{t_0}, \mathrm{x}_{t_1}} \; h_{t_0}\left(x_{t_0}\right) + h_{t_1}\left(x_{t_1}\right), \tag{52}$$

subject to,

$$x_{t_1} - x_{t_0} \leq \mu(t_1 - t_0),$$

where $\mu \in \mathbb{R}$. This is, in principle, identical to the problem posed in equation (51), except we are now solving for two parameters, x_{t_0} and x_{t_1}, as opposed to a single unknown, x. In addition, we have added a slight complication to our constraint. That is, the difference in our two parameter values will be related, in a linear

[56] Please note that the following discussion is very fast and loose. There are a variety of issues relating to the continuity, differentiability, and convexity of the functions in question that become critical; moreover, the nature and form of the constraints is of some importance. We are abstracting from these details, given the illustrative nature of this discussion and the fact that these technical details add little to the intuition we are trying to develop.

fashion, to the size of the time interval; it is standard to constrain the dynamics of our parameter values through time. Thus, we have generalized the basic problem. We can, in fact, continue this generalization by simply considering an arbitrarily large number of time periods. That is, we may generalize to $t \in \{t_0, t_1, ..., t_T\}$ as follows,

$$\min_{\mathrm{x_{t_0}, x_{t_1}, \ldots, x_{t_T}}} \sum_{i=0}^{T} h_{t_i}(x_{t_i}) \tag{53}$$

subject to,

$$x_{t_i} - x_{t_{i-1}} = \mu(t_i - t_{i-1}),$$

for $i = 0, 1, ..., T$ and where $\mu \in \mathbb{R}$. Essentially, instead of solving for two parameter values, we are instead solving an optimization problem for a vector of time-dependent parameters $\{x_{t_0}, x_{t_1}, ..., x_{t_T}\}$. Another way of thinking about this problem is that we seek a sequence of values that will control our function, h, in such a way that it attains its minimum value over $\{t_0, t_1, ..., t_T\}$. We have moved to the realm of a discrete-time optimal-control problem. To eliminate some of the notational burden, let us move from this discrete-time setting to continuous time by letting the number of discrete-time increments, which we consider over the interval $[0, T]$, tend to infinity. This change has a subtle but important impact on the form of our problem. Equation (53) will take the following form,

$$\min_{\mathrm{x(t)}} \int_0^T h(x(t))\,dt \tag{54}$$

subject to,

$$dx(t) = \mu dt,$$

for $t \in [0, T]$ and where,

$$\mu \in \mathbb{R},$$
$$dx(0) = \delta.$$

This is a problem of optimal control in a deterministic setting. Let us pause to make a few observations. First of all, note that the solution to this optimization is the unknown function, $x(t)$, rather than a sequence of values. Thus, we are asked to find an optimal function, or process, $x(t)$ such that h is minimized over some time interval, $[0, T]$. The second thing to note is that our condition, or constraint, on the dynamics of $x(t)$ has become, in continuous time, a *differential equation*. This is a very typical result. We are, essentially, trying to find a function $x(t)$ that will allow us to guide h to its minimum value over the time interval, $[0, T]$, and we are also given its dynamics in the form of a differential equation. Consider a practical example from engineering.

Let our function h represent the fuel usage of a satellite and $x(t)$ represent our control over its trajectory. In solving this problem, therefore, we may be trying to find the optimal flight path for our satellite—constrained by some physical limits on our controls as summarized by our differential equation—that minimizes our fuel consumption. This is a more complicated situation than that presented in equation (51) and, indeed, one requires techniques from an area of mathematics termed dynamic programming to find its solution.

Now, we take the critical step forward to the basic form of our problem. It will involve changing the characteristics of our differential equation, $dx(t)$. In particular, it asks what happens when the dynamics of our control process, as described by $dx(t)$, has a random component. That is, we define a probability space, $(\Omega, \mathcal{F}, \mathbb{P})$, and let

$$dx(t) = \mu dt + \sigma dW(t), \tag{55}$$

where,

$$\mu, \sigma \in \mathbb{R},$$

and $\{W(t), \mathcal{F}(t), t \geq 0\}$ is a standard, scalar *Wiener process*.[57] There are other ways to introduce randomness into the model but this is the most common and, in some ways, the cleanest approach. To be formal, $dx(t)$ has become a *stochastic differential equation*. This seriously changes the situation. Now, we have a function h that not only depends on time but also depends on a function, $x(t)$, whose dynamics are random. In this case, we cannot with certainty know all the possible values of h over the time interval $[0, T]$. Indeed, all we can say is what we expect, on average, to happen over the time interval. In the context of our satellite example, this is equivalent to saying that our controls are subject to some noise in the form of a measurement error or interference of some kind. The result is that we can, at best, talk about minimizing the expected value of this integral over the time period.[58] This is represented in the underlying manner,

$$\min_{\mathrm{x(t)}} \mathbb{E}\left[\int_0^T h\left(x(t)\right) dt \,\middle|\, \mathcal{F}_t\right] \tag{57}$$

subject to,

$$dx(t) = \mu dt + \sigma dW(t),$$

[57] One may think of this, in a heuristic sense, as being the continuous time equivalent of the discrete *random-walk* process.

[58] In a more formal setting, we would write this as the Lebesque integral of a slightly modified equation (54) with respect to the probability measure, $\mathbb{P}$. This is represented as follows,

$$\min_{\mathrm{x(t)}} \int_A \int_0^T h(x(t)) \, dt \, d\mathbb{P}, \text{ for all } A \in \mathcal{F}_t \tag{56}$$

for $t \in [0, T]$ and where,

$$\mu, \sigma \in \mathbb{R},$$
$$dx(0) = \delta$$

At last, we have a *stochastic* optimal-control problem. The task is to find the unknown function $x(t)$, which in a general setting may also be a function of time, that *steers* the function g to its highest *expected* value over the time interval in question. To a reader familiar with microeconomic theory, this is similar to economic agents facing various consumption bundles and attempting to maximize their expected utility over some time horizon. This is a more involved question than that posed in equation (51) and *dynamic programming* is the area of mathematics that deals with existence, uniqueness, and corresponding techniques for the solutions to this class of problems.

Appendix B: Model Estimation, Calibration, and Diagnostics

In this study, we rely on a number of parametric models to describe the evolution of the business cycle, the government's financial requirements, and the term structure of interest rates.[59] This raises two important questions. First, how does one determine the parameter for use in these models? The short answer is that the parameters are either estimated, using an econometric technique, or calibrated to the research analyst's assumptions about the future. Second, once these parameters are determined, how does the research analyst ensure that the parameters are actually consistent with both historical data and their assumptions regarding future events? The short answer to this second, important question is to simulate the data, given the parameter set, and examine its properties.

In this appendix, we will briefly describe the basic approach to dealing with these two issues. We begin with the first step, which involves estimation and calibration. It is performed in a sequence of four steps.

Step 1 - The Term Structure The estimation of the two-factor CIR model, used to describe term-structure dynamics, is performed using a Kalman filter estimation procedure. The technical details of this approach are discussed in substantial detail in Bolder (2001). In the full stochastic model, we extend the two-factor CIR model to permit the steepness of the term structure of interest to evolve over time in a manner that is consistent with the business cycle. Loosely speaking, this is achieved by letting one of the parameters of the term-structure model, termed the market price of risk of the first state variable, or λ_1, vary through time as a convex combination of two extreme values, λ_1^e and λ_1^r. These values are *not* estimated but calibrated. More specifically, we choose λ_1^e and λ_1^r to keep the average dynamics of the term structure of interest rates the same as in the simple stochastic environment (i.e., a stochastic model without business cycle dynamics). To conceptualize this idea, imagine two states of the world. With a constant-parameter model, the estimation procedure will provide an average of the two states, weighted by their occurence, over the estimation interval. With a time-varying parameter model, conversely, you will have two sets of parameters that describe the two states. The average result over the estimation period will, nevertheless, be the same. The actual approach to considering the *average* dynamics is identical to the approach used in performing our model diagnostic; this will be discussed later in this appendix. The results of this esimation/calibration procedure are outlined in the first two columns of Table B1.

Step 2 - The Business Cycle Our model of the business cycle, which is based on Hamilton (1989), is estimated using the non-linear filter suggested in the original Hamilton (1989) paper. This technique is outlined in Bolder (2002). These parameters are outlined in the final two columns of Table B1.

Step 3 - Government Financial Requirements The final step is the calibration of the parameters of the

[59] The details of these models are thoroughly described in Bolder (2001, 2002).

Table B1: **Term-Structure and Business Cycle Model Parameter Estimates**: In its first two columns, this table summarizes the term-structure model parameters for the two-factor CIR model. It also, in the final two columns, outlines the parameters for the AR(4) model using quarterly Canadian GDP data and the Hamilton (1989) constant transition probability model.

Term-structure model		**Business cycle model**	
Parameter	**Estimate**	**Parameter**	**Estimate**
κ_1	0.980	p	0.959
κ_2	0.119	q	0.535
θ_1	0.030	ϕ_1	0.177
θ_2	0.012	ϕ_2	0.474
σ_1	0.074	ϕ_3	0.301
σ_2	0.075	ϕ_4	-0.097
λ_1^e	-0.319	σ	0.725
λ_1^r	-0.134	μ_0	0.282
λ_2	-0.124	μ_1	2.126

financial requirement process. We opted for a simple calibrated process for the government's financial position. This demonstrated our explicit preference for ease of interpretation of the model parameters relative to a complicated econometric specification. The calibration is performed relative to a set of assumptions. In this study, for instance, we specified a desire for a sequence of financial requirements that are, in expectation, very close to zero with normal variation of plus or minus $1 billion. To achieve this, we selected different parameter sets, simulated their dynamics conditional on these parameter choices, and used the graphs in Figure B1 to compare the behaviour of the process to our previously-stated assumptions. The resulting parameters are shown in Table B2.

Table B2: **Financial-Requirement Model Parameter Estimates**: This table describes the set of calibrated parameters, used in this study, for the government's finanical requirements.

Parameter	**Parameter estimate**
β	-0.450
α	0.700
γ	1.000
ξ	1.000

Having worked through the previous sequence of steps, we have a set of potential candidates for the parameterization of our joint model for the business cycle, the term structure, and the government's financial position. We still need to satisfy ourselves that the dynamics of this model is consistent with historical behaviour. This is particularly true with the term structure of interest rates. The idea is quite simple. Consider the example of the term structure of interest rates. We take the original data that we used to

Figure B1: **Financial-Requirement Model Diagnostics**: This figure outlines four separate graphs that we use to ensure that the parameterization of the financial requirements process, $\{F_t, t \geq 0\}$ is consistently calibrated to our assumptions about its future behaviour.

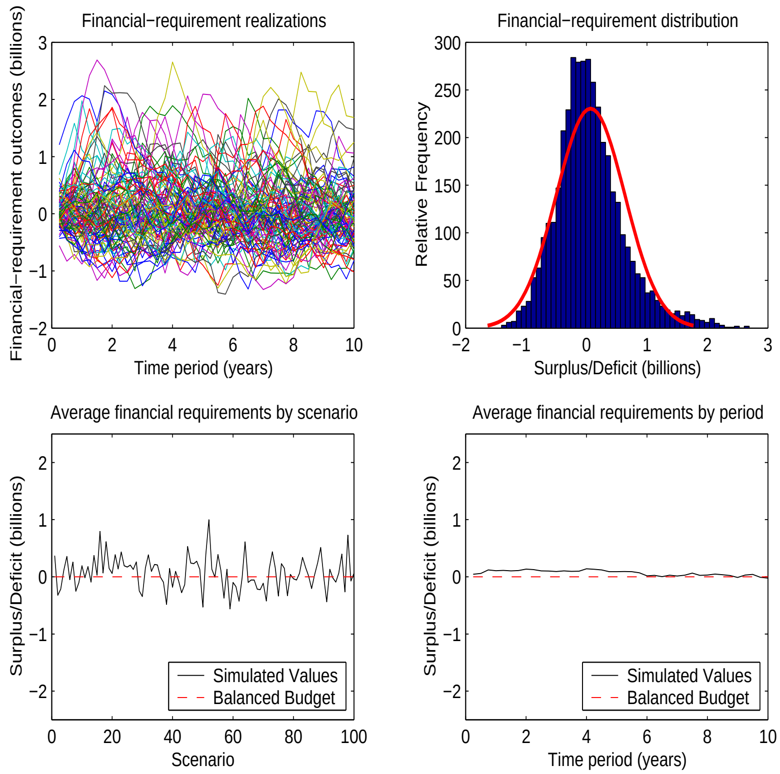

estimate our term-structure model and examine a number of its features, including its level, volatility, and the empirical distribution of short-term zero-coupon rates. We then, for a given parameterization, generate a large number of simulated term-structure outcomes. Then, we compare the features of the simulated data with the actual data and compare the results. We would expect that our model does a good job of capturing the features of the data, but some aspects of the model may be superior to others. This diagnostic procedure serves, therefore, to highlight both the strengths and weaknesses of our term-structure model. Figure B2 outlines a suite of six graphs that we use in our diagnostics of the term-structure dynamics associated with a given parameterization of our model. We discuss each of the individual graphs in order.

Level of the Curve The first graph compares the average level of the zero-coupon curve, over the estimation period, with the average level of a fairly large number of simulated zero-coupon curves. The idea is that we would like, on average, that when we simulate zero-coupon curves—which are so instrumental in constructing our debt-cost distributions—that they are similar to observed zero-coupon curves over, say, the last 10 years. This does not imply, of course, for any given realization that the zero-coupon curve cannot take a wide variety of forms.

Distribution of Curve Steepness The second graph provides some perspective on the relative frequency of the steepness of the term structure of interest rates. Inspection of the histogram reveals that, on average, the difference between the 10-year zero-coupon rate and the 3-month zero-coupon rate is approximately 180 basis points. Over the range of simulations, it can increase to almost 400 basis points and fall (i.e., invert) to -200 basis points.

Figure B2: **Interest-Rate Model Diagnostics**: This figure illustrates six separate graphs that we use to ensure that the parameterization of the term-structure of interest rates is consistent with historical behaviour.

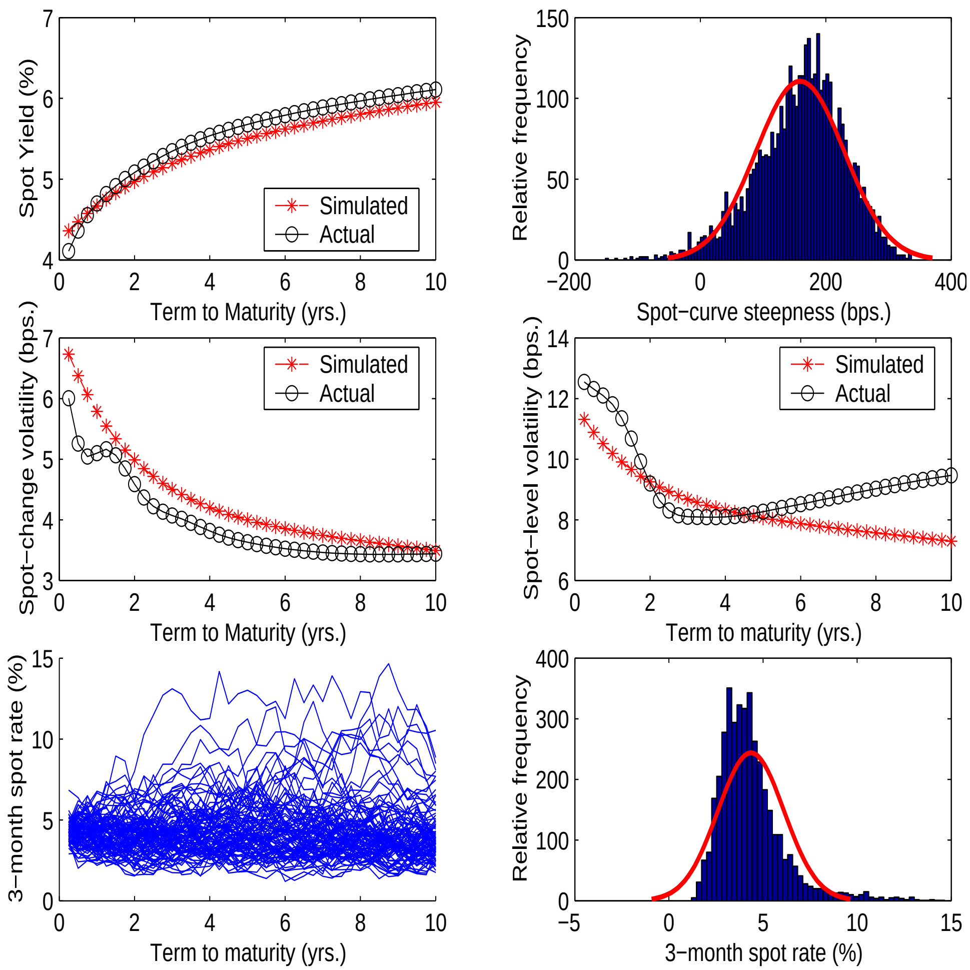

Volatility of Rate Changes The next two graphs in Figure B2 attempt to examine the volatility structure of the simulated term structures. The third graph compares the volatility of changes in actual zero-coupon rates—for maturities from 1 month to 10 years—to the associated volatility of the changes in simulated zero-coupon rates. The actual realized volatility of zero-coupon rate changes decreases gradually from seven basis points at the short end to three basis points at the 10-year maturity. This is a general stylized empirical fact about the term structure of interest rates.

Volatility of Rate Levels The fourth graph focuses on the volatility of the *level* of zero-coupon interest rates. That is, it compares the volatility of the level of actual zero-coupon rates—for maturities from 1 month to 10 years—with the associated volatility of the level of simulated zero-coupon rates. We note that the volatility of actual zero-coupon levels is approximately 13 basis points at one month, falls to about eight basis points at the two-year zero-coupon maturity, and then gradually increases to 10 basis points at 10 years. This quadratic form is difficult for our model to capture, but it does provide a reasonable fit to the actual data.

Range of Short-Rate Outcomes In the fifth graph, each individual sample path for the 3-month zero-coupon rate is illustrated. This provides some insight into the range of different outcomes that can occur across the entire range of simulations. Note, for example, that there are a small number of sample paths where the 3-month zero-coupon rate exceeds 10 per cent, but generally it remains bounded between 2 and 5 per cent.

Distribution of Short-Rate Outcomes The final graph provides another perspective on the range of simulated outcomes for the 3-month zero-coupon rate. In particular, it provides a histogram describing the relative frequency of various 3-month zero-coupon realizations across all simulations. Observe that the 3-month zero-coupon rate exhibits a positive and positively skewed empirical distribution. This is consistent with the non-negativity of the CIR model.

Appendix C: Variance Reduction

In this appendix, we will introduce a methodology that could conceivably help us reduce the computational expense of our current software platform: variance reduction. These are methods used to improve the efficiency of an estimator (computed by simulation) for a given amount of work. The idea relates to the fact that we can approximate the following integral,

$$\int_0^1 f(x)dx, \tag{58}$$

with the subsequent sum,

$$\frac{1}{n}\sum_{i=1}^{n} f(U_i), \tag{59}$$

where $U_i \sim \mathcal{U}[0,1]$. This result should look very similar to the basic idea behind numerical integration. The key point is that,

$$\mathbb{E}[f(U)] = \int_0^1 f(x)g(x)dx = \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} f(U_i), \tag{60}$$

where, $g(x) = 1$ is the probability density function for the uniform distribution.[60] Of course, we can't actually let n tend to infinity, as suggested in equation (53), so our estimators of this integral will always have some error, or variance. That is,

$$\text{var}\left(\frac{1}{n}\sum_{i=1}^{n} f(U_i)\right) > 0, \tag{62}$$

for all n. It turns out, however, that there are some clever tricks we can use that will actually reduce this variance for a given number of simulations, n. In fact, there are three main techniques that we will consider briefly in the following sections. Note that the underlying presentation is quite general and some thought is required to determine how these approaches might be applied to our problem.

C.1 Antithetic variables

Consider the following estimator for the integral in equation (58),

$$\hat{\theta} = \frac{1}{2}\left(f(U_1) + f(U_2)\right), \tag{63}$$

[60] The general case is,

$$\int_a^b f(x)dx = \lim_{n\to\infty} \frac{b-a}{n}\sum_{i=1}^{n} f(a + (b-a)U_i), \tag{61}$$

where $U_i \sim \mathcal{U}[0,1]$.

where $U_1, U_2 \sim \mathcal{U}[0,1]$. The variance of this estimator is defined as follows,

$$\text{var}(\hat{\theta}) = \frac{1}{2}\left(\text{var}(f(U_1)) + \text{cov}(f(U_1), f(U_2))\right). \tag{64}$$

Observe that if $\text{cov}(f(U_1), f(U_2)) < 0$, then we have reduced the variance of this estimator. It turns out that, for monotone f, the minimal covariance between U_1 and U_2 arises when we set $U_2 = 1 - U_1$. Thus, in the antithetic variance approach, we estimate equation (58) with the following sum,

$$\frac{1}{n}\sum_{i=1}^{n}\frac{1}{2}\left(f(U_i) + f(1-U_i)\right). \tag{65}$$

C.2 Control variates

The following expression holds immediately (or, rather, vacuously),

$$\int_0^1 f(x)dx = \int_0^1 g(x)dx + \int_0^1 \left(f(x) - g(x)\right)dx. \tag{66}$$

This implies the following estimator,

$$\hat{\theta} = \int_0^1 g(x)dx + \frac{1}{n}\sum_{i=1}^{n}\left[f(U_i) - g(U_i)\right]. \tag{67}$$

This is not a particularly useful extension, unless $g(x)$ is analytically integrable and the variance of the difference between $f(x)$ and $g(x)$ is smaller than the variance of $f(x)$ alone. In fact, the variance of this difference is,

$$\text{var}(f(U_i) - g(U_i)) = \text{var}(f(U_i)) + \text{var}(g(U_i)) - 2\text{cov}(f(U_i), g(U_i)). \tag{68}$$

We can see that this technique is effective when the covariance between f and g is large and positive.

C.3 Importance sampling

As in the previous technique, we require an easily integrable function, g, where there exists $c \in \mathbb{R}$ such that,

$$\int_0^1 cg(x)dx = 1. \tag{69}$$

This allows us to rewrite the integral in equation (58) as,

$$\int_0^1 f(x)dx = \int_0^1 \frac{f(x)}{cg(x)}cg(x)dx = \mathbb{E}\left[\frac{f(U)}{cg(U)}\right], \tag{70}$$

where $U \sim \mathcal{U}[0,1]$. The corresponding estimator is,

$$\hat{\theta} = \frac{1}{n}\sum_{i=1}^{n}\frac{f(Z_i}{cg(Z_i)}, \tag{71}$$

where $Z_i \sim cg(z)$. Finally, the variance of this estimator is,

$$\text{var}(\hat{\theta}) = \frac{1}{n}\text{var}\left[\frac{f(U)}{cg(U)}\right]. \tag{72}$$

In short, we have altered the underlying probability measure (from the Lebesque measure) to something else that redistributes the weight given to each observation. This approach is effective when our probability density function, $cg(x)$ is a close fit to the original function of interest, $f(x)$.

C.4 Generating non-central χ^2 random variates

The generation of random variates, in a computer setting, always begins with the simulation of uniform random variates. The uniform random variables are put through some form of transformation to arrive at a draw from the desired distribution, be it Gaussian, Gamma, Beta, or Cauchy. Many software packages provide built-in functions that perform this transformation in a manner that is transparent to the user. To use a given variance-reduction technique, one must work directly with the original uniform random variates. We must work through the nature of the transformation required to get to the random draws that are needed in our application.[61] In particular, to simulate the CIR model, we need to be able to take random draws from the non-central χ^2 distribution.

The theoretical foundations of constructing a non-central χ^2 distribution are described in Bolder (2002, Appendix B). Concisely put, if we have two random variables, X and Y, such that,

$$X \sim \mathcal{N}(\sqrt{\beta}, 1), \tag{73}$$

$$Y \sim \chi^2(\alpha - 1), \tag{74}$$

then,

$$XY \sim \chi^2(\alpha, \beta). \tag{75}$$

Or, in words, XY follows a non-central χ^2 distribution with α degrees of freedom and non-centrality parameter β.

The consequence is that, to simulate from a non-central χ^2 distribution, we need to simulate X from equation (73), Y from equation (74), and then simply take their product. Recall that, in each case, we need to begin with a uniform random variate and apply the necessary transformation. Fortunately, the generation of X is relatively straightforward. First, we ask the computer to provide us with a uniform random variate defined on the unit interval,

$$U \sim \mathcal{U}(0, 1). \tag{76}$$

[61] This section draws heavily from Fishman (1995).

We then apply the Box-Muller transformation to U; this tranformation is described in section C.5. The result is a standard normal random variate, which we will call V. That is,

$$V \sim \mathcal{N}(0,1). \tag{77}$$

Now, it follows from the properties of the normal distribution that,

$$X = V + \sqrt{\beta} \sim \mathcal{N}\left(\sqrt{\beta}, 1\right), \tag{78}$$

for $\beta \in \mathbb{R}$.

The generation of Y is more involved. As a first step, we would prefer to work with the Gamma distribution rather than the χ^2. There are several facts about the Gamma distribution that we will require in our analysis.

Fact 1 Given the form of Y described in equation (74), an equivalent representation is

$$Y \sim \text{Gamma}\left(\frac{\alpha-1}{2}, 2\right), \tag{79}$$

where $\alpha - 1$ is a positive integer. The first parameter is known as the *shape* parameter, and the second parameter is termed the *scale* parameter.

Fact 2 If $X_i \sim \text{Gamma}(\gamma_i, \delta)$ for $i = 1, ..., n$ then,

$$\sum_{i=1}^{n} X_i \sim \text{Gamma}\left(\sum_{i=1}^{n} \gamma_i, \delta\right). \tag{80}$$

Fact 3 If $X \sim \text{Gamma}(\gamma, 1)$ and $\delta \in \mathbb{R}$ then,

$$Y = \delta X \sim \text{Gamma}(\gamma, \delta). \tag{81}$$

This implies that we can always set the scale parameter equal to unity.

Fact 4 Given a random variable of the form, $X \sim \text{Gamma}(1, \delta)$, an equivalent representation is,

$$X \sim \text{Exp}(\delta). \tag{82}$$

Or, in words, a Gamma random variable with a shape parameter equal to unity and a shape parameter of δ is equivalent to an exponential random variable with parameter δ.

Fact 5 If $U \sim \mathcal{U}(0,1)$ is a uniform random variate defined on the unit interval, then by the method of inverse transform,

$$X = -\ln(U) \sim \text{Exp}(1). \tag{83}$$

One consequence of this fact is that it reduces our problem of simulating from equation (74), to a situation of generating a random draw from

$$Y \sim \text{Gamma}(\gamma, 1), \tag{84}$$

where $\gamma \in \mathbb{R}$ is an arbitrary real number. Moreover, we can always represent γ as the sum of n and α where these two values are the integer and real part of γ, respectively. For example, if $\gamma = 4.35$, then $n = 4$ and $\alpha = 0.35$. The plan, therefore, is to generate two Gamma random variates. The first is,

$$X_1 \sim \text{Gamma}(n, 1), \tag{85}$$

where $n \in \mathbb{N} \setminus \{0\}$ and,

$$X_2 \sim \text{Gamma}(\alpha, 1), \tag{86}$$

where $\alpha \in (0, 1)$. It follows from our first, second, and third facts that,

$$\begin{aligned} Y = 2(X_1 + X_2) &\sim \text{Gamma}(n + \alpha, 2), && (87)\\ &\sim \text{Gamma}(\gamma, 2), && (88)\\ &\sim \chi^2(2\gamma). && (89) \end{aligned}$$

To actually generate, from a set of uniform random variates, a variable from equation (85), we employ the fifth fact. That is, we ask our computer to generate $U_1, ..., U_n$ uniform random variates defined on the unit interval and construct X_1 as,

$$X_1 = -\ln \left(\sum_{i=1}^{n} U_i \right), \tag{90}$$

from which it follows that $X_i \sim \text{Gamma}(n, 1)$, as desired.

To generate X_2 is not as straightforward. First, we use our computer to simulate three uniform random variates, U_1, U_2, and U_3 defined on the unit interval. Then, we set

$$V = -\ln(U_3) \sim \text{Exp}(1). \tag{91}$$

The next step is to construct the following random variable,

$$W = \frac{U_1^{\frac{1}{\alpha}}}{U_1^{\frac{1}{\alpha}} + U_2^{\frac{1}{\alpha - 1}}}. \tag{92}$$

If we have that the condition,

$$U_1^{\frac{1}{\alpha}} + U_2^{\frac{1}{\alpha - 1}} < 1, \tag{93}$$

then,

$$VW \sim \text{Gamma}(\alpha, 1). \tag{94}$$

as desired. If the condition in equation (93) does *not* hold, then we continue to generate U_1 and U_2 until it is satisfied. This is termed Johnk's algorithm.

C.5 The Box-Muller transformation

This technique is very useful for transforming independent uniformly distributed random variables into independent random variables from a standard normal distribution. Traditional techniques do not work in this instance, because we do not have an analytical expression for the inverse of the cumulative distribution function of a standard normal distribution. This ingenious algorithm uses the usual transformation,

$$f_{Y_1,Y_2}(y_1, y_2) = f_{X_1,X_2}(x_1, x_2) \cdot |J|, \tag{95}$$

where $|J|$ is the determinant of the Jacobian matrix. The success of this method hinges on two cleverly selected transformations, y_1, y_2. Given two random drawn values, x_1 and x_2, taken from a uniform distribution on the unit interval, $(0, 1)$, consider the following transformations:

$$\begin{aligned} y_1 &= \sqrt{-2\ln x_1}\cos 2\pi x_2, \\ y_2 &= \sqrt{-2\ln x_1}\sin 2\pi x_2. \end{aligned} \tag{96}$$

The claim is that y_1 and y_2 are independent standard normal variates. This may appear somewhat questionable, but let us take a closer look. First, let us solve for x_1 and x_2 in terms of our transformation variables. This will require a bit of cunning. Consider $y_1^2 + y_2^2$,

$$\begin{aligned} y_1^2 + y_2^2 &= (\sqrt{-2\ln x_1}\cos 2\pi x_2) + (\sqrt{-2\ln x_1}\sin 2\pi x_2), \\ y_1^2 + y_2^2 &= -2\ln x_1 \cos^2 2\pi x_2 - 2\ln x_1 \sin^2 2\pi x_2, \\ y_1^2 + y_2^2 &= -2\ln x_1(\cos^2 2\pi x_2 + \sin^2 2\pi x_2), \end{aligned} \tag{97}$$

and recall that,

$$\cos^2 2\pi x_2 + \sin^2 2\pi x_2 = 1, \tag{98}$$

which implies that,

$$\begin{aligned} y_1^2 + y_2^2 &= -2\ln x_1, \\ x_1 &= e^{\frac{-y_1^2 - y_2^2}{2}}. \end{aligned} \tag{99}$$

Now, to find x_2, let us take a look at the ratio of y_1 and y_2,

$$\begin{aligned}
\frac{y_1}{y_2} &= \frac{\sqrt{-2\ln x_1}\cos 2\pi x_2}{\sqrt{-2\ln x_1}\sin 2\pi x_2}, \\
\frac{y_1}{y_2} &= \frac{\cos 2\pi x_2}{\sin 2\pi x_2}, \\
\frac{y_1}{y_2} &= \tan 2\pi x_2, \\
\arctan\left(\frac{y_1}{y_2}\right) &= \arctan(\tan 2\pi x_2), \\
x_2 &= \frac{1}{2\pi}\arctan\left(\frac{y_1}{y_2}\right).
\end{aligned} \tag{100}$$

It remains to calculate the Jacobian matrix, which will have the following form:

$$\begin{bmatrix} \frac{\partial x_1}{\partial y_1} & \frac{\partial x_1}{\partial y_2} \\ \frac{\partial x_2}{\partial y_1} & \frac{\partial x_2}{\partial y_2} \end{bmatrix} \tag{101}$$

Let us look at each of these partial derivatives:

$$\begin{aligned}
\frac{\partial x_1}{\partial y_1} &= -y_1 e^{\frac{-y_1^2-y_2^2}{2}}, \\
\frac{\partial x_1}{\partial y_2} &= -y_2 e^{\frac{-y_1^2-y_2^2}{2}}, \\
\frac{\partial x_2}{\partial y_1} &= \frac{1}{2\pi}\left(\frac{-y_2}{y_1^2}\right)\frac{1}{1+(\frac{y_2}{y_1})^2}, \\
\frac{\partial x_2}{\partial y_2} &= \frac{1}{2\pi}\left(\frac{1}{y_1}\right)\frac{1}{1+(\frac{y_2}{y_1})^2}.
\end{aligned} \tag{102}$$

Our Jacobian is thus,

$$\begin{bmatrix} -y_1 e^{\frac{-y_1^2-y_2^2}{2}} & -y_2 e^{\frac{-y_1^2-y_2^2}{2}}, \\ \frac{1}{2\pi}\left(\frac{-y_2}{y_1^2}\right)\frac{1}{1+(\frac{y_2}{y_1})^2} & \frac{1}{2\pi}\left(\frac{1}{y_1}\right)\frac{1}{1+(\frac{y_2}{y_1})^2} \end{bmatrix}, \tag{103}$$

and the determinant is, therefore,

$$\begin{aligned}
|J| &= \left(-y_1 e^{\frac{-y_1^2-y_2^2}{2}}\right)\cdot\left(\frac{1}{2\pi}\left(\frac{1}{y_1}\right)\frac{1}{1+(\frac{y_2}{y_1})^2}\right) - \left(-y_2 e^{\frac{-y_1^2-y_2^2}{2}}\right)\cdot\left(\frac{1}{2\pi}\left(\frac{-y_2}{y_1^2}\right)\frac{1}{1+(\frac{y_2}{y_1})^2}\right), \\
&= \left(-\frac{1}{2\pi}e^{\frac{-y_1^2-y_2^2}{2}}\frac{1}{1+(\frac{y_2}{y_1})^2}\right) - \left(-\frac{1}{2\pi}\frac{y_2^2}{y_1^2}e^{\frac{-y_1^2-y_2^2}{2}}\frac{1}{1+(\frac{y_2}{y_1})^2}\right), \\
&= -\frac{1}{2\pi}e^{\frac{-y_1^2-y_2^2}{2}}\frac{1}{1+(\frac{y_2}{y_1})^2}\left(1+\left(\frac{y_2}{y_1}\right)^2\right), \\
&= -\frac{1}{2\pi}e^{\frac{-y_1^2-y_2^2}{2}}, \\
&= -\left(\frac{1}{\sqrt{2\pi}}e^{\frac{-y_1^2}{2}}\right)\left(\frac{1}{\sqrt{2\pi}}e^{\frac{-y_2^2}{2}}\right).
\end{aligned} \tag{104}$$

This happy result is exactly what we were trying to establish. If we return to our original formula and recall that the density function of the uniform distribution on $(0,1)$ is 1, we have that[62]

$$f_{Y_1,Y_2}(y_1,y_2) = f_{X_1,X_2}(x_1,x_2)\cdot|J| = -\left(\frac{1}{\sqrt{2\pi}}e^{\frac{-y_1^2}{2}}\right)\left(\frac{1}{\sqrt{2\pi}}e^{\frac{-y_2^2}{2}}\right). \tag{105}$$

That is, y_1 and y_2 are independent standard normal variates, as was desired.

[62]Recall that the density of uniform (a,b) is defined as $f(x) = \frac{1}{b-a}$, $\forall\, x \in [a,b]$.

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. ***Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue*****.**

2003

2003-9	Bank Lending, Credit Shocks, and the Transmission of Canadian Monetary Policy	J. Atta-Mensah and A. Dib
2003-8	Comparing Alternative Output-Gap Estimators: A Monte Carlo Approach	A. Rennison
2003-7	Testing the Stability of the Canadian Phillips Curve Using Exact Methods	L. Khalaf and M. Kichian
2003-6	Valuation of Canadian- vs. U.S.-Listed Equity: Is There a Discount?	M.R. King and D. Segal
2003-5	Shift Contagion in Asset Markets	T. Gravelle, M. Kichian, and J. Morley
2003-4	Are Distorted Beliefs Too Good to be True?	M. Misina
2003-3	Modélisation et prévision du taux de change réel effectif américain	R. Lalonde and P. Sabourin
2003-2	Managing Operational Risk in Payment, Clearing, and Settlement Systems	K. McPhail
2003-1	Banking Crises and Contagion: Empirical Evidence	E. Santor

2002

2002-42	Salaire réel, chocs technologiques et fluctuations économiques	D. Tremblay
2002-41	Estimating Settlement Risk and the Potential for Contagion in Canada's Automated Clearing Settlement System	C.A. Northcott
2002-40	Inflation Changes, Yield Spreads, and Threshold Effects	G. Tkacz
2002-39	An Empirical Analysis of Dynamic Interrelationships Among Inflation, Inflation Uncertainty, Relative Price Dispersion, and Output Growth	F. Vitek
2002-38	Oil-Price Shocks and Retail Energy Prices in Canada	M. Chacra
2002-37	Alternative Public Spending Rules and Output Volatility	J.-P. Lam and W. Scarth

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: http://www.bankofcanada.ca

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : http://www.banqueducanada.ca