



DOCUMENTS DE TRAVAIL DE LA DGTPE

2007/10 - November 2007

What are the consequences of active management of average debt maturity in terms of cost and risk?

Jean-Paul RENNE

What are the consequences of active management of average debt maturity in terms of cost and risk?

Jean-Paul RENNE

The views expressed in this working paper are those of the author. The purpose of disseminating this paper is to stimulate debate and to elicit comments and criticism.

At the time of writing, Jean-Paul Renne was in charge of modelling at Agence France Trésor.

Direction Générale du Trésor et de la Politique Économique – Ministère de l'Économie, des Finances et de l'Emploi

Télédoc 647 - 139, rue de Bercy – 7572 Paris Cédex 12 – France – Tél : 33 1 44 87 18 51 – Fax : 33 1 53 18 36 25

E-mail : jean-paul.renne@aft.gouv.fr – Site WEB : <http://www.minefi.gouv.fr>

SOMMAIRE

Non-technical summary.....	5
Introduction.....	7
I- Methodology.....	8
II- Strategies used.....	10
III- Results of the “baseline” simulations.....	12
IV- Sensitivity of results to model assumptions.....	15
V- Back-testing.....	19
References.....	22
Annex 1 – The Model.....	23
Annex 2 – Risk measures.....	29

RÉSUMÉ

Depuis 2001, l'Agence France Trésor gère la durée de vie moyenne de la dette. Dans une situation « normale » caractérisée par une courbe des taux pentue, donc des taux longs plus élevés et des taux courts plus faibles mais plus volatils, réduire cette durée moyenne devrait permettre de diminuer la charge d'intérêt moyenne sur longue période, toutes choses égales par ailleurs. En contrepartie, cela entraîne une augmentation de la variabilité de cette charge à court terme. Le pilotage de la durée de vie moyenne doit donc passer par une évaluation quantitative du compromis entre la charge d'intérêt moyenne et la variabilité de celle-ci. Ce document propose une telle évaluation. L'approche repose sur un modèle macro- financier estimé sur données historiques et utilisé pour générer aléatoirement un grand nombre de scénarios futurs envisageables. L'application des différentes stratégies de financement sous ces différents scénarios nous renseigne sur la moyenne des charges d'intérêt (mesure de coût) et la variabilité de celle-ci (mesure de risque) induites par chacune des stratégies.

ABSTRACT


Since 2001, Agence France Trésor has been managing the average maturity of debt. In “normal” circumstances, characterised by a rising rate curve, with higher long rates and weaker but more volatile short rates, reducing this average maturity should make it possible to reduce interest costs on average over a long period, all things being equal. On the other hand, this increases the variability of this expense. Average debt maturity management therefore requires quantitative assessment of the compromise between average interest payments and payments variability. This paper aims at providing such an assessment. More precisely, the objective of the proposed approach is to compare the performances of financing strategies implying different average debt maturities. A macro-finance model is estimated on historical data and further used in order to stochastically generate future macro-finance scenarios. Each financing strategy is then applied under these scenarios, which provides us with average interest payments (cost measure) and the variability of these payments (risk measure) associated with this strategy.

Non-technical summary

In 2001, acting on instructions from the Minister of Finance and after consultation with Parliament, Agence France Trésor initiated a programme to manage the average maturity of government debt using interest rate swaps. One advantage of such swaps is that they can reduce average debt maturity without affecting the liquidity of the longest-dated securities, for which there is increasing demand¹.

With a rising yield curve (on average), paying a short-term rate to receive a long-term rate is one strategy that can reduce average interest payments, if it is applied over the very long term. Figure 1 shows that the slope of the yield curve, as measured by the spread between the yield on 10-year Treasury bonds and the yield on 2-year Treasury notes (BTAN), was positive 78% of the time between 1989 and 2000, while the spread between the 10-year yield and the 3-month yield was positive 73% of the time. Since the introduction of the euro, the spread has been positive 100% of the time. On the other hand, replacing payments of the same coupon rate over 10 years with payments that vary more frequently will increase the variability of interest payments, which raises problems for oversight of the budget.

Figure 1 – Yield curve slopes



If the increase in the variability of interest payments is slight compared to the reduction in average interest payments, reducing average debt maturity through swaps could be seen as an effective strategy. The risk-taking induced by this strategy needs to be measured properly and left up to the discretion of political authorities. An AFT study based on simulations made with an econometric model was carried out in 2001. This study led to the conclusion that such a strategy should be recommended.

In 2001, the Minister of Finance proposed a framework for using the strategy. The framework calls for suspending the strategy if the volatility of interest rates becomes unfavourable. The Minister of Finance

¹ Some of the growing demand for very-long-term interest-rate instruments comes from European life insurance companies and pension funds. Regulatory changes mean that such entities are required to adopt stricter asset-liability management, which involves extending the average maturity of their assets to match that of their liabilities (see BIS, 2007).

suspended reduction of average debt maturity in accordance with the framework in July 2002, but the “short” swaps entered into in 2001 and 2002 were renewed in order to limit exposure to rates of 1 year or less².

This paper proposes a new quantification of the cost/risk considerations at the heart of average debt maturity management. A model of the same type as the one that AFT developed in 2001 was estimated, incorporating recent data to provide a better account of current debt management conditions. Some aspects of the model were also reviewed. Finally, we conducted several tests of the robustness of the results.

The main findings are as follows.

Using swaps to reduce average debt maturity by one year under “average” rate conditions, meaning that the swaps are entered into at a time when rates are at their steady state values, historically estimated at 4.5% for short-term rates and 5.5% for long-term rates:

- reduces average annual interest payments by nearly 1bn euros (with regard to the 41bn euros in interest payments included in the 2008 budget bill),
- increases average annual interest payment variation by some 0.45bn euros (from 1.45bn euros to 1.9bn euros),
- raises the Cost-at-Risk –i.e the amount that interest payments will be less than with a probability of 90%– by 0.8bn euros, from 53.6bn to 54.4bn euros,
- does not increase, and even reduces, variability of the budget balance (primary budget balance + interest payments) because of the positive correlation between the cyclically adjusted budget balance and short-term rates. The intuitive explanation of this result is that, at the low point in the cycle, the budget deficit grows, but short-term interest rates tend to fall.

The robustness of these results was tested by altering several assumptions:

- When the average size of the macro-financial shocks affecting the economy in a simulation is increased, the effect that reducing average debt maturity has on risk is expanded.
- The results are not greatly affected by altering the type of distribution that shocks to the economy follow or the average curvature of the yield curve (with a constant average slope between 3 months and 10 years).
- If the model is estimated on a more recent period (1995-2007, as opposed to 1986-2007 for the baseline model), the results are relatively more favourable for a reduction of average debt maturity. In this case, the Cost-at-Risk arising from interest payments decreases as average debt maturity falls.

The expected gains in terms of average cost from a policy to reduce average debt maturity depend significantly on the assumption about the average slope of the yield curve. Naturally, the gains decrease as the average slope of yield curve flattens out. On the other hand, the effects in terms of risk are not very sensitive to the assumption about the average slope of the yield curve.

² The swaps programme relied on “long” swaps where a 10-year fixed rate was received and a 6-month floating rate was paid and “short” swaps where a 6-month floating rate was received and a 2-year fixed rate was paid.

Introduction

Agence France Trésor has been managing average debt maturity since 2001. Under “normal” circumstances, with a rising yield curve, where long-term rates are higher and short-term rates are lower, but more volatile, a reduction in average debt maturity should make it possible to reduce average interest payments over a long period, all else being equal. On the other hand, it would also make interest payments more variable. Therefore, management of average debt maturity calls for a quantitative assessment of the trade-off between the average level of interest payments and the variability of the payments.

This paper proposes a framework for quantitative analysis aimed at assessing the performances of financing strategies implemented over a long period. The approach presented is general enough to assess a wide variety of strategies, but it is used here for a specific analysis of the consequences of managing average debt maturity.

The use of models to assess prospective public debt management strategies has grown rapidly over the last decade. The World Bank and the IMF³ have stressed the potential that this approach has to capture the issues relating to public debt management more fully. The models used by public debt managers from several countries have given rise to various publications. Of particular note are the contributions by Bolder (2002 and 2003) from Canada, by Bergström, Holmlund and Lindberg (2002) from Sweden, and by Pick and Anthony (2006) from the United Kingdom. In addition, the Central Bank of Denmark’s annual publication, *Danish Government Borrowing and Debt*, presents simulation results. The Dutch State Treasury Agency (2007) also recently presented its new risk analysis model. These analyses are based on simulation of interest payments in future years⁴. In general, the simulations are taken from a model that reproduces the interrelated dynamics of macroeconomic and financial variables that are important for public debt management.

The rest of this paper is organised as follows. The first section describes the approach in general terms⁵. The second section gives details about the strategies that are compared using the model. The third section presents the results of the simulations made using the “baseline” model. The fourth section analyses how different assumptions influence the results. And the fifth section proposes a back-testing exercise.

³ See Guidelines for Public Debt Management, IMF and World Bank (2001) or Valendia (2002).

⁴ For a specific presentation of the method for assessing financing strategies using simulations, readers may refer to Black and Telmer (2000).

⁵ Annex 1 provides a more detailed presentation of the macro-financial model.

I- Methodology

This section answers three questions about (1) financing strategies, (2) simulated scenarios, and (3) strategy performance measures in order to explain the underpinning of the approach.

I-1 What is a financing strategy?

In each period, the government needs to call on the financial markets to cover its borrowing requirement. The government can borrow by issuing several types of securities (Treasury bills, notes and bonds). These securities have different maturities and their redemption can be for a real or notional value. Outstanding negotiable public debt as of 30 October 2007 broke down into 89.2bn euros in Treasury bills (BTF), 196.9bn euros in Treasury notes (BTAN), and 647.3bn euros in Treasury bonds (OAT).

A *financing strategy* means a “rule” that the public debt manager plans to follow for a given period. More specifically, the strategy means the way the debt manager allocates the borrowing requirement for each new period to the different securities available. The simplest strategies, which are the ones applied here, call for allocating the borrowing requirement to the various securities in constant proportions over time. Systematic use of such strategies means that the debt portfolio moves towards an allocation that is specific to each strategy. For example, if the strategy consists of issuing 10-year bonds to cover the entire borrowing requirement, the portfolio will naturally end up containing only 10-year bonds, regardless of the initial portfolio allocation.


I-2 Measuring the performance of different strategies

First of all, it must be noted that strategies may only be compared in relation to scenarios of changes in macro-financial conditions. This means that we can observe the time series of interest payments associated with each strategy that would occur under *one* central scenario, such as the one proposed by the Consensus Forecast. However, under another scenario, the performances of the various strategies might not be the same and their ranking could be different. Therefore, rather than assessing the performances of the strategies under a *single* given scenario, even if it is a central scenario, it is better to observe the different performances of each strategy under a *set* of foreseeable scenarios reflecting the “range of possible macro-financial scenarios” for the coming years. At the same time, when we apply each strategy under all of the scenarios, we obtain a set of interest payment time series that provide a picture of the “range of possible interest payments” associated with each financing strategy. We can then derive the cost and risk measures from this set of trajectories.

Figure 2 illustrates this approach. Suppose we are studying a given strategy. After applying this strategy under each of the future scenarios considered, we obtain the same number of associated interest payment time series. We are then able to calculate the average trajectory. This then provides us with information about the expected average interest payment if the strategy is applied systematically.

Furthermore, the dispersion of the interest payment time series provides us with information about the risk associated with the strategy. A financing strategy will be deemed “risky” if the simulated interest payment time series include some that deviate significantly from the average trajectory, or if they show big variations from one year to the next (these two concepts of “instantaneous” and “cumulative” risk are explained in Annex 2).

Figure 2 – Notions of cost and risk associated with a financing strategy



I-3 Scenarios to be considered

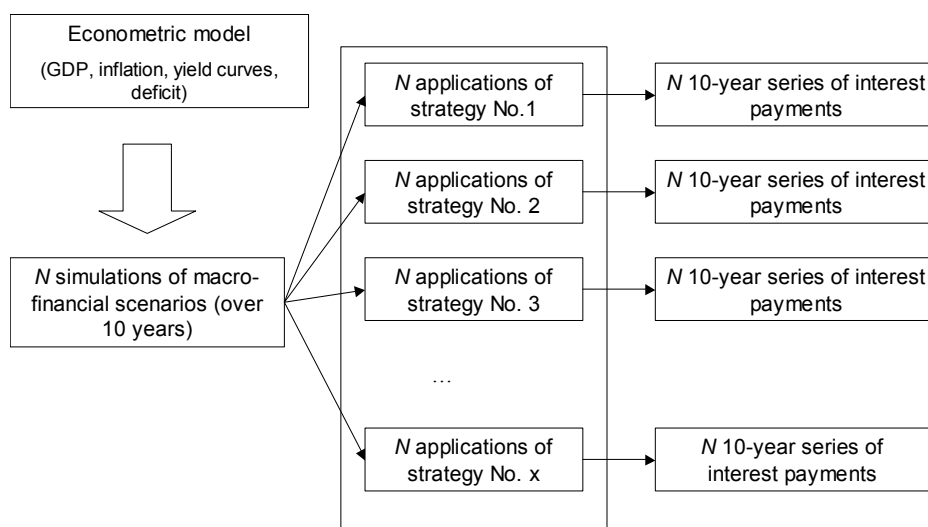
Financing strategies have to be applied under a set of foreseeable macro-financial scenarios. How should such scenarios be selected? One possibility would be to apply a few key scenarios, such as a central scenario, a high scenario and a low scenario. However, the resulting trajectories would necessarily be “smooth” over the coming decades and could not show how macro-financial conditions are subject to unforeseeable shocks in each period. Yet, this characteristic is critical for analysing how a strategy translates such shocks into interest payments, which is precisely how we define the performance of the strategy in terms of risk.

One alternative, which is the solution used in this document, is to simulate a very large number of scenarios in which macro-financial variables are randomly affected by shocks in each period. The characteristics of the shocks, such as standard deviation and correlations, correspond to the characteristics that can be estimated from historical data⁶. Some of the modelling elements may be changed with regard to the estimates if there is reason to think that breaks have occurred recently. We then need to simulate a large number of scenarios in order to cover the set of foreseeable scenarios sufficiently. Figure 3 provides a stylised visual representation of the approach.

It should be noted that, for the same reasons, back-testing a strategy in no way provides a guide to the future, since the past trajectory of the economy is only one outcome among many other possible occurrences, and the likelihood that it will happen again in the future is infinitesimal (see section V).

⁶ The way the shocks are drawn randomly is a key point of the approach. Consequently, the distributions followed in drawing the shocks must be tested for robustness. These tests are discussed in section IV. In the baseline simulations, the sample of shocks is drawn from estimated past shocks (bootstrap technique).

Figure 3 – Principle of the approach



Key: The model is used to simulate a large number ($N = 10,000$) of macro-financial scenarios over a 10-year period. Each strategy is applied under each of these scenarios. Thus, after the simulations, we obtain N 10-year series of interest payments for each strategy, from which we derive a cost measure (average interest payments) and a risk measure (variability of interest payments).

II- Strategies used

As mentioned above, only simple strategies are considered here. These strategies consist of systematically allocating the borrowing requirement to the securities available to the debt manager in constant proportions over the simulation period⁷. Each strategy is, therefore, perfectly defined by this allocation. Even in this narrow category, there are still an infinite number of strategies, There are as many strategies as there are ways to allocate 100% to the eight types of securities considered in the model⁸. Naturally, we are able to study only a few of these strategies.

In this study, we limit ourselves to a family of strategies that differ only in the specified proportions of 3-month Treasury bills (BTF), 2-year Treasury notes (BTAN) and 10-year Treasury bonds (OAT). More specifically, we assume that, in the first period, swaps involving these three types of interest rates are entered into in order to obtain the desired debt structure and, thus, the desired average debt maturity instantly. Interest rate swaps are then renewed systematically until they mature. This means that the proportions of 1-year Treasury notes, 5-year Treasury notes, 30-year Treasury bonds and inflation-indexed securities under the different strategies studied remain constant at their current levels, but the proportions of the three other securities used in the model describe the set defined as follows:

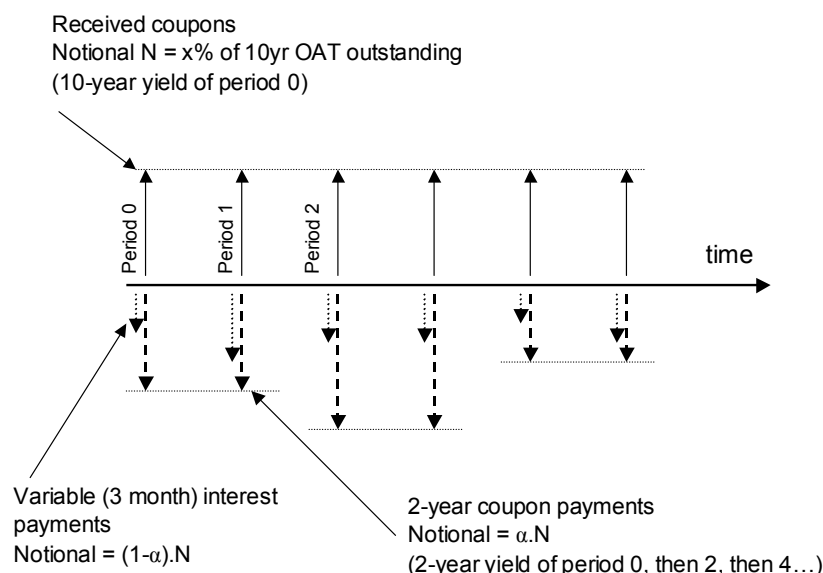
⁷ However, redemptions are financed by issuing the same securities as those coming to maturity (in other words, only the government deficit is financed using the specified allocation). This is done to avoid having the longest-dated maturities gradually dominate the debt (see Bolder, 2003).

⁸ The eight types of securities considered in the model are: 3-month Treasury bills (BTF), 1-year, 2-year and 5-year Treasury notes (BTAN), 10-year Treasury bonds (OAT) (nominal and inflation-indexed) and 30-year Treasury bonds (nominal and inflation-indexed). Introducing 15-year and 50-year Treasury bonds would be unlikely to change the results.

The proportion of 10-year Treasury bonds varies from 60% to 110% of its current level⁹.

This variation is financed by altering exposure to 3-month Treasury bills (BTF) and 2-year Treasury notes (BTAN) so that the relative proportions of the latter range from “all 3-month” to “all 2-year” (α denotes the proportion of 2-year Treasury notes).

Figure 4 – Application of the swap policy



Some of these strategies fall outside of the initial swaps policy framework of Agence France Trésor, under which, roughly speaking, coupons maturing within 2 years are paid and coupons maturing in 10 years are received. In this case, Figure 4 does not show 3-month Treasury bill coupons ($\alpha = 1$). The market does not offer standard swaps that make it possible to obtain these cash flow streams directly (swaps where a 10-year rate is received and a 2-year rate is paid), so Agence France Trésor enters into two types of transactions:

- Swaps where a 10-year fixed rate is received and the 6-month Euribor is paid (A swaps, hereafter);
- Swaps where the 6-month Euribor is received and a 2-year fixed rate is paid (B swaps, hereafter).

Under the swaps policy implemented in 2001, the notional amounts are the same, the 6-month Euribor coupons cancel each other out and Agence France Trésor receives 10-year coupons and pays 2-year coupons¹⁰.

The second means of differentiating between strategies is the proportion of 3-month Treasury bills relative to 2-year Treasury notes. It does not match the swap policy set out above exactly, where a 10-year coupon is received and a 2-year coupon is paid. Here we are testing strategies based on notional amounts of A and B swaps that might be different. This distinction is noteworthy nonetheless because it enables us to assess to what extent a change in average debt maturity has effects that are dependent on the way the change is implemented¹¹.

⁹ For example, an 80% proportion of 10-year Treasury bonds means that interest-rate swaps are entered into to reduce the payment of a 10-year fixed rate on a notional amount equal to 20% of outstanding 10-year Treasury bonds.


¹⁰ In practice, 6-month Euribor positions can be left open for the time it takes to renew them.

¹¹ There are several ways of reducing the average maturity of a debt portfolio. In this case, one way would be to reduce the proportion of 10-year treasury bonds relative to 2-year Treasury notes, and another way would be to reduce the proportion of 2-year Treasury notes issued and increase the proportion of 3-month Treasury bills issued.

It also enables us to determine the influence of mismatching of long swaps and short swaps.

These strategies result in debt portfolios with average maturities ranging from 5.3 years to 7.5 years (versus the average maturity of around 7.0 years, after swaps, as of 30 September 2007). Figure 5 shows the average maturities of the strategies studied in relation to the two degrees of freedom defined above.

Figure 5 – Average maturity of the strategies tested



III- Results of the “baseline” simulations

Each set of simulations is produced using a model and therefore depends on the underlying assumptions of that model. The results presented in this section are those relating to the “baseline” model, which is presented in greater detail in Annex 1. The sensitivity of the results to several model assumptions is discussed in the following section.

The stationary values chosen are 2% annual GDP growth, 2% annual inflation, a 4.5% short-term interest rate and a 5.5% 10-year interest rate¹². These values are also used as the initial conditions for the simulations.

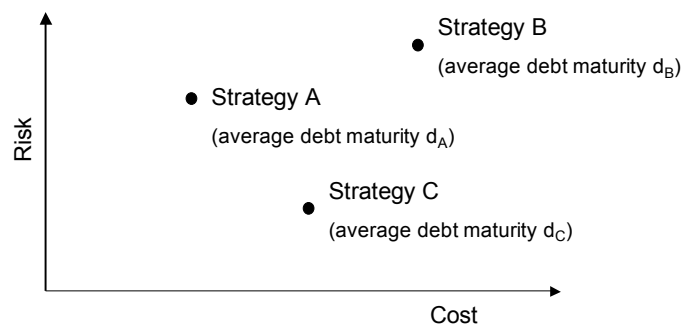
The simulations (10,000 ten-year scenarios) can be used to derive cost and risk measures for each strategy. These two measures are used as coordinates to plot the different strategies on a cost/risk plane. Box 1 provides details about this way of presenting the results and how it should be interpreted.

¹² The stationary values are the average values of the series that would be obtained if the model were simulated over a very long period. See Annex 1 for more details about the choice of these values.

Box 1 – Presentation of the results

A chart is used to facilitate comparisons of the different strategies. The strategies are plotted along two axes. One axis corresponds to the expected average cost when the strategy is applied and the other axis corresponds to the risk associated with the strategy.

Figure 6 – Interpretation of the cost-risk plane



The figure above shows an example of the distribution of the strategies on a cost-risk plane. Compared to strategies A and C, strategy B is not efficient because it carries a higher average cost and greater risk. On the other hand, it is impossible to say whether strategy A is better than strategy C, since C carries lower risk but a higher average cost. In order to decide between A and C, we need to know how much risk we are willing to take to lower the average cost. This is called risk aversion. It is not up to the debt manager to make such choices, which are political in nature.

In conclusion, this type of chart cannot be used to define the optimum debt management strategy, but rather to eliminate the strategies that are dominated by the others carrying lower average costs *and* lower risks. The set of strategies that are not dominated by the others, which in our simplified example is the straight line from A to C, is called the *efficiency frontier*.

The simulations show that the average interest payments increase as a function of average debt maturity (see Figure 15). More specifically, we assessed an average difference in interest payments of 0.13 annual GDP points between the shortest average debt maturity studied (5.3 years) and the longest (7.5 years). This means that using interest-rate swaps to reduce debt maturity by one year would reduce the average annual interest payments by nearly 1bn euros in the first 10 years.

With regard to the impact that a reduction in average debt maturity would have on risk, the results are qualitatively different, depending on whether the risk is seen as the variability of *interest payments* or as the variability of the *budget balance* (the primary balance plus interest payments). An examination of the cost-risk planes in Figure 15 leads to the following conclusions about the impact of a reduction in average debt maturity on risk with constant relative proportions of 3-month Treasury bills and 2-year Treasury notes:

A reduction in average debt maturity increases the variability of interest payments.

The simulations suggest that a reduction of one year in average debt maturity increases the average annual variation in interest payments by approximately 30% (more than 0.45bn euros, or from 1.45bn euros to 1.9bn euros) and increases the Cost-at-Risk by 0.04 GDP points, or 0.8bn euros (raising the Cost-at-Risk from 53.6bn euros to 54.4bn euros).

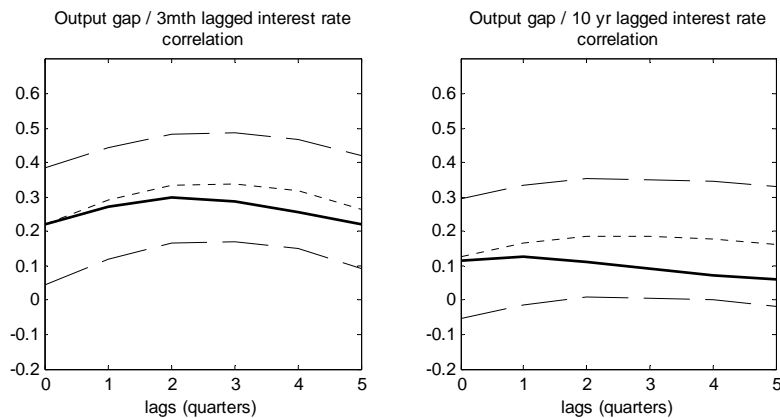
On the other hand, a reduction of average debt maturity from 7.5 years to 5.3 years does not have a significant effect on the variability of the budget balance.

Reducing average debt maturity from 7.5 years to 5.3 years leads to changes in the risk measures of about 1%, whether we measure in terms of average annual variation or the Cost-at-Risk of the budget balance. The Cost-at-Risk¹³ ranges from 2.92% to 2.95% of GDP and the average annual variation ranges from 0.927% to 0.938% of GDP.

Therefore, a reduction in average debt maturity tends to increase the variability of interest payments, but it has virtually no impact on the variability of the budget balance. In other words, as average debt maturity decreases, the increase in the variability of interest payments is offset by a stronger negative correlation between this variability and that of the primary deficit.

There are two explanations for this phenomenon. First, the correlation between short-term interest rates and the primary budget balance is stronger than that between long-term interest rates and the primary budget balance (the primary deficit is negatively correlated to the business cycle, see Annex 1.B). Figure 7 below shows the correlations between the output gap and the 3-month lagged interest rate, and the correlation between the output gap and the 10-year lagged interest rate. Secondly, short-term debt benefits more from this correlation because it is renewed more often. When long-term debt is issued, only the first interest payments are sensitive to the correlation between interest rates and the budget balance (at the time of issue). When short-term debt is issued, interest payments benefit fully from the correlation.

Figure 7 – Correlations between the output gap and the interest rates



Key: The dashed lines correspond to the bounds of the 90% confidence interval of the correlations estimated from historical data (dotted lines). The thick line corresponds to the correlations simulated using the model.

It seems that, starting with a situation where all 3-month Treasury bills have been replaced by 2-year Treasury notes, an increase in the proportion of outstanding 3-month Treasury bills relative to 2-year Treasury notes reduces both average interest payments and the annual variability of interest payments. The effect on cost stems from the slope of the yield curve between 3-month maturities and 2-year maturities, which is positive on average. The effect on the variability of interest payments is less intuitive because

¹³ Here, the Cost-at-Risk is the amount that interest payments will be less than with a probability of 90% (in 10 years). For example, Figure 12 shows that when average debt maturity is 7.5 years, interest payments are less than 2,98 GDP points in 90% of the cases.

issuing short-term debt increases refinancing risk. The results of the simulations show that this effect is dominant once the proportion of 3-month Treasury bills is greater than that of 2-year Treasury notes. A limited amount of 3-month Treasury bills seems to help smooth the annual interest payments by preventing the “bumps” caused by renewal of 2-year Treasury notes¹⁴.

IV- Sensitivity of results to model assumptions

We carried out further simulations after altering some of the model assumptions to test the robustness of the results. In this section we examine the influence of the following elements:

The scale of simulated macro-financial shocks (standard deviation),

The form of the distribution used to draw a sample of the shocks at random (normal distribution or “fat-tail” distribution, see Box 2),

The parameters of the yield curve (with a constant 3-month to 10-year slope),

A change in the average slope of the yield curve,

Estimation period of the model,

The number of lags used in the VAR

IV-1 The scale of the macro-financial shocks

In these simulations, we increase the standard deviations of the shocks affecting the macroeconomic variables by 20% compared to the values used in the baseline model. This reflects greater macro-financial volatility in coming years.

The results of these simulations are shown in Figure 16. There is very little change in terms of cost. The qualitative results with regard to risk are not affected but they are quantitatively “expanded”. For example, a reduction in average debt maturity leads to a 450m-euro increase in average annual variability of interest payments under the baseline scenario, but in this scenario, the increase is some 600m euros.

IV-2 The distribution types of the shocks

In the central model, the sample of shocks is drawn from the estimated residual values (bootstrap technique, see CBO, 2001 or Bardoux et al., 2006). This means that it is implicitly assumed that “historical” shocks cover their defining set for the historical estimation period in a way that is representative of their distribution. One drawback of this method is that it excludes extreme shocks, if there are none during the estimation period. In other words, if we draw a sample of residual values from the historical values, no shock can be greater than the estimated shocks. On the other hand, if the sample of shocks is drawn following a normal distribution, there is a non-null probability that a shock drawn may be greater than any value, no matter how large.

¹⁴ Let’s assume that a given amount needs to be financed in year 0. If it is financed with 2-year Treasury notes, the annual interest payments will be fixed for the first two years, but the interest rates will have time to fluctuate substantially between two refinancings. This could potentially give rise to strong variability of interest payments between year 2 and year 3. On the other hand, if the same amount is financed with 3-month Treasury bills, interest rates will fluctuate less between refinancings (if they are not too volatile) and average interest payments over the period would be smoother on average over the period, despite the financing risk, which is higher than in the previous case.

We have produced two variants here. In the first, the simulations are carried out on the assumption that the samples of shocks are drawn following a normal distribution, and the covariance matrix is derived when the VAR is estimated. However, there are an infinite number of different distributions providing the same covariance matrix, and the normal distribution is only one of them. More specifically, there are distributions that give a greater probability to extreme phenomena, which has to be counterbalanced by a smaller probability of moderate and intermediate shocks, if we want to keep the standard deviation constant. These distributions are said to have “fat tails” (see Box 2). The second variant consists of drawing a sample of shocks following this type of distribution.

Figures 17 and 18 suggest that there are only marginal changes in the results with regard to interest payments and average annual variation. On the other hand, even though the Cost-at-Risk measures are qualitatively the same as those in the baseline simulations, they are higher overall when the sample of shocks is drawn following a normal distribution or a fat-tail distribution.

Box 2 – Fat-tail distribution

In the baseline simulations, we assume that the sample of shocks is drawn uniformly from the empirical distribution plotted by the estimated shocks (bootstrap technique). In other words, the characteristics of the simulated shocks, such as the standard deviation, are exactly the same as those of the estimated shocks.

New simulations have to be made to test the sensitivity of the results to this choice. The first example considers normal distributions showing the same standard deviations and variance-covariance matrix as the estimated shocks. Naturally, normal distributions are only one example of distributions and there are an infinite number of distributions showing the same standard deviation.


One potentially important characteristic of the distribution used for this simulation exercise is the *kurtosis*. Intuitively, this characteristic accounts for the probability of extreme phenomena occurring relative to the occurrence of “average” phenomena.

Such extreme events do occur on financial markets. For example, the New York Stock Exchange varied by more than 7% on 48 occasions between 1916 and 2003, whereas, if stock prices followed a Gaussian distribution, such events would occur only once every 300,000 years.

Two densities may show the same mean and variance, but different *kurtosis*. The chart below shows a normal density compared to a density with the same standard deviation ($=1$), but with kurtosis equal to 9 (the kurtosis of a normal distribution is equal to 3)¹⁵.

¹⁵ This density with a kurtosis equal to 9 is the density followed by a random variable X resulting from the following procedure: a random variable P is drawn at random following a uniform distribution $[0,1]$. If P is smaller than a given probability p , X is drawn following a normal distribution $N(-x_0, \sigma_2)$; if P is between p and $1-p$, X is drawn following a normal distribution $N(0, \sigma_1)$ and if P is higher than $1-p$, X is drawn following a normal distribution $N(x_0, \sigma_2)$. The parameters p , x_0 , σ_1 and σ_2 are calibrated so that the standard deviation of X is 1.

Figure 8 – Fat-tail distribution




IV-3 The parameters of the yield curve

The yield curve in the model has been extrapolated from the 3-month interest rate and the 10-year interest rate. For each date, the other segments of the curve are calculated from these two rates using a function that depends on two parameters denoted β_2 and τ (see Annex 1.C). These parameters are set to achieve the best fit for a given period in the past and then they are kept constant during the simulation period. Figure 9 shows the influence of these parameters on the average yield curves. Each curve shown in the figure corresponds to a given β_2 and τ pairing. Each curve is then plotted for the average values of the rates (4.5% for the 3-month rate and 5.5% for the 10-year rate).

For example, simulations were made with the pairing of $\beta_2 = -3$ and $\tau = 1$, versus $\beta_2 = 1.97$ and $\tau = 1.72$ in the baseline model (see Figure 19).

In this configuration of the yield curve, the 2-year rate is lower on average than the 3-month rate, which is not a realistic assumption. It follows that reducing average debt maturity by increasing the proportion of outstanding 3-month Treasury bills would increase average interest payments. Nevertheless, the influence of a reduction of average debt maturity achieved by reducing the proportion of outstanding 10-year Treasury bonds, with the same allocation of 3-month Treasury bills and 2-year Treasury notes, does not produce any qualitative change in the results.

Figure 9 – Average yield curves for different yield curve parameters




IV-4 Change in the average slope of the yield curve

The “baseline” simulations are made using the initial conditions, where the slope is positive, rising by 1 point between the short-term interest rate and the 10-year rate. This condition is satisfied on average over the simulation period.

Here, we analyse the performances of the swaps programme under the apparently more adverse assumption of decrease in the average slope of the yield curve over the simulation period. We test two cases: the average slope of the yield curve is 25 basis points in the first case and 50 basis points in the second case on average, as opposed to 100 basis points in the baseline model. If we assume that the stationary value of short-term rates remains at 4.5%, then the stationary values of long-term rates are 4.75% and 5.0% respectively. Figure 10 shows the corresponding yield curves¹⁶. The average yield curve of the baseline model is shown as a dotted line for the purposes of comparison.

Figure 10 – Average yield curves



¹⁶ In both cases, we used $\beta_2 = 0$ in the yield curve parameters. If the value of β_2 used is that in the baseline model (-1.97), then the interest rate on 2-year Treasury notes is lower on average than the rate on 3-month Treasury bills, which is not the case over a long period.

Figures 20 and 21 show the cost/risk planes obtained after making simulations with each of the models. Naturally, the expected gains from a reduction in average debt maturity in terms of lower average interest payments are smaller in these cases. If the long-term interest rates are at 4.75% or 5.0%, the respective gains per year are 0.20bn euros or 0.45bn euros on average for a 1-year reduction in average debt maturity. These gains are to be compared to the gain of 1bn euros in the baseline model. This means that the average gain is approximately proportional to the spread between 3-month rates and 10-year rates.

The qualitative results are practically the same with regard to volatility measures. On the other hand, the Cost-at-Risk measures are less favourable for a reduction of average debt maturity. Generally speaking, the average interest payment and Cost-at-Risk measures are naturally lower than in the baseline model because the stationary long-term rates are lower in these simulations. In the baseline model, the stationary long-term interest rate is 5.5%, as opposed to 4.75% and 5.0% in the two models tested here.

IV-5 Estimation period of the model

The baseline model was estimated for the period from 1986Q1 to 2007Q1. How would the results be affected if the estimation period were shortened? This corresponds to the assumption that the most recent years, rather than the last 20 years as a whole, are more representative of the future dynamics of macroeconomic variables. Indeed, it could be argued that the French economy underwent major structural changes in the recent period, with deregulation, globalisation and the changeover to the euro.

Therefore, we re-estimated the model for the period from 1995Q1 to 2007Q1. It was then used to simulate new scenarios. The results are shown in Figure 22.

The main change in the results stems from the impact on the Cost-at-Risk linked to interest payments. In the baseline model, a reduction in average debt maturity increases the Cost-at-Risk linked to interest payments, but, in this model, the Cost-at-Risk is reduced. This means that the shortening of the estimation period results in a stronger argument for reducing average debt maturity. These simulations suggest that we can obtain lower cost and risk, if risk is measured in terms of the Cost-at-Risk linked to interest payments.

However, the estimation covers only 49 quarters, which is a small number for the proper estimation of a VAR model. Furthermore, given that the simulations cover the next 10 years, using only the last 10 years as a reference period is no doubt unsatisfactory. Consequently, the results of the simulations must be interpreted with care.

IV-6 The number of lags used in the VAR

The current variable values in the model used to represent economic changes depend on the past or “lagged” values of the same variables (see Box 3).

The baseline model uses two lags, or two quarters, for the VAR model. This choice is based on the results of econometric tests. Yet, the tests used did not produce a consensus about the choice of the two lags. More specifically, 2 of the 5 tests recommended the use of a single lag (see Annex 1.A).

Figure 23 shows the simulation results when the VAR model is estimated with a single lag. The results show little change, which confirms the robustness of the results.

V- Back-testing

Back-testing is a quantitative method in which the performances that would have been obtained if the different strategies had been applied in the past are observed. In other words, back-testing attempts to answer

the question, “what would have happened if this or that strategy had been implemented over the last n years?”

It is important to understand that the ranking of the strategies produced by back-testing is less relevant as a decision-making aid than the ranking obtained in section III. As mentioned above, just because a strategy performs satisfactorily for a given deterministic past scenario, it does not mean that it will always produce satisfactory results in a future period, where economic changes are random. However, back-testing is still helpful, if only because it enables us to see if the results are of the same order of magnitude as the results suggested by the random simulations.

The strategies tested differ only with regard to their relative proportions of 10-year Treasury bonds and 2-year Treasury notes in order to simplify interpretation of the results; we do not examine the influence of the allocation between 2-year Treasury notes and 3-month Treasury bills. The average debt maturities associated with the different strategies range from 5.6 years to 7.4 years.

Table 1 – Back-testing results

	Average debt maturity	Performance measure	1985-2006	1990-2006	1995-2006	2000-2006
Strategy 1	5.6 years	Average interest payments (% of GDP)	2.37	2.38	2.09	1.87
		Average variation (% of GDP)	0.16	0.18	0.13	0.04
Strategy 2	6.0 years	Average interest payments (% of GDP)	2.45	2.45	2.18	1.93
		Average variation (% of GDP)	0.16	0.18	0.11	0.04
Strategy 3	7.0 years	Average interest payments (% of GDP)	2.63	2.62	2.41	2.07
		Average variation (% of GDP)	0.16	0.18	0.11	0.04
Strategy 4	7.4 years	Average interest payments (% of GDP)	2.69	2.68	2.49	2.12
		Average variation (% of GDP)	0.16	0.18	0.11	0.04

Key: Applying Strategy 2, with average debt maturity of 6.0 years, over the period from 1995 to 2006 would have resulted in average annual interest payments equal to 2.18% of GDP.

Table 1 shows the results of simulations made over several periods¹⁷. In each of the 4 periods considered, reducing average debt maturity from 7 years to 6 years decreases the average annual interest payments by

¹⁷ The initial debt structure at the start of each simulation, before entering into swaps to achieve the desired average debt maturity, is the prevailing structure at the first quarter of 2007. More specifically, the debt structure by redemption amounts is expanded to correspond to the debt/GDP ratio on the first date considered (e.g. 30.3% in 1985). Furthermore, the matrix of future interest payments as seen from the first quarter of the simulation is equal to the matrix for the first quarter of 2007. It would have been better to use the exact debt structure from the first quarter of the simulation, with the redemption matrices and rates of the first quarter considered, but such data are hard to construct for past years. Consequently, it is impossible to compare the average interest payments as shown in Table 1 with those observed historically.

between 0.14 GDP points, for the period from 2000 to 2006, and 0.23 GDP points for the period from 1995 to 2006. These values are in line with the results of the stochastic simulations that suggest a decrease of 0.13 GDP points in the annual interest payments following a similar reduction in average debt maturity.

Table 1 also shows that, for the periods considered, reductions in average debt maturity would have had very little impact on the variability of interest payments.

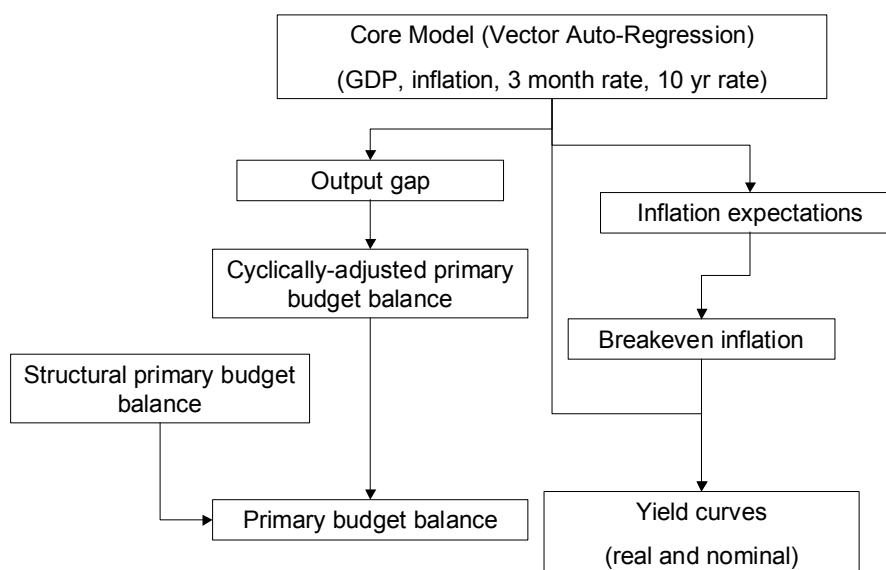
References

- Bardoux, F., D. Cintioli, and M. Marchiori, 2006, Risque de portefeuille, Value-at-Risk, les atouts de la simulation historique, *Asset Management*, December 2005 - January 2006.
- Benassy-Quéré, A., L. Boone and V. Coudert, 2003, *Les taux d'intérêt*, Ed. La Découverte.
- Bergström, P. and A. Holmlund, 2000, A simulation model framework for government debt analysis, *mimeograph*, The Swedish National Debt Office.
- Bergström, P., A. Holmlund and Lindberg, 2002, The SNDO's simulation model for government debt analysis, *mimeograph*, The Swedish National Debt Office.
- BIS, 2007, Institutional investors, global savings and asset allocation, *CGFS Papers No. 27*, 2007.
- Black, R. and C. Telmer, 2001, Liability Management Using Dynamic Portfolio Strategies, *Algo Research Quarterly* vol. 2.
- Bolder, D., 2002, Towards a More Complete Debt Strategy Simulation Framework, *Bank of Canada Working Paper No. 2002-13*.
- Bolder, D. 2003, A Stochastic Simulation Framework for Debt Strategy, *Bank of Canada Working Paper No. 2003-10*.
- Cappiello, L. and S. Guéné, 2005, Measuring Market and Inflation Risk Premia in France and in Germany, *ECB Working Paper Series No. 436*.
- CBO, 2001, Uncertainty in Social Security's Long-Term Finances: A Stochastic Analysis, Congress of the United States, Congressional Budget Office.
- Coupet, M., 2006, Révisions de la croissance potentielle de l'économie française à moyen-long terme [Translated in English: "Estimates of French medium- to long-term potential growth revised"] *Trésor Eco* No. 2, 2006.
- Danish Nationalbank, 2006, *Danish Government Borrowing and Debt*, Chapter 9, Management of the Central Government's interest rate risk.
- Dutch State Treasury Agency, 2007, Risk Management of the National Debt: Evaluation of the 2003-2007 Policy & 2008-2011 Policy, September 2007, Ministry of Finance.
- European Commission, 2002, *Public Finances in EMU*. No. 3/2002.
- Hamilton, J. 1994, *Time Series Analysis*, Princeton University Press.
- Hodrick, R. and E. Prescott, 1997, Postwar U.S. Business Cycles: An Empirical Investigation, *Journal of Money, Credit, and Banking*.
- IMF, World Bank, 2003, *Guidelines for Public Debt Management*.
- Laubscher, N. F., 1959, Note on Fisher's Transformation of the Correlation Coefficient, *Journal of the Royal Statistical Society*.
- Nelson, C. and A. Siegel, 1985, Parsimonious Modeling of Yield Curves for U.S. Treasury Bills, *NBER Working Paper No. 1594*.
- Pick, A. and M. Anthony, 2006, A simulation model for the analysis of the UK's sovereign debt strategy, *UK DMO research papers*.
- Rennison, A., 2003, Comparing Alternative Output Gap Estimators: A Monte-Carlo Approach, *Bank of Canada Working Papers No. 03-8*.
- Sack, B. and R. Elsasser, 2004, Treasury Inflation-Indexed Debt: A Review of the U.S. Experience. *FRBNY Economic Policy Review*, May 2004.
- Sims, C. A., 1980, Macroeconomics and Reality, *Econometrica* vol. 48.
- Stock, J. H. and M. W. Watson, 2001, "Vector Autoregressions", *The Journal of Economic Perspectives*, vol. 15.
- Van den Noord, P., 2000. The size and role of automatic stabilizers in the 1990s and beyond, *OECD Economics Department WP No. 230*.
- Velandia, A., 2002, A Risk Quantification Model for Public Debt Management. *World Bank working paper*.

Annex 1 – The Model

Figure 11 shows the different “blocks” that make up the model as a whole. First, the core model defines the dynamics of the four variables: real GDP, inflation, the 3-month interest rate and the 10-year interest rate. Secondly, the auxiliary models describe the primary budget balance, using the output from the central block. Thirdly, while the core model describes changes in two interest rates (3-month and 10-year), other models help complete the nominal and real yield curve. This Annex provides more detail about these three points.

Figure 11 – Model Structure



Annex 1.A Core Model (VAR)

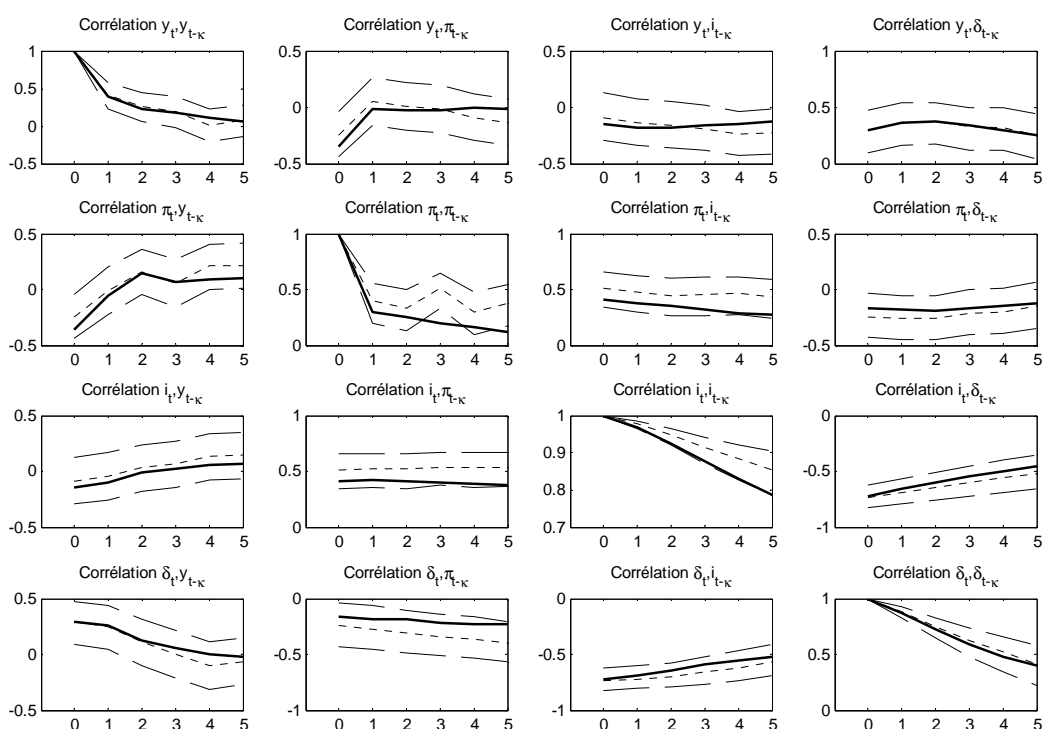
This model is to be used for simulating “foreseeable” macro-financial trajectories in the coming years. Consequently, it should be estimated over a period that is consistent with the time horizon under consideration. More specifically, if the horizon is 10 years out, the estimation period should be at least as long. In addition, it should cover a period that includes a whole number of business cycles to prevent bias that could skew the results in one direction or the other. The notion of the output gap is used to give the approximate dates of business cycles. More specifically, the output gap is obtained by filtering the real GDP series with a Hodrick-Prescott filter¹⁸, which is widely used for this purpose (see e.g. Rennison, 2003). Ultimately, the period from 1986Q1 to 2007Q1 was selected.

The core model is based on VAR (vector auto-regression, see Box 3) modelling. Four variables feature in the core model: GDP growth (y_t), inflation (π_t), the 3-month interest rate (i_t) and the 10-year interest rate ($i_t + \delta_t$). This type of modelling is widely used because it is simple and it can capture correlations between the different variables. Figure 12 uses examples to show that the model correctly reproduces the historical correlations between the four variables from 1986Q1 to 2007Q1¹⁹.

¹⁸ Hodrick and Prescott (1997).

¹⁹ It shows that most of the simulated correlations fall within the 90% confidence intervals of the historical correlations.

Figure 12 – Comparison of simulated and historical simulations



Key: Each of the 16 figures shows the correlation of one of the four variables considered in the VAR (GDP growth, inflation, short-term rates and 3-month/10-year spread) with one of the four lagged series (κ is the number of lags, in quarters). The dashed lines are the bounds of the 90% confidence interval of the correlations measured in the historical time series. Consequently, the presence of a thick line, showing the correlations measured in the simulated data, between the two bounds indicates good model performance.

The data were taken from the *Feri* database. Interest rate data were taken from the Banque de France and the GDP data and consumer price index data were taken from INSEE²⁰. The structural primary budget balance series comes from the OECD *Economic Outlook* database.

It is relatively simple to use the maximum likelihood method to estimate an unrestricted VAR, meaning a VAR with no restrictions on the variables, other than the choice of the number of lags. This is because it only involves using least ordinary squares to estimate the parameters of each of the VAR equations for each endogenous variable one by one.

²⁰ Real GDP is obtained by using the consumer price index to deflate GDP at current prices. GDP growth is calculated by the quarterly change in the log of real quarterly GDP. Inflation is the quarterly change in the log of the consumer price index.

Box 3 –VAR Modelling

We chose to model the joint dynamics of the four core variables (GDP growth, inflation, 3-month interest rate and 3-month/10-year spread) as a VAR (vector auto-regression). As Hamilton (1994) suggests, this type of modelling is widely used to analyse economic systems. The work of Sims (1980) contributed actively to the development of this methodology. This work suggests that VARs are fairly robust to the Lucas Critique²¹. For Stock and Watson (2001), this methodology is still one of the most effective for analysing and forecasting time series.

Let X_t be the endogenous variable vector (in this case, the variables are quarterly GDP growth, quarterly inflation, the 3-month interest rate and the 3-month/10-year spread). VAR modelling explains each element of this vector by its lags and those of the other variables selected in the VAR, plus a random innovation to the variables. The model, which is called VAR(p) when p lags of the endogenous variables appear in the equations, can be written:

$$X_t = c + \sum_{i=1}^p A_i X_{t-i} + \varepsilon_t \quad (*)$$

where the dimensions of the A_i matrices are 4 x 4 and the innovations ε_t are independent and identically distributed following a normal distribution $N(0, \Sigma)$.

By using L to denote the lag operator, equation (*) can be rewritten (with $A_0 = c$)

$$X_t = \left(\sum_{i=0}^p A_i(L) \right) X_t + \varepsilon_t,$$

If $B(L)$ is the inverse of the polynomial $I - A(L)$, we get the representation of the infinite moving average of equation (*):

$$X_t = \sum_{i=0}^{\infty} B_i(L) \varepsilon_{t-i}.$$

We deduce the unconditional averages and variances of vector X from this equation:

$$Moy(X) = B_0 \quad Var(X) = \sum_{i=0}^{\infty} B_i \Omega (B_i)^T.$$

The number of lags selected is two, as suggested by the Sequential Modified Probability Ratio Test, the Final Prediction Error Test and the Akaike Information Criterion Test. The Schwarz Information Criterion and the Hanna-Quinn Information Criterion pointed to the selection of one lag. However, the latter selection results in residual autocorrelation, revealed by Portmanteau Tests and Lagrange Multiplier Tests. Therefore, we selected two lags for the baseline model, because these two types of tests suggested that there is no residual autocorrelation in this case. Yet, alternative simulations were made using a model incorporating a VAR with only one lag in order to test the robustness of the results of the lag selection tests (see section IV-6).

We conducted Chow Break-Point Tests for each of the four univariate equations of the VAR. We selected two break-points for these tests. The first is in the middle of the sample period (1995Q1) and the second

²¹ A model is subject to the Lucas Critique if its equations are no longer valid when agents change their behaviour. Take the example of a model estimated on a given period in the past. Such a model may not be suitable for assessing the impact of a future tax measure, since agents may alter their behaviour after the measure is introduced and this change may not be reflected in the equations estimated on the basis of past data, thus creating a bias in the assessment.

corresponds to the establishment of the ECB (1998Q2). Overall, the tests point to the stability of the model parameters²². In addition, the constants of the model equations were modified to give the simulated variables the following values on average over the simulation period, which can be seen as their long-term values: 2% annual real GDP growth, 2% annual inflation, a 4.5% 3-month interest rate (annualised), and a 1% spread between the 10-year rate and the 3-month rate. The real GDP growth rate of 2% corresponds to the potential long-term growth rate estimated by the Treasury and Economic Policy Directorate General (see Coupet, 2006). The 2% inflation rate corresponds to the upper bound of the ECB's target, the short-term rate of 4.5% corresponds approximately to the average rate over the last 15 years, as is the case for the 10-year/3-month spread too²³.

One of the reasons for the decision to use this procedure for modifying the unconditional averages of the series was the observation that the unconditional averages, unlike the correlations between variables, are incorrectly estimated in finite samples²⁴.

Annexe 1.B Primary budget balance

The cyclically-adjusted primary budget balance is modelled as a simple function of the output gap²⁵. We can use an assessment of the elasticities of fiscal revenue and expenditure to cyclical variations (see Van den Noord, 2002), to deem that the proportion of the primary budget balance linked to fluctuations in the business cycle (cyclically-adjusted budget balance) is equal to half the output gap (expressed in GDP points). This model is a generally-accepted approximation in Europe (see European Commission, 2002). The structural annual budget balance is modelled as an auto-regression process with one lag, or AR(1). This simple model is consistent with the small number of observations available (36 years from 1971 to 2006). In addition, the estimation results suggest the robustness of the estimation (Durbin Watson Statistic of 1.7, normally distributed residuals according to the Jarque-Bera Test and no residual autocorrelation shown by the Breusch-Godfrey Test).

Furthermore, the constant of the auto-regression process is modified to achieve approximate stabilisation of the budget balance expressed in GDP points over the period²⁶.

We need to have a measure of potential output in order to obtain the output gap. Potential output is obtained by filtering the GDP series. We use the Hodrick-Prescott Filter for this purpose. The smoothing parameter selected is 1,600, which is the value that is usually selected for filtering quarterly data.

²² Out of the 24 Break-Point Tests (4 equations x 2 break-points x 3 statistics per test) only one rejects the null hypothesis of no break-point.

²³ By way of comparison, the historical averages of the 3-month rate and the spread over the estimation period of the VAR (1986Q1 to 2007Q1) are 5.70% and 0.85% respectively. That being the case, the resulting 6.55% level of the 10-year rate seems high with regard to the long-term value suggested by the Ramsey Rule, for example. Benassy-Quéré, Boone and Coudert (2003) arrive at a value of 3.6% for the real long-term rate, which works out to a nominal rate of 5.6% with average inflation at 2%.

²⁴ For this purpose, we simulated a large number n of macroeconomic trajectories over 20-year periods, using the model estimated from historical data. We then re-estimated the VAR for each of the n trajectories. This gave us the distributions of the n unconditional averages, which are calculated on the basis of the matrices from the estimation of the VAR (see Box 3). These distributions show large standard deviations, which suggests that they cannot be estimated satisfactorily. Consequently, it seems preferable to select long-term values that are not estimated.

²⁵ The output gap is the difference between GDP and potential GDP expressed as a percentage of GDP.

²⁶ Using s to denote the primary budget balance and g to denote outstanding debt expressed in GDP points, we show, for example, that the stabilising primary budget balance, meaning the balance that is compatible with a constant debt-to-GDP ratio, stands at approximately $dx(g-r)$ where r is the average nominal interest rate paid on debt and g is the nominal GDP growth rate. For example, if r is greater than g , a primary budget surplus is required to stabilise the debt-to-GDP ratio.

Annexe 1.C Yield curves

The VAR only models two points on the yield curve: the 3-month interest rate and the 10-year rate. The rates for the other points on the nominal yield curve are obtained as a function of these two rates for each simulated period (see Box 4).

Figure 13 shows an example of a simulation over 20 periods, or 5 years, in the two lower quadrants. It shows how the model makes it possible for different forms of yield curves to occur, including inverted yield curves. On average, we calculate that the yield curve is inverted, meaning that the 3-month rate is higher than the 10-year rate, in 21% of the simulated periods. This is a lower percentage than the one mentioned in the introduction, since short-term rates were higher than the 10-year rates in 27% of the quarters. However, this percentage can be largely explained by the period leading up to the EMS crisis in the early 1990s. At the time the Banque de France had to increase its short-term rates to maintain the parity of the franc against the Deutschmark. The German currency was boosted by hikes in the Bundesbank's key rates aimed at choking off the inflationary pressures stemming from reunification. This type of mechanism is unlikely to reoccur in the next 10 years, since the euro is a floating currency. This means it would be better for the simulated period to include fewer periods with inverted yield curves than the historical period mentioned above²⁷.

Box 4 – Yield curve model

The VAR simulates four variables: GDP, inflation, the 3-month interest rate and the spread between 10-year rates and 3-month rates. An additional model is required to obtain the entire yield curve.

It has been firmly established by empirical work that most of the variability of interest rates over time stems from only two of the factors that characterise the yield curve: its level and its slope. In other words, if we have a measure of the level of the yield curve and another measure of its slope, we can derive the whole yield curve with a very small mean error for all the different rates.

We use this characteristic in our model to postulate a parametric form for the yield curve. We then estimate the parameters to minimise the differential between the modelled rates and the observed rates for 12 benchmark maturities between 3 months and 30 years over the last 15 years²⁸.

More specifically, the parametric form selected is a Laguerre polynomial. Nelson and Siegel (1985) use this form in their approach. However, Nelson and Siegel were modelling the yield curve for zero-coupon bonds with a Laguerre polynomial, whereas we are modelling yields on bonds with coupons. The equation giving each rate for maturity m at a date t is written:

$$y_{t,m} = \beta_{0,t} + \beta_{1,t} \exp\left(-\frac{m}{\tau}\right) + \beta_{2,t} \exp\left(-\frac{m}{\tau}\right) \exp\left(-\frac{m}{\tau}\right) \quad (\Delta)$$

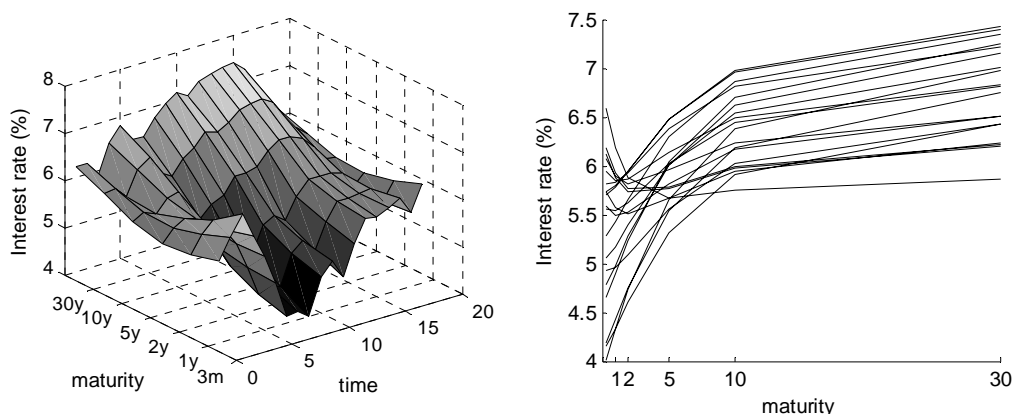
Once the parameters τ and β_2 have been selected and we know the observed 3-month rate and 10-year rate, we can calculate the two factors that vary over time ($\beta_{0,t}$ and $\beta_{1,t}$). The entire yield curve at the date t is then deduced from (Δ).

²⁷ In the present exercise, putting the percentage of periods with inverted yield curves at 21% could even be deemed conservative, meaning that 21% is probably an overestimation of the percentage of periods with inverted yield curves in the next 10 years.

²⁸ The maturities selected are 3 months, 6 months, 1 year, 2 years, 3 years, 4 years, 5 years, 7 years, 10 years, 15 years, 20 years and 30 years. The estimation is based on monthly data for benchmark government securities (source: Bloomberg).

After determining the parameters τ and β_2 that give the best fit, we can use this method to explain (on average over the 12 maturities selected) more than 98% of interest rate variance over the last 15 years.

Figure 13 – An example of yield curve estimation



Key: The two charts present the same simulations. The first chart shows the distortion of the yield curve over time and the second plots all of the yield curves on the same plane.

The method explained in Box 4 concerns the construction of *nominal* yield curves for each period based on the 3-month and 10-year interest rates. Yet, the model also considers inflation-indexed securities. The real yield curve is obtained by subtracting the breakeven inflation rate from the nominal rates. The breakeven inflation rate incorporates agents' inflation expectations, plus a term including the inflation risk premium and other, second order, terms reflecting various differences between an inflation-indexed bond and a nominal bond (see Sack and Elsasser, 2004).

We assume that economic agents' inflation expectations are based on knowledge of the model underlying economic dynamics²⁹. We further assume that the inflation risk premium is constant at 20 basis points³⁰.

²⁹ For details about the formation of expectations using a VAR model, see Hamilton (1994), Chapter 11.

³⁰ Much empirical research is aimed at assessing the inflation risk premium. A recent paper (Cappiello and Guéné, 2005) suggests that the premium is relatively stable over time and that its order of magnitude in France is 20 basis points.

Annex 2 – Risk measures

This annex presents the indicators selected to assess the performances of the different strategies.

After applying different strategies in N scenarios (of n years) simulated using a macro-financial model, we obtain the same number N of interest payment time series (over $4xn$ quarters) for each strategy. We then carry out two transformations before extracting the performance measure data for the financing strategies.

First, the quarterly interest payment time series need to be transformed into annual data, which are better for assessing debt management performance. Secondly, the annual interest payment amounts are divided by the corresponding annual GDP figures.

We use c_t^j to denote the interest payments for the year t (expressed as a % of GDP) with a given strategy under scenario j . The measure of the average interest rate payment up to the year t under scenario j , \bar{c}_t^j is given by:

$$\bar{c}_t^j = \frac{1}{t} \sum_{i=1}^t c_i^j .$$

Then, the cost measure associated with the strategy at horizon t , after simulating the N scenarios, is given by:

$$\bar{c}_t = \frac{1}{N} \sum_{j=1}^N \bar{c}_t^j$$

(Consequently, the cost measure is the average of the averages of the N scenarios over t years.)

There are many risks that can be considered (see Bolder, 2003). We selected two different risk measures in the model. The first, which can be qualified as an “instantaneous” risk, relates to the volatility of interest payments from year to year. The second is a cumulative risk and it relates to the dispersion of interest payments at a given time horizon.

The risk measure reflecting the volatility of interest payments is the mean standard deviation of the annual variations in interest payments calculated for each scenario. This measure reflects how much interest payments can vary from one year to the next. When it is measured up to a horizon t , it is denoted vol_t and calculated as follows:

$$vol_t = \frac{1}{N} \sum_{j=1}^N \sqrt{Var_{i \in \{1,t\}}(c_i^j - c_{i-1}^j)} \quad \text{with} \quad Var_{i \in \{1,t\}}(x_i) = \frac{1}{t} \sum_{i=1}^t x_i^2 - \left(\frac{1}{t} \sum_{i=1}^t x_i \right)^2$$

The risk measure relating to the dispersion of interest payments is conditional at a given horizon t . This measure is similar to the notion of *CaR* (Cost-at-Risk) used by the Danish central bank (Danish Nationalbank, 2006). It corresponds to the amount of interest payments that will be exceeded only with a 10% probability. Naturally, the higher this value, the riskier the strategy considered is deemed to be. We distinguish between absolute *CaR* and relative *CaR*. The latter measure corresponds to the former, minus the expected average interest payment. Absolute *CaR* at a horizon t is given by the following equation:

$$CaR_t(10\%) = Sup\{z \mid P(\bar{c}_t < z) \leq 90\%\}$$

Strategies may appear risky, according to one of the measures, and much less so, according to the other. For example, a strategy that tends to generate smooth interest payment time series, but with large fluctuations over time (large and infrequent variations or small and very frequent variations) will be deemed to be risky by the Cost-at-Risk measure and not risky by the volatility measure. This is illustrated by Figure 14.

Figure 14 – Examples of strategies that are deemed risky according to one measure but not the other




Figure 15 – Simulation results – Baseline Model




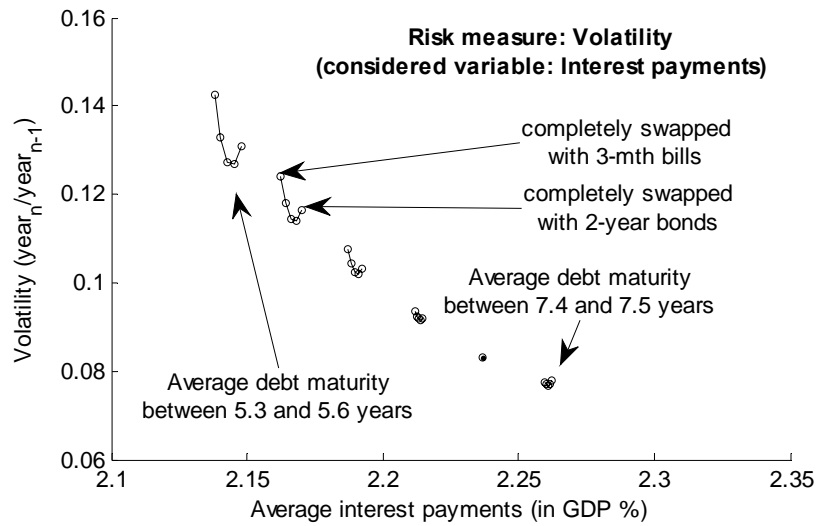


Figure 16 – Simulation results – Size of shocks







Figure 17 – Simulation results –Samples of shocks drawn following a normal distribution

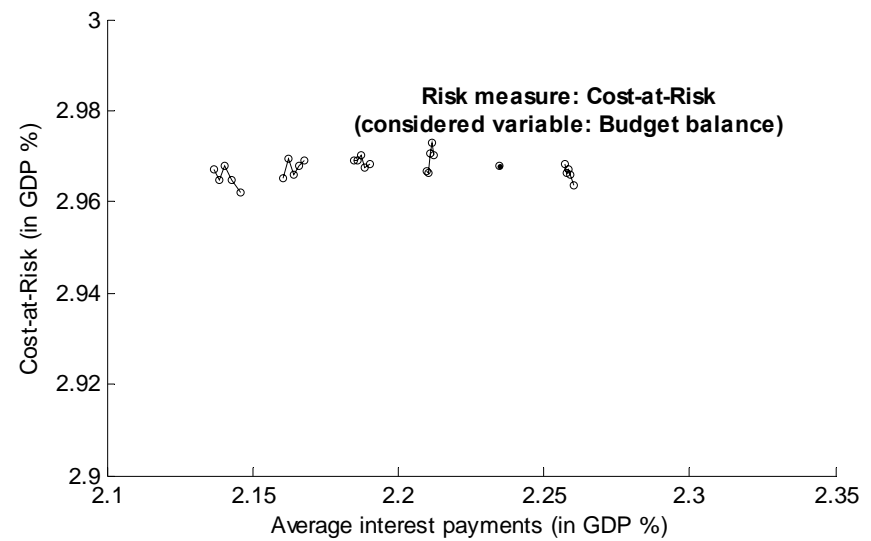
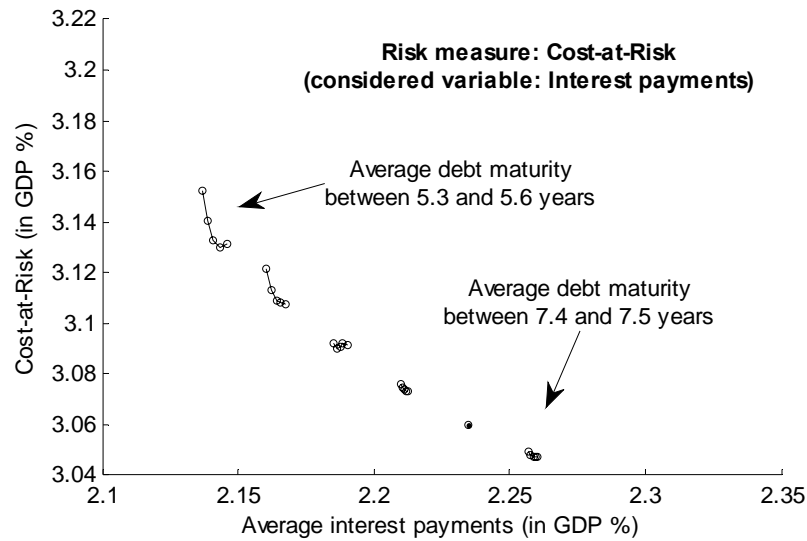




Figure 18 – Simulation results – Samples of shocks drawn following a fat-tail distribution


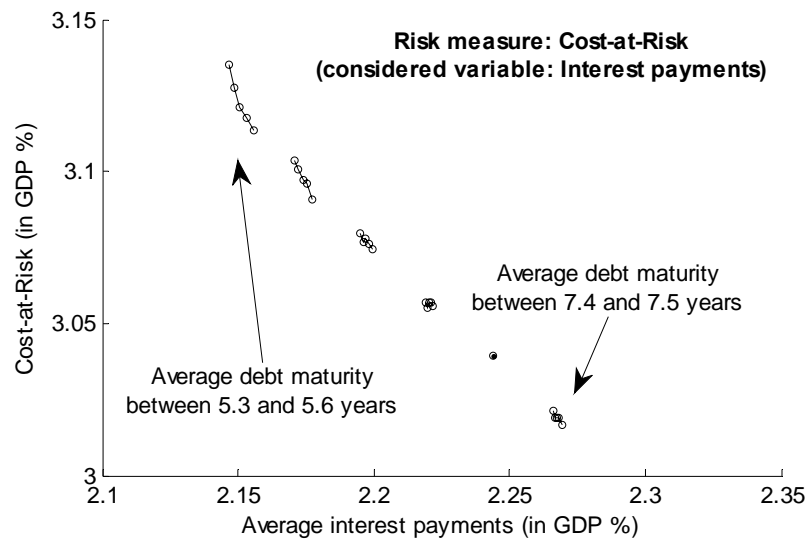

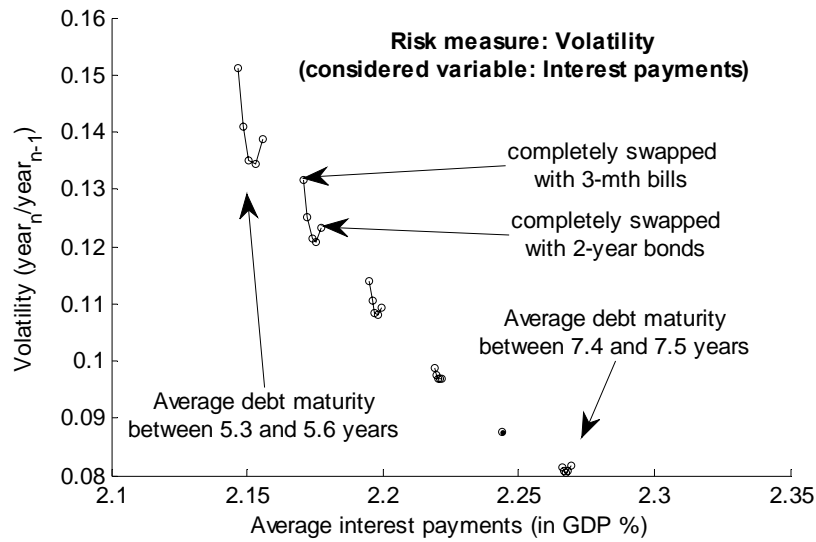


Figure 19 – Simulation results – Parameters of the yield curve

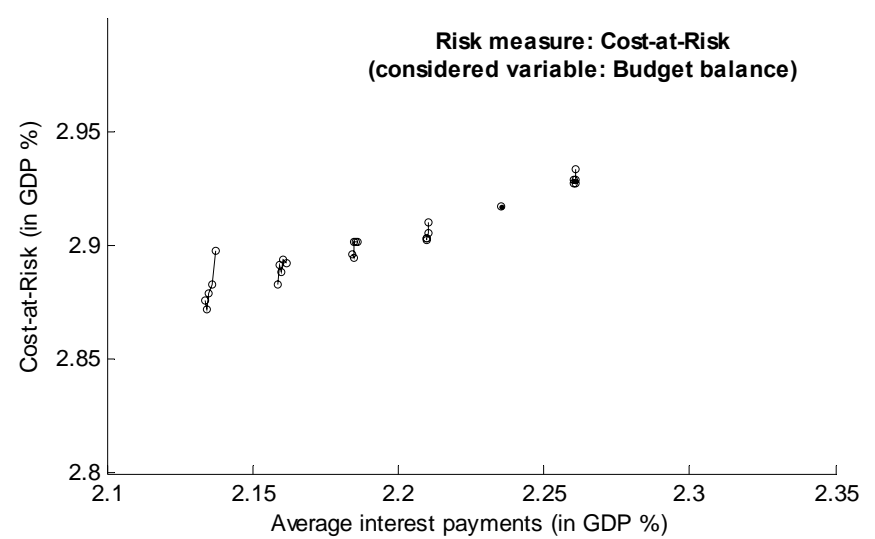
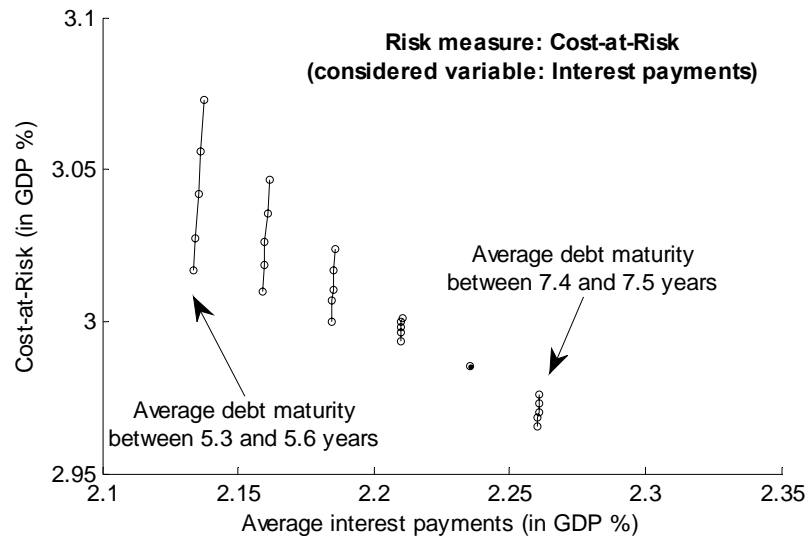




Figure 20 – Simulation results – Change in the slope of the yield curve (25 basis points)


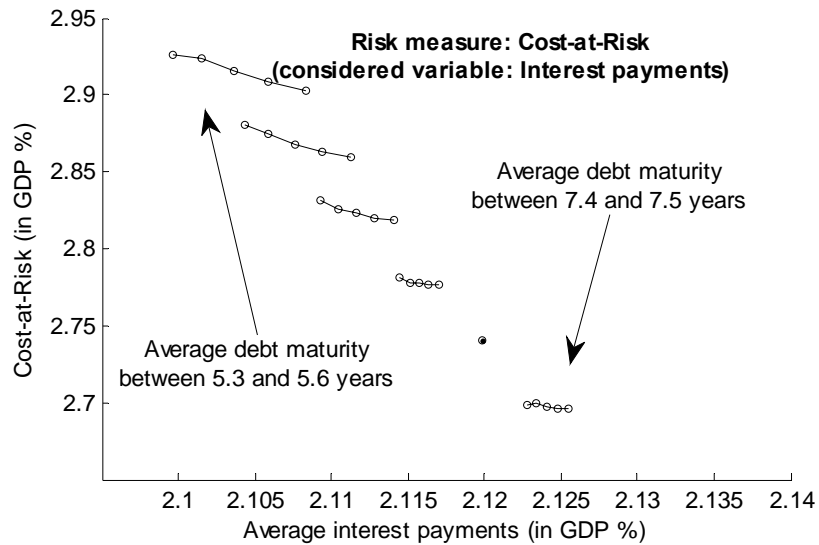




Figure 21 – Simulation results – Change in the slope of the yield curve (50 basis points)

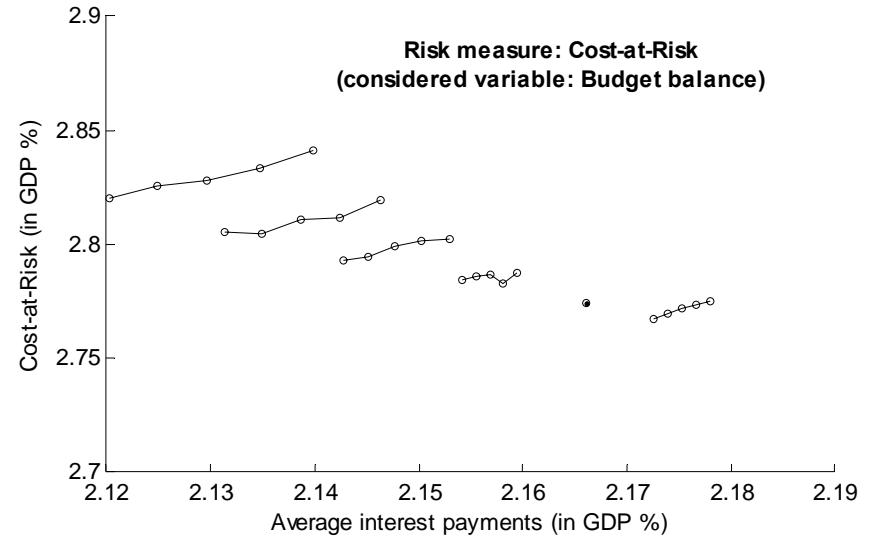

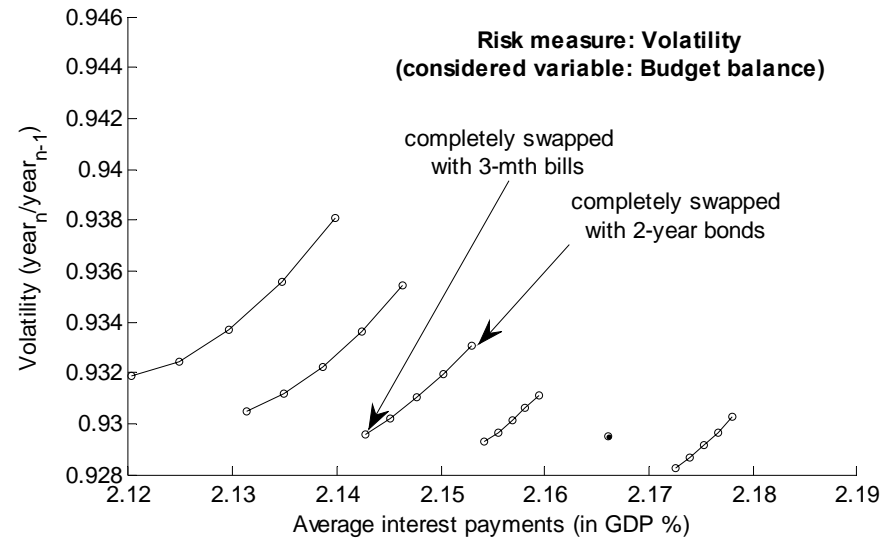



Figure 22 – Simulation results – Estimation period: 1995Q1 to 2007Q1


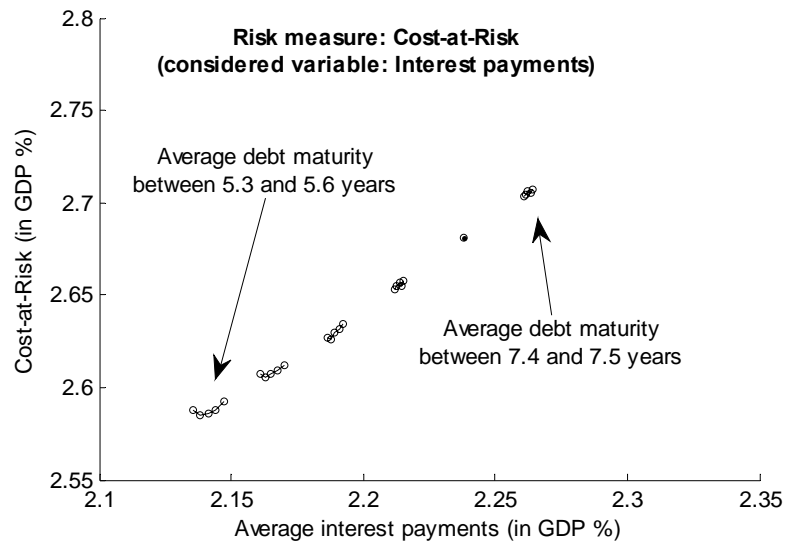

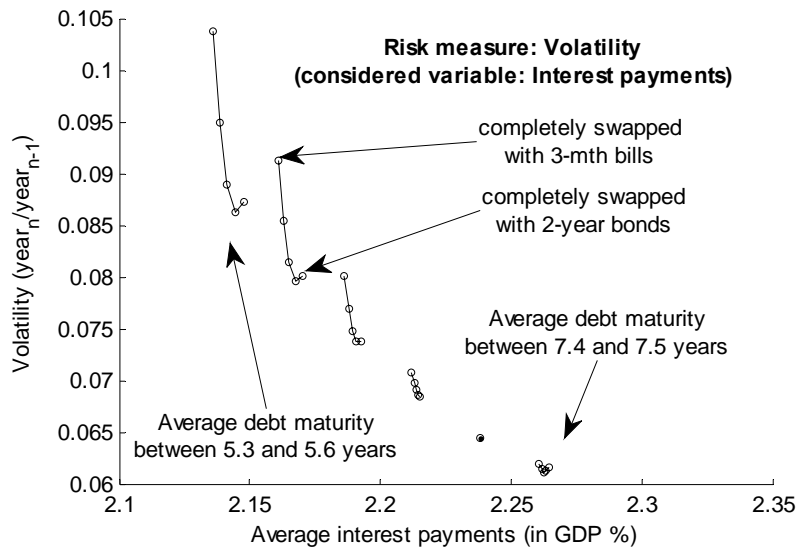


Figure 23 – Simulation results – Number of lags selected in the VAR: 1

