A Simulation Model Framework for Government Debt Analysis*

Pål Bergström & Anders Holmlund†

November 30, 2000

Abstract

This report describes a model that the Swedish National Debt Office has developed for the study of strategic decisions on the Swedish government debt portfolio. The framework is a Monte Carlo simulation model where different portfolio strategies are examined over a ten-year horizon. The objective is to find a portfolio which, in accordance with the objective of the Debt Office, seeks to minimise costs with due regard to risk. Different possible portfolio strategies, in terms of currency compositions and target durations, are confronted with modelled stochastic financial variables. The model described in this report is still under construction, and the specification is far from final. One important issue is therefore to describe the weaknesses of the present model and the different aspects where the model could and should be improved.

*This document has benefitted from discussions with, and comments from, several colleagues at the SNDO, most noteably Lars Hörngren. Comments from Greg Makoff and Bruno Momont, both Schroeder Salomon Smith Barney, London, are also gratefully acknowledged. Section 4 in this document is to a large extent identical to a skillful translation made by Victor Kayfetz, Scan Edit, based on the corresponding section in the Swedish guideline proposal documnet. All remaining errors and shortcomings are, however, our own.

†The authors are affiliated with the Swedish National Debt Office. To the extent that the text in this document coincides with the guideline proposal of the Debt Office, the views expressed in this report reflects the views of the debt office. Whenever the documents do not coincide, this document does not necessarily reflect these views. Please send any comments and suggestions to anders.holmlund@rgk.se.

1 Introduction

The purpose of this report is to describe the model that the Swedish National Debt Office has developed for studying the effects of the choice of duration and currency composition in the Swedish government debt portfolio. Neither claiming originality nor generality, we think this document could serve as a starting point for discussions with both academics and practitioners in the field of government debt management. This is especially important since studies in this field are fairly scarce[1]. The framework is a Monte Carlo simulation model where different portfolio strategies are examined over a ten-year horizon. The objective is to find a portfolio which, in accordance with the objective of the Debt Office, seeks to minimise costs with due regard to risk. Different possible portfolio strategies, in terms of currency composition and target durations are confronted with modelled stochastic financial variables. Put slightly differently, the model may be said to consist of two components: a stochastic simulation model that generates factors such as interest rates, exchange rates and the borrowing requirement, and a model that simulates changes in the debt portfolio as well as various costs associated with this, based on the simulated risk factors and an assumed borrowing strategy. These two models are described in separate sections below. Thereafter the results of this year's analysis are presented. This section is by and large identical with the corresponding section in the guideline proposal, and reproduced here mainly for convenience.[2]

Even though substantial progress has been made since a similar model was used to analyse last year's guidelines, it is worth pointing out that this model is still under construction. One important issue is to describe the weaknesses of the present model and the different aspects where the model could and should be improved. The modelling team at the Debt Office would in this context like to invite comments from researchers and practitioners active in the field of debt strategy modelling on the basis of the model presented in this report.

[1]Recent studies in the academic field include Missale(1999), Missale(2000) and Missale(2001). The main results which we arrive at in our specific simulation framework are qualitatively similar to some of the more general conclusions presented in the mentioned studies.

[2]See SNDO "Central Government Debt, Proposed Guidelines", September 2000, available at www.rgk.se.

2 Simulation of stochastic components

In this section, we sketch the framework for the simulation of the stochastic components in the debt analysis model. Three economies are modelled; in addition to the Swedish economy, the EMU and the US economies. A core issue when analysing different debt strategies is the interplay between the funding requirement and the costs for funding in any given time period. The approach taken here is to try to capture the correlation between the financial variables and the borrowing requirement in the framework of a highly stylised and parsimonious macroeconomic simulation model. The ambition has been to keep the modelled variables at a minimum. For each of the three economies the following five random variables are modelled:

- Inflation
- Real GDP
- Short term interest rates
- Long term interest rates
- Exchange rates

In addition, for the Swedish economy we model the

- Central government borrowing requirement

2.1 Stochastic processes

For the majority of variables in the model, the modelling workhorse is a first order autoregressive (AR(1)) process, which, for a generic random variable x can be stated as

$$x_t = \alpha + \rho x_{t-1} + \varepsilon_t \quad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right). \tag{1}$$

If we need to form model consistent expectations based on this process, we can easily form n-step ahead forecasts from the AR(1) process using the formula

$$E\left[x_{t+n}\right]_t = \alpha \sum_{i=0}^{n-1} \rho^i + \rho^n x_t. \tag{2}$$

Of particular interest is the case where t becomes very large, since this case could be interpreted as the long-run equilibrium of the series and used as such when calibrating the model. Letting t tend to infinity we have the result that[3]

$$E\left[x_T\right] = \frac{\alpha}{1-\rho}. \tag{3}$$

A second model building block, which is used throughout, is the Markov chain switching regimes model. In essence, this means that the observable variables, i.e. the variables listed above, follow AR(1)-processes which have different parameters for different states. In the SNDO model, this state is supposed to capture the business cycle, i.e. whether the economy in question is in a boom or in a recession in a certain time period. To fix ideas, let the state variable s follow a first-order Markov process, i.e.

$$P\left\{s_t = j | s_{t-1} = i\right\} = p_{ij}.$$

For the present setting with two states, the transition probabilities are completely determined by a 2X2 matrix

$$\begin{aligned} \mathbf{P} &= \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}, \\ p_{11} + p_{12} &= p_{21} + p_{22} = 1. \end{aligned}$$

All the coefficients, including the transition probabilities, can be estimated using the Hamilton(1989) procedure applied to empirical data. When parameterising a model more informally, we could make use of the relation between transition probabilities and the expected duration of each state[4], i.e.

$$D_1 = \frac{1}{1-p_{11}},\ D_2 = \frac{1}{1-p_{22}}.$$

[3] The derivation of this standard time series result can be found in Appendix A.

[4] For a proof for N states, see Kim & Nelson(1999), pp.71ff.

Simulating a Markov chain and associated "observed" simulated random variables for a given parameterisation is straightforward. If we create one series of error terms following the process $\epsilon \sim NID\left(0, \sigma_\epsilon^2\right)$ and one series of states s, by generating a Bernoulli random number series determined by the transition matrix P, we can use (1) to create a series of an observed variable. It can be shown that the long-run equilibrium of the regime switching process AR(1) process is given by

$$x^* = \frac{\mu_1}{1-\beta_1}\left(1-p_{22}\right) / \left(2-p_{11}-p_{22}\right) + \frac{\mu_2}{1-\beta_2}\left(1-p_{11}\right) / \left(2-p_{11}-p_{22}\right). \tag{4}$$

The intuition behind this expression is that the long-run equilibrium equals the equilibrium, x^*, given by the parameters μ and β in each of the two regimes, multiplied by the steady state probability of being in each of the two states, respectively. A proof of this is sketched in appendix A.

We will now turn to each of the components modelled, briefly sketch the structure of each modelled process and, thereafter, state and briefly comment on parameters chosen.

2.2 Inflation

Inflation is modelled as an AR(1)-process which is not regime dependent:

$$\pi_t = \alpha + \rho\pi_{t-1} + \varepsilon_t \quad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right).$$

It might seem like a natural suggestion to make inflation regime dependent. The modelling choice made here is, however, to assume that the central banks are successful in stabilising inflation and that the impact of the state of the economy on inflation is entirely captured by the reaction function used below to model the short interest rate. The inflation we model will therefore be stable around the inflation target irrespective of the development of the economy.

2.2.1 Parameterisation

	Inflation		
	SWE	EMU	USA
Average:	2.0%	1.5%	2.5%
Parameter	SWE	EMU	USA
α	0.00008	0.00006	0.00010
ρ	0.95	0.95	0.95
σ_ε^2	0.04%	0.03%	0.05%

The parameterisation, which is displayed in the table above, is chosen so that the unconditional expectation is reflecting that each central bank's inflation target is expected to be met in the long run. The choice for Sweden and EMU is consistent with the policy rules declared by the central banks, even if the target is less explicit in the EMU case. For the US, the modelled long-run inflation target is more arbitrary, reflecting an assumed higher tolerance towards inflation than in the two other economies. Note that the parameterisation above (and in the following) is used for simulating monthly changes in the variable in question, and that these changes must be aggregated to a yearly figure in order to arrive at the long-run rate presented in the table.
Data would suggest a fairly persistent inflation process with ρ somewhere around 0.95 for an AR(1) process fitted to historical CPI data, which is the choice made for all three economies. The error term variance corresponds to a yearly standard deviation of the AR(1) process of 0.4%, 0.3% and 0.5% for the three economies respectively. The uncertainty modelled is thus mainly in line with consensus forecasts.

2.3 Real GDP

Real GDP growth is assumed to follow a regime switching AR(1) process, e.g.

$$y_t = \mu_s + \beta y_{t-1} + \epsilon_t \; \epsilon \sim NID\left(0, \sigma_\epsilon^2\right) \quad |\beta| < 1, \; s = 1, 2, \tag{5}$$

where y_t is the *percentage growth rate* of real GDP. The two states, 1 & 2, correspond to a boom state and a recession state in the economy, respectively. This is achieved by different values for the $\mu-$parameters. The different persistence of growth in the two regimes could of course have been modelled by a state-dependent β, but for the sake of simplicity, we have only altered the $\mu-$parameter across the regimes. The shifts between the two states become more gradual the closer the $\beta-$parameter is to unity.

2.3.1 Parameterisation

	GDP Growth		
	SWE	EMU	USA
Average:	2.4%	2.4%	3.1%
Parameter	SWE	EMU	USA
μ_1	0.00015	0.00014	0.00017
μ_2	-0.00009	-0.00005	-0.00008
β	0.95	0.95	0.95
σ_ε^2	0.05%	0.05%	0.05%

The main assumptions made about business cycles and growth are that long-run real growth is expected to be higher in the US than in Sweden and the EMU area. There is also a difference stemming from the assumption on relatively long booms in the US, compared to Sweden and EMU. The only difference between the growth processes in Sweden and EMU is that there are more pronounced business cycles in Sweden, an assumption stemming from the idea that one separate economy, like Sweden, should experience bigger swings to the economic cycle compared with the conglomerate of countries in the EMU. The speed of adjustment between states are assumed to be identical for the three economies.

2.4 Short interest rates

The link between growth, inflation and the short interest rate[5] is modelled as a stylised central bank policy rule, a so-called Taylor rule:

$$\begin{aligned} i_t^T &= r^* + \pi_t^e + \theta\left(\pi_t^e - \pi^*\right) + \lambda\left(Y_t - Y_t^*\right), \qquad (6) \\ Y_t^* &= Y^0\left(1+y^*\right)^t, \\ Y_t &= Y^0 \prod_{i=1}^{t}\left(1+y_i\right), \end{aligned}$$

where r^* is the equilibrium real interest rate, π_t^e is the expected inflation, Y_t^* potential and Y_t actual GDP. The expected inflation is modelled by adaptive expectations, i.e. by assuming that expected inflation is equal to current inflation. For simulated paths of π and y, we use (6) to construct simulated paths for the Taylor rate. In the long run $\pi_t^e = \pi^*$ and $Y_t - Y_t^*$, and hence $i_t^T = r^* + \pi_t^e$. The "observed" short rate is then modelled as an error correction version of the standard AR(1) process

$$i_t = \alpha + \beta i_{t-1} - \gamma\left(i - i^T\right)_{t-1} + \varepsilon_t \quad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right).$$

The result of this specification is that the short rate does not equal the Taylor rate in every period. The speed of adjustment towards the Taylor rate is governed by the γ-parameter. The general idea of this setup is that central banks tend to exercise some kind of interest rate smoothing, and that there is hence a higher degree of persistence in the short rate, than in the Taylor rate.

[5]When translating the short rate into the strategy modelling part, it is assumed to be equivalent to the three month interest rate.

2.4.1 Parameterisation

	Short rates		
	SWE	EMU	USA
Average:	5.0%	4.5%	5.5%
Parameter	SWE	EMU	USA
θ	0.5	0.5	0.5
λ	0.15	0.15	0.15
α	0.003	0.002	0.00275
β	0.95	0.95	0.95
γ	0.1	0.1	0.1
σ_ε^2	0.08%	0.06%	0.06%
i_{LR}^T	3.0%	3.0%	3.0%

The Taylor rate process is assumed to be identical in all three countries, which is also the case for the adjustment speed of the short rate. The long-run differences in the short rate between the economies are assumed to be consistent with the differences in inflation targets. This is a restriction imposed by means of the parameterisation, and is not a corollary of the model's structure. The error term in the Swedish short rate is expected to have a somewhat higher variance than in the other two processes, once again reflecting the smaller size of the Swedish economy.

2.5 Long interest rates

The long term interest rate, the ten-year rate, is modelled by a regime switching AR(1) process:

$$l_t = \eta_{s_{t+i}} + \phi l_{t-1} + \varepsilon_t \;\; \varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right),$$

where l_t is the spread between the three-month and the ten-year rate. The peculiarity of the process is that the intercept, η, switches i periods hence. Each regime is associated with a certain spread. This spread is however modelled as forward looking with perfect foresight, and therefore anticipates the regime shift by i periods. This way, it is possible to model, e.g. yield curves with slopes that increase during the end of a recession.

2.5.1 Parameterisation

	Long rate spread		
	SWE	EMU	USA
Average:	0.74%	0.74%	0.60%
Parameter	SWE	EMU	USA
η_1	0.0002	0.0002	0.00015
η_2	-0.0001	-0.0001	-0.00005
ϕ	0.98	0.98	0.98
i	6	6	6
σ_ε^2	0.17%	0.11%	0.11%

Yield curve spreads are modelled so that recessions are associated with flatter yield curves compared to the booms. The anticipation period mentioned is six months for all three economies. Spreads are also more persistent than, e.g. short rates. Long run spreads are assumed to be somewhat bigger in Sweden and the EMU than in the US. Once again the Swedish spread is assumed to be more volatile than the foreign counterparts.

2.6 Exchange rates

The base for modelling exchange rates is an augmented AR(1)-process for the *real* exchange rate:

$$E_t^{i,j} = \tau + \psi E_{t-1}^{i,j} - \nu \left(y_t^i - y_t^j\right) - \omega \left(\left(i^i + l^i\right)_t - \left(i^j + l^j\right)_t\right) + \varepsilon_t \;\; \varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right),$$

where $E_t^{i,j}$ is the exchange rate between economies i and j. The $\nu-$parameter captures short-run effects from different growth rates. If growth is stronger in one of the economies, its currency will tend to appreciate against the weaker economy. In the long run, the growth difference expression will cause trends in the real exchange rate if $y_{LR}^i \neq y_{LR}^j$. This is in a sense consistent with the literature on real equilibrium exchange rates, where these may trend according to differences in productivity. If one country continuously exhibits stronger growth, a reasonable explanation could be differences in productivity. In the same fashion, the $\omega-$ parameter captures differences in the long interest rates. Finally nominal exchange rates are formed by adding inflation to the changes in the real exchange rates. Long-run differences in inflation will induce the same type of long run trend as growth and interest rate differences caused in the real exchange rates.

2.6.1 Parameterisation

	Exchange rates	
	EUR/SEK	USD/SEK
Long Run:	8.00	9.00
Parameter	EUR/SEK	USD/SEK
τ	0.080	0.090
ψ	0.99	0.99
ν	10	10
ω	0.25	0.25
σ_ε^2	0.07%	0.07%

The process governing the real exchange rate is very persistent, and the degree of mean reversion is therefore quite low. The adjustment parameters, which are set on an ad-hoc basis, are symmetric across exchange rates, as are all other parameters, except for the equilibrium rates which are set to initially equal the rates that prevailed during the period in which the analysis was performed (June 2000). Starting out from these initial values, the trends described above, due to differences in growth rates, interest rates and inflation, will affect exchange rates.

2.7 Central government borrowing requirement

The borrowing requirement of the Swedish central government is modelled according to a policy rule stating that a certain share of nominal GDP, γ, should be used to amortise government debt each year. This assumption reflects the official target of a surplus in public sector financial savings. This condition is to be met measured over the entire business cycle, and not necessarily in every single time period. The average borrowing requirement is equivalent to a borrowing requirement (B) in period t given by

$$B_t = \gamma Y^0 \left(1 + y^* + \pi^*\right)^t ,$$

where $Y^{\mathbf{0}}$ is some initial level of nominal GDP, y^* is given by (4) and $\pi^* = \frac{\alpha}{1-\rho}$. The quantity $Y^{\mathbf{0}}(1 + y^* + \pi^*)^t$ could be viewed as equilibrium nominal GDP in time period t. Multiplying this expression by a constant gives a steadily growing amount which is equal to the amortisation. Such a steady state amortisation rate is, however, not very useful for our purposes. Therefore, we introduce an ad-hoc link to the activity in the economy by modelling (B) as

$$B_t = \gamma Y^{\mathbf{0}}(1 + y^* + \pi^*)^t - \delta(y_t - y^*) + \varepsilon_t,\ \varepsilon \sim NID\left(0, \sigma_\varepsilon^2\right).$$

The parameter δ captures the sensitivity of the government finances to business cycle fluctuations. ε_t could be included to capture random fluctuations in the BR, if this variable is thought to be more variable than what could be explained by business cycle volatility. Note that $E[B_t] = -\gamma Y^{\mathbf{0}}(1 + y^* + \pi^*)^t$ since $E[y_t] = y^*$, so despite the fluctuations, the target amortisation will be met on average.

2.7.1 Parameterisation

	Borrowing req.
Parameter	SWE
γ	0.04
δ	2200
σ_ε^2	0.03%

The parameter γ implies a borrowing requirement of 0.5% of GDP per annum, reflecting a stylised assumption about the division of the targeted surplus between central government and other parts of the public sector. This will imply that the debt increases as a nominal amount, but decreases in real terms and especially as a share of GDP.

3 Simulation of portfolio strategies

This part of the note describes the strategy simulation part of the model, in which portfolios are rolled over and costs and risks computed. First, a couple of general topics about strategy simulation are discussed. Then, the implementation of the strategy simulation in the SNDO model is described in detail. Finally the cost and risk measures used are commented briefly.

3.1 General topics

The strategy simulation implemented in the SNDO model can be described as a continuous rebalancing procedure: in each period maturing debt and the net borrowing requirement, derived from the model described in setction 2, is funded in a way that makes the portfolio meet certain criteria. We thus keep track of all individual cash flows of the simulated portfolio, as will be described in this section.

Working with a strategy simulation of this kind raises two initial questions, which will be addressed below:

1. Which portfolio should we use as the starting point of our strategy simulation?
2. In what terms should the strategies we want to analyse be defined?

3.1.1 The point of departure for the strategies simulation

One could argue that the natural starting point for the strategy simulation would be our actual portfolio. This is the portfolio that we have built up over many years, imposing constraints in the short term. Furthermore, starting with the actual portfolio and its maturity structure would in a sense be consistent with the idea to use our own or other institutions actual forecasts for economic and financial variables in the first periods of our model economy, rather than purely simulated variables.

On the other hand, starting with the actual portfolio tends to make the results from different strategies less clear cut, since the common starting portfolio will influence the results. The shorter the simulation horizon, the larger this influence will be. As our simulation horizon is set to ten years and the current portfolio has a duration of three years, the current portfolio would indeed not be dominating for more than 30 per cent of the simulation horizon. However, since the aim of the simulation studies is to analyse the *long-term* cost and risk characteristics of different debt portfolios rather than short-term effects of moving from one portfolio to another, we still think it is a better idea to work with steady state portfolios. Another advantage with this is that the we can be sure that the differences between alternative strategies reflect different costs and risks associated with different durations and debt allocations, and are not the result of transition costs.

Once a desired debt composition is found, one could involve the actual portfolio, and in a second step of the simulation exercise analyse different ways of restructuring the debt from the initial state to the desired one. This has however not been part of this year's analysis.

3.1.2 Formulation of portfolio strategies

According to the Debt Office's experience, there are two main alternatives for the formulation of strategies in debt simulation models of this kind. The first is to work with refinancing strategies, stating how the borrowing requirement is to be refinanced across the yield curve in each period. In simulation models refinancing strategies are often static, i.e. the borrowing requirement is distributed over different maturities according to a fixed template. They could however also be dynamic. An example of a what we mean by a dynamic refinancing strategy would be to issue in long maturities only when the long-term yields or the slope of the yield curve is below some specified level. In the model we developed last year, we worked solely with static refinancing strategies.

The second way of formulating strategies is to work with strategies defined in terms of a duration target and a target for the allocation between different types of debt (in our case domestic and foreign currency debt). This is the approach that we have been using in this year's modelling work. The reason for choosing this approach, rather than the static refinancing approach, is that we think that this is a natural way of thinking about portfolios, and it is also the way we tend to characterise our portfolio internally and externally (for instance when communicating with the Swedish Ministry of Finance).

Using this "portfolio approach" for the strategies implies that the borrowing requirement in each period will be refinanced in a way that makes the portfolio meet the duration and allocation targets to the largest possible extent. This could be compared to the static refinancing approach, in which borrowing requirement are always refinanced in the same way, and duration and debt allocation fluctuates

Even though an approach in which you formulate strategies in terms of duration and allocation is more intuitive than working with refinancing strategies, it has still got some less appealing effects, primarily concerning the foreign currency debt:

Assume that we in one strategy have set the target for the FX debt to 20 per cent. In a period when the domestic currency depreciates, the proportion of the foreign debt increases and we would have to buy back debt to meet the target. This buyback would occur at a time when the domestic currency is weak. Similarly, an appreciation in the domestic currency would result in a need to increase foreign exposure, at a time when the krona is strong. If we have mean reversion tendencies in exchange rates, this effect will systematically make foreign currency debt less favourable.

One way of avoiding this effect would be to define the target allocation given the foreign exchange rates prevailing at the start of the simulation. For example, if the foreign currency part of the debt is 10 per cent at the start and the krona then depreciates 20 per cent, the target would shift to 12 per cent. This method would be equivalent to assuming that the amount of foreign currency units (dollars, euros etc.) is constant over time. With this approach, however, the level of FX debt relative to krona debt would increase or decrease if the currency is trending.

The way the FX target is handled in the model is that the total refinancing need is distributed over the three currencies according to the allocation specified by the strategy. If the strategy is to have 30 per cent foreign currency, then 30 per cent of the gross borrowing requirement will always be met in dollars and euros. Thus, the allocation over currencies is very similar to the kind of static refinancing strategies discussed above. Since a certain part of the portfolio rolls over each month, the level of FX debt will be continuously adjusted towards the strategy target. In this way, the systematic excess cost of foreign exchange is avoided, at the same time as the level in relation to the debt as a whole is kept fairly constant.

3.2 The strategy simulation

The strategy simulation is done on top of the simulated economy. Each strategy is run a thousand times, and each individual run stretches ten years into the future, with monthly steps. Economic and financial variables are simulated according to the description in section 2 for one simulation run at the time.

The goal of the analysis is to examine the effects of different duration choices and shares of foreign currency debt in rough terms. The strategies investigated are therefore clearly differentiated and extend over a relatively wide range of alternatives. We have also found it sufficient to have a small number of strategies. The share of foreign currency debt in total government debt has been allowed to vary between 0 and 45 per cent, in stepd of 15 percentage points. The shares of EUR and USD debt have been set at 70 and 30 per cent, respectively, approximately equivalent to the current structure of the foreign currency debt. The duration figures are two, three and four years, respectively. In all strategies, the duration target is the same for all three categories of debt. This leads to the following twelve strategies:

FX debt (share)	FX debt (dur)	SEK debt (dur)
0 %	2	2
15 %	2	2
30 %	2	2
45 %	2	2
0 %	3	3
15 %	3	3
30 %	3	3
45 %	3	3
0 %	4	4
15 %	4	4
30 %	4	4
45 %	4	4

3.2.1 The initial portfolio

The initial portfolio is held in a two dimensional matrix. This matrix has two columns for each currency, one containing the debt principals and one containing the coupon payments. Since we are modelling a three economy world, we thus have six columns in our initial portfolio. The 120 rows of the matrix correspond to monthly maturity buckets, ranging from one month to ten years. Thus the amounts in the first row of the initial portfolio matures during the first month (step) of the simulation run, the amounts in the second mature in the second step and so on.

At the start of each simulation run, all outstanding debt is distributed evenly over the maturity buckets, in a way that makes the portfolio duration and currency allocation the same as the target for the strategy that is being simulated. The sum of the debt outstanding at the start of the simulation is set to 1300 billion SEK, approximately the same as the actual current size of the debt. The initial coupon of all debt has been set to six per cent. It is thus implicitly assumed that all debt that is outstanding at the start of the simulation run has been raised at the same interest rate. This simplifying assumption leads to some loss of generality, since the initial portfolio affects the long duration portfolios during a larger part of the simulation

3.2.2 The strategy simulation step by step

In this section, each step in the portfolio simulation is described. In the SNDO simulation model, these steps are implemented through a number of routines. In each monthly step of the simulation run, debt matures and coupons are paid. The task for the strategy simulation is to refinance maturing debt in a way that keeps the portfolio on the strategy target, and to keep track of costs, risks, debt to GDP ratio and other variables that might be of interest.

The first number calculated in each simulation step is the sum of coupon payments in that period. This is calculated simply as the coupon payment in each currency, times the simulated FX rates for that particular period.

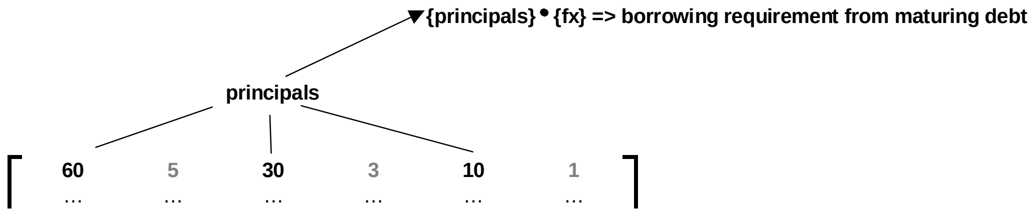

Then, the borrowing requirement for the specific period is calculated in the same way: the amounts maturing in each currency in this period are found in columns 1, 3 and 5 of the first row of the portfolio matrix. These amounts are converted to the domestic currency at the simulated FX rates.

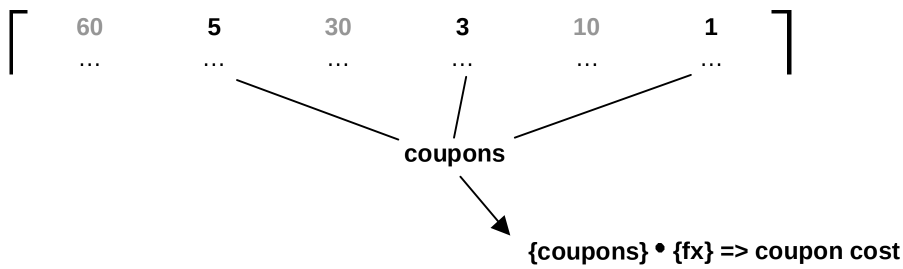

This gives the borrowing requirement stemming from the debt itself. There is also a net borrowing requirement from the budget in the simulated economy. This budget deficit/surplus is added to the maturing debt, giving the total refinancing need. Finally, the total refinancing need is distributed over the three currencies according to the proportions in the strategy, as described above.

The next step in the strategy simulation is to let the shortest bonds mature. This is done by simply moving the contents of all rows in the portfolio matrix up one row, and setting the amounts in the last row to zero. The portfolio matrix will now contain the portfolio composition after the maturities of the current period have occurred, but before any new bonds have been issued.

The borrowing requirement in each currency should now be refinanced in a way that makes the portfolio hit the target, if possible. To be able to determine at what duration point we should borrow, we need to know the duration and market value of the debt outstanding in each of the three currencies. To calculate this we need three sets of discount factors, one for each currency. For each currency, we needdiscount factors for each of the maturity buckets. In the economic/financial simulation, we simulate only two points on the yield curve, the three-months rate and ten-year rate. To get the discount factors, we therefore first linearly interpolate yields for each maturity bucket, and then convert them to discount factors in the standard way:

$$df = \frac{1}{(1+r)^t}.$$

As we have the cashflows of the portfolios organised in maturity buckets, and know the time to maturity of each bucket, it is straightforward to calculate the market value, mv, of each part (currency) of the debt in the portfolio as

$$mv = \sum cf_t * df_t,$$

where t runs from 1 to 120 and denotes the time to maturity of each bucket; cf denotes the cashflow in bucket t and df is the discount factor associated with maturity t. With the same notation, the duration of each sub portfolio (Swedish krona, euro and dollars) is calculated as[6]

$$duration = \frac{\sum cf_t * df_t * t}{mv}, \tag{7}$$

Note that the durations of the three sub-portfolios of the debt are not weighted, since in the model duration targets are specified per currency, not for the portfolio as a whole[7]. Denote the duration of the debt before new funding d_0, the market value of the debt before new funding mv_0 and the borrowing requirement br. In order to meet the duration target d^*, the borrowing requirement has to be funded at a certain average duration d_a, which can be calculated as:

$$\begin{aligned} d^* &= \frac{d_0 mv_0 + d_a br}{mv_0 + br} \Rightarrow \\ d_a &= \frac{d^* (mv_0 + br) - d_0 mv_0}{br} \end{aligned}$$

All of the borrowing requirement is assumed to be covered by issues of par bonds, i.e. bonds with an integer number of years to maturity and a coupon making the present value equal to the nominal amount. In each currency, par bonds maturing between one year and ten years from the current date can be used. Bonds are throughout assumed to pay annual coupons.

Given the discount factors calculated from the simulated interest rates, we calculate the coupon c of a par bond maturing in N years using the following

[6] Technically, this equation gives the Fisher-Weill duration, since we use the discount factors corresponding to zero coupon rates to discount the cashflows. The more frequently used Macauley duration uses discount factors produced by the yield-to-maturity of the specific bond. The difference between these two measures is small: with an upward sloping yield curve, Fisher-Weill method will give a slightly lower duration than the Macauley method.

[7] In this year's simulations however, all sub-parts of the debt have been given the same duration target.

equation:

$$\begin{aligned} 100 &= c\sum_{i=1}^{N} df_i + 100 df_N \Rightarrow \\ c &= \frac{100\,(1-df_N)}{\sum_{i=1}^{N} df_i}. \end{aligned}$$

The duration of the par bonds are calculated using (7) . For each currency we now have the duration and coupons of par bonds maturing between one and ten years from now. We also know at which duration point we need to borrow in each currency to meet the duration target. Since that duration point is unlikely to coincide with the duration of a single par bond, new funding is done in two par bonds in each currency. Which two bonds that will be used depends, in principle, on the shape of the yield curve: the model will allocate the borrowing in a way that minimises the average coupon of the funding. If the yield curve is concave (steeper in the short end than in the long end), the shortest and the longest par bond will be used. If the curve would be convex (steeper in the long end than in the shorter end), as is often the case when the yield curve is inverted, the two bonds nearest below and above the duration at which we want to fund are chosen. This principle is illustrated below:

The amounts to be borrowed in each of the two bonds are calculated in the same way as standard barbell calculations. Let d denote duration and N nominal amounts (equalling clean and dirty prices as we are working with par bonds!). d^* is the duration needed to meet the target and subindices S and L denote the short and long par bond, respectively. The sum of nominal amounts in the two bonds must of course equal the borrowing requirement br. The amounts that should be allocated to the short and the long bond are then given by:

$$\begin{aligned} br &= N_S + N_L \Rightarrow N_L = br - N_S \ , \ \ d^* = \frac{d_s N_s + d_L\,(br - N_S)}{br} \Rightarrow \\ d^* br &= N_S\,(d_S - d_L) + d_L br \Rightarrow N_S = \frac{br\,(d^* - d_L)}{(d_S - d_L)} \end{aligned}$$

In some cases, the duration needed to meet the duration target is outside the range of par bond durations. All funding is then done in the shortest or the longest par bond, as appropriate. Naturally, the duration target cannot be exactly met in those cases.

At the end of this step, we have two, or in some cases only one, new bond in each currency. The maturities and coupon rates are known. All new bonds are added to the portfolio, and thereby changes the amounts in the maturity buckets in which the new bonds pay coupons or mature.

In most cases, the gross borrowing requirement is larger than zero - we need funding to refinance maturing bonds and/or a budget deficit. However, it can also happen that the gross borrowing requirement is less than zero. This happens when the simulated economy shows a budget surplus larger than the sum of maturing bonds in the portfolio. In those situations, the roll-over procedure becomes more complicated. Instead of issuing new bonds, we now have to buy back bonds. Buybacks must be done in a way that makes the portfolio duration meet the strategy target and can of course only be done utilising the bonds that are currently in the portfolio. Furthermore, the individual cash flows of a bond cannot be bought back separately, but must be handled as a whole.

In the model, however, cash flows from all bonds are stored in maturity buckets, i.e. in any maturity bucket we could have coupons and principal amounts from a number of different bonds. This is illustrated below in a simplified example with three bonds and yearly maturity buckets:

Bond matrix

Bond	Tenor	Nominal	Coupon
1	2 yrs	100	8%
2	3 yrs	50	5%
3	5 yrs	100	7%

Cash flow matrix

Maturity bucket	Principals	Coupons
1	0	17.5
2	100	17.5
3	50	9.5
4	0	7
5	100	7

For example, in the first maturity bucket we have no principal, since the shortest bond matures in two years. However, since all bonds pay annual coupons, we will have 17.5 (= 8%*100 + 5%*50 + 7%*100) in the coupon column .

The first step in the buyback situation is therefore to revert the cash flow matrix into a matrix of "average bonds" per maturity bucket. The procedure starts at the bottom line of the cashflow matrix and then works its way up to the first row. Using the last line of the portfolio above we have 100 in the principals column and 7 in the coupons column. Since this is the last row, we know that all coupons paid in this bucket must belong to the bond maturing in that bucket. The coupon rate is therefore 7 per cent. The second last row has a zero in the principals column, and is therefore ignored. In the third maturity bucket, the coupons paid sum to 9.5. However, 7 of them come from the five year bond. Consequently, the bond maturing in year three pays a 9.5 - 7 = 2.5 coupon, giving a coupon rate of 2.5/50 = 5 per cent. Finally, the bond maturing in year two pays a 17.5 - 9.5 = 8 coupon, corresponding to a coupon rate of 8%. In this way, the bond matrix underlying any cash flow matrix can be calculated recursively. Once the coupons of each bond are known, it is a simple matter to calculate the duration of the same bonds.

The durations and coupons of all the bonds in the portfolio are now known. Unfortunately there is no simple way to determine which bonds to use for the buyback: as long as the required duration is between the shortest and the longest duration in the current portfolio, one can find a large number of combinations of bonds that in the right proportions will make the portfolio meet the target duration. At worst, we would have 3600 possible combinations. Going through all of them and finding out which one would cost the least in terms of paid out premiums would be computationally inefficient. In addition to this, many of the combinations are likely to imply that more should be bought back in an individual bond than is currently outstanding. We have therefore implemented a simpler and more robust algorithm, which works also in the cases where the required duration falls outside the durations in the current portfolio.

For each bond, the absolute difference between the bond's duration and the required duration is calculated. The bonds are then sorted by these differences, in ascending order. The top row thus contains the bond whose duration is closest, in absolute terms, to the required one. If the amount outstanding in that bond is larger than the amortisation needed, then only this bond will be affected by the buyback. If the amount outstanding is instead smaller than the amortisation needed, the bond on the second row will be used as well, and so on until enough bonds have been bought back.

In extreme cases, all of the outstanding debt may be bought back. When no bonds are left to buy back, any remaining surpluses are deposited, one month at the time.

All bonds that are completely or partly bought back are stored in a matrix, containing maturity, coupon and amount bought back. These nominal amounts are then subtracted from the cash flow matrix.

Since coupon bonds seldom trade at par, bond buybacks are done at a premium or a discount. Thus, mark-to-market gains or losses are realised when bonds are bought back. Although the focus of the SNDO simulation exercise is not on mark-to-market costs, it would be inappropriate not to take these realised mark-to-market effects into account. Therefore, we add the difference between the dirty price at which the bond is bought back and its nominal amount to the coupon cost for the period. If, for example, the coupon payments in one period amount to 10, and bonds for a nominal amount of 40 are bought back at 42, the total cost for the period would be 12 (= [42 - 40] + 10). Note that the mark-to-market cost is not periodised in any way. This cost is hence incurred in the period when it occurs.

4 Results

4.1 Cost measures

In the analysis, the Debt Office used two cost measures, a nominal one where costs are measured in SEK, and a more real-term one where costs are expressed as a share of GDP, called the debt cost ratio. For both measures, risk is expressed as the percentile distance between the 50th and 95th percentile in the simulated cost allocation. A 95 per cent percentile range of, say, 20 per cent, can be interpreted as meaning there is a 5 per cent probability that the portfolio in question will turn out to have a cost more than 20 per cent above the median. If the percentile distance is 40 per cent, there is a 5 per cent probability that the outcome will deviate by more than 40 per cent from the average. The larger the percentile distance, the higher the risk of the portfolio. Expressing risk in this way rather than in terms of standard deviation in allocation makes it possible to focus on the side of the risk that is relevant, namely that government debt costs will be significantly higher than expected.

In both the nominal and the real cost measure, costs are defined as cash flow costs. A cash flow cost measure just takes costs that are actually paid into account, whereas the main alternative, a mark-to-market cost measure, looks at the changes in the present value of the portfolio from one period to another. Cash flow cost measures by definition capture only realised costs, whereas mark-to-market cost measures captures unrealised costs aswell.

There are arguments in favor of both measures. As stated above in this paper, as well as in the guidelines proposal, the main focus in the analysis performed by the Debt Office is on cash flow costs. The reason for this is that our aim is to minimise the costs that are actually affecting the budget surplus. Since the debt is long term, swings in mark-to-market value are in most cases not materialised. Therefore, mark-to-market costs are viewed as less relevant for the decision on the long term composition of the debt.

4.1.1 Nominal costs

With this cost measure, debt costs are estimated in SEK terms, period by period. In other words, all coupon payments are translated using the simulated exchange rate for each period, plus any realised exchange rate gains or losses on loans that have been repurchased.

The chart below shows the results from the simulations of the twelve strategies, using the assumptions in the basic parameterisation. Roughly speaking, costs are primarily affected by the choice of duration, while the choice of the share of foreign currency debt mainly affects risk. A portfolio with borrowing denominated only in SEK and with a shorter duration has a lower expected cost but a somewhat higher risk, which is consistent with the results that the Debt Office presented last year. This result is mainly dependent on the assumption that on average, yield curves have an upward slope, plus the fact that a more short-term portfolio is refinanced more frequently and is therefore more affected by interest rate volatility.

The effects of exchange rates on interest payments are somewhat less than the effects from the choice of duration. Based on the given parameterisation, there is no cost advantage in foreign currency loans either. The somewhat lower EMU interest rates are eaten up by the depreciation of the krona against the euro that results from the lower inflation rate in EMU. Even if there is a corresponding effect from an appreciating dollar, the euro effect dominates, since 70 per cent of the foreign currency debt is EUR-denominated. From a risk standpoint, however, there is a minimum at 15 per cent foreign currency debt for portfolios with a duration of 2 or 3 years.

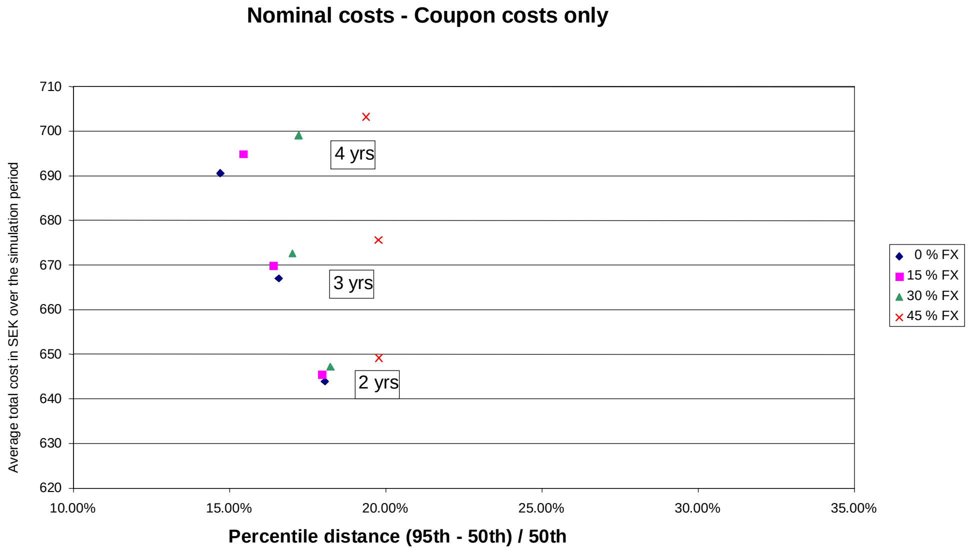

It is worth noting that the risk picture that emerges when only taking into account coupon payments is based on strong assumptions. Ignoring the exchange rate effects on the face value of bonds implies that the government is issuing perpetual bonds in foreign currencies which may then remain outstanding forever. Since the stock of foreign currency loans will then, in principle, be unchanged, this also implies that the current debt level is optimal, not as a share of total debt, but in terms of nominal foreign currency amounts.

Since it is difficult to believe that the current stock of foreign currency debt could be optimal in any sense, and that it moreover could be expected to remain optimal over time, this measure underestimates actual risk. It may therefore be essential to include changes in market values when assessing the risks of foreign currency loans. In qualitative terms, taking into account the effects of exchange rates on face values implies a substantial increase in the risk of foreign currency loans, since these are far larger than the coupon amounts in foreign currency. In simulations where market values have a full impact, foreign currency loans are so risky that the optimal share of such loans would be 0 per cent. Letting market values have a full impact does not provide a realistic picture of the risk either, however, since it will probably never become necessary to repurchase foreign debt within a short time interval. In other words, this approach overestimates risk.

By way of summary, nominal cost measures indicate that exposure to foreign currencies results in sizeable risks and that diversification gains can only justify a limited foreign currency exposure. Even with a partial focus on coupon costs from a risk standpoint, it is difficult to justify a larger share of foreign currency debt than 15 per cent. The bigger an emphasis one then places on the effects of exchange rates on face value, the smaller the optimal proportion of foreign currency debt will be.

4.1.2 Costs expressed as a share of GDP

The debt cost ratio, where nominal coupon costs are stated as a share of GDP, can be justified by the fact that budget balance can be assumed to be correlated to growth and that lower debt costs as a proportion of GDP thus imply a smaller need for adjustments in the government budget in order to meet interest payments. This cost measure, in all its simplicity, is a step towards a more ALM-based approach, in keeping with the discussion in section 2.2.2. of the guideline proposal. From this standpoint, the debt cost ratio seems like a more interesting measure. As above, the risk measure is a percentile distance in the allocation of debt cost ratios across the entire simulation horizon. The results are presented in the figure below.

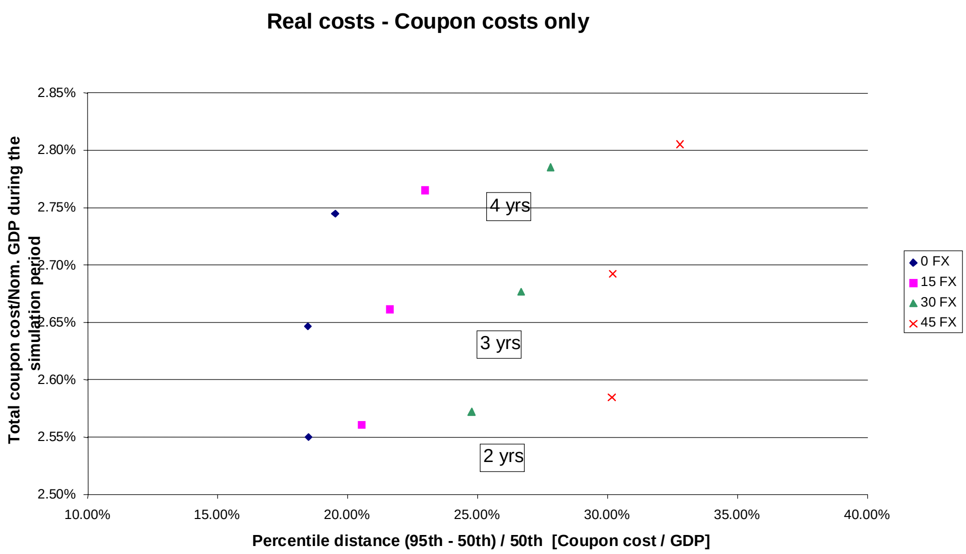

It is worth noting, by way of introduction, that the trade-off between costs and risks that existed in the nominal results does not occur. A shortening of duration not only leads to lower costs, but also to lower risk. This perhaps counterintuitive result is explained by the interest rate process in the model. High growth leads to a closing of the output gap and to an increase in short-term interest rates, assuming that the Riksbank (Swedish central bank) follows the Taylor rule. Since the yield curve on average has a constant slope for a given regime, parallel shifts in the yield curve mainly occur. Given a shorter-term portfolio, a relatively larger share of the debt will be refinanced during each period, which in turn leads to a higher correlation between coupon cost and the general interest rate situation during the simulation period. This, plus the fact that interest rates and GDP are highly correlated via the Taylor rule, results in a higher correlation between coupon cost and GDP for shorter-term portfolios. The consequence of this is a less volatile debt cost ratio with short-term debt. Lower costs for a shorter-duration portfolio follow, as earlier, from the assumption of upward-sloping average yield curves.

Enlarging the share of foreign currency debt means both higher risk and higher cost. Higher cost is mainly a product of the model's parameterisation, while higher risk is a product of the model's structure. Economic growth and thus interest rates in EMU and the US are independent of each other and independent of interest rates in Sweden. A large share of foreign currency debt thus contributes unequivocally to greater variability in the debt cost ratio, since the correlation between coupon cost and GDP decreases the larger the share of foreign currency debt is. If one compares the impact on risk level of longer duration with the impact of a larger share of foreign currency debt, the effect of foreign currency debt is substantially larger, based on the assumptions in the basic parameterisation. Saying that the economic cycle in Sweden is independent of the EMU and US economies is a strong assumption. The impact of foreign currency borrowing on the volatility of the real cost measure is a product of one's assumptions about the correlations between foreign interest rates and the Swedish economic cycle. Qualitatively, however, the finding that higher foreign currency debt leads to greater risk in this respect should still be robust, since no perfect correlation exists.

A more real-term approach thus implies that short duration and a small share of foreign currency debt leads to both lower costs and lower risk, given the model's assumptions. It is, however, appropriate to emphasise that this strong result is related to a rather strong implicit assumption in the model. In the model world, both fiscal policy and monetary policy constantly enjoy the full confidence of participants in the securities and foreign exchange markets. All explosive scenarios are excluded. All financial variables return, sooner or later, to their expected values. Above all, the result showing low real risk in short-term borrowing is sensitive to this assumption. Given a short-term borrowing strategy, a crisis of confidence in fiscal and monetary policy could lead to a need to refinance a large proportion of the portfolio at a time of unfavourable interest rates.

As for the risk inherent in foreign currency borrowing, it may be reasonable to imagine that a crisis of confidence could adversely affect the SEK exchange rate, which - all else being equal - would increase the risk of foreign currency debt. Easing the assumption of constant full confidence would thus weaken the result of low risk in short-term borrowing, while strengthening the result of higher risk due to a larger share of foreign currency debt.

4.1.3 Sensitivity analysis

As a simple sensitivity analysis, an alternative parameterisation is presented here, in which Swedish economic growth is substantially weaker than in the basic parameterisation. This is achieved by assuming that the boom and recession regimes are equally long. Given this parameterisation, average growth is only 0.7 per cent annually. The other parameters and mechanisms are identical to those of the basic model. The figure below presents both nominal and real-term cost measures.

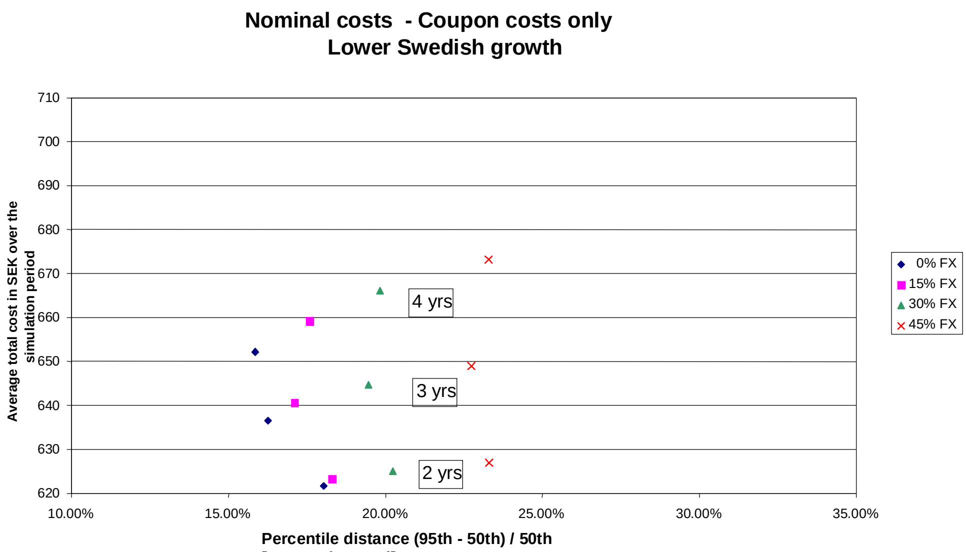

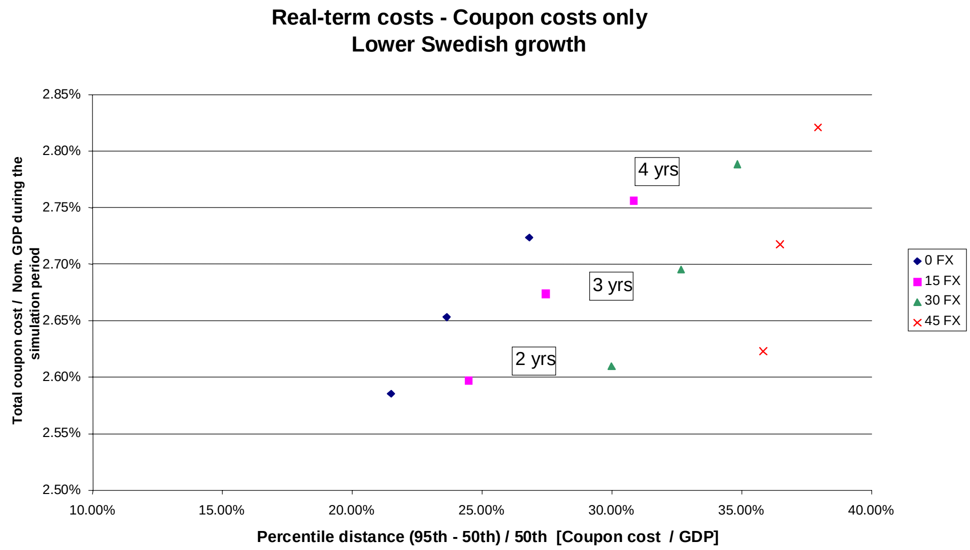

Most of the qualitative results from the basic model remain. For both cost measures, a larger share of foreign currency debt implies higher cost. This is due to the relatively weaker economic growth in Sweden and the resulting general depreciation of the krona. The result that a short duration leads to a higher correlation with GDP is strengthened here. This is related to the fact that the number of changes of regime is smaller and the correlation between interest rates and GDP generally larger. Again, the assumption of unwavering faith in economic policy is decisive. This assumption may be viewed as even stronger in the gloomy growth scenario above.

4.2 Summary and conclusions

The analysis using the described model in preparation for this year's proposed guidelines can be summarised briefly as follows:

- Given the assumption that yield curves slope upward on average, cost savings in nominal terms can be obtained by means of a relatively short duration, without allowing risk, expressed as variation in interest payments, to become unacceptably large. Given a real-term cost measure, the debt cost ratio, this result is even stronger, since the interest costs in the model co-vary with central government revenues.

- Given a nominal cost measure, the analyses indicate that foreign currency debt leads to greater variation in costs and only limited diversification gains. Nor does foreign currency borrowing yield any expected cost advantage. The Debt Office model results can justify a foreign currency share of about 15 per cent based on a nominal cost measure. Using a debt cost ratio as the measure, foreign currency debt appears to be a more risky alternative. Since the cost differences are small, this measure indicates that foreign currency debt should be brought down to zero.

- These results are based on the assumption that both monetary and fiscal policy enjoy full confidence. Easing this assumption would probably imply a longer duration, but a smaller proportion of foreign currency debt would also be appropriate, especially on the basis of a debt cost ratio.

It should be emphasised that the Debt Office model is an analytical tool that is still under development. Other assumptions can and should be studied in order to improve the understanding of the model's characteristics and the sensitivity of the simulation results. The structure of the model may also need to be reappraised and tested more thoroughly. Despite these qualifications, the Debt Office believes that in its current condition, the model illustrates some essential characteristics and mechanisms of Swedish government debt and the underlying economy.

References

[1] HAMILTON, J.D. (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica*, 57(2), 357-384.

[2] HAMILTON, J.D. (1994), *Time Series Analysis*, Princeton, New Jersey: Princeton University Press

[3] KIM, C-J and NELSON, C.R. (1999), *State Space Models with Regime Switching*, Cambridge, Massachusetts: MIT Press

[4] MISSALE, A. (1999), *Public Debt Management*, Oxford: Oxford University Press

[5] MISSALE, A. (2000), "How Should the Debt be Managed? Supporting the Stability Pact.", Forthcoming in Brunila, A., Buti, M. and Franco, F.: *The Stability and Growth Pact - The Architecture of Fiscal Policy in EMU.*

[6] MISSALE, A. (2001), "Opitimal Debt Management with a Stability and Growth Pact.", Fortcoming in *Public Finance and Management.*

A Appendix

In this appendix, a few of the results presented in the main text are derived in a slightly more rigorous way.

A.1 Derivation of equation (3)

The expected mean conditional on the information in period t can be derived as follows

$$
\begin{aligned}
x_t &= \alpha + \rho x_{t-1} + \varepsilon_t \quad \rho < |1|, \varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right), \\
x_1 &= \alpha + \rho x_0 + \varepsilon_1, \\
x_2 &= \alpha + \rho x_1 + \varepsilon_2 \\
&= \alpha + \rho\left(\alpha + \rho x_0 + \varepsilon_1\right) + \varepsilon_2 \\
&= \left(1+\rho\right)\alpha + \rho^2 x_0 + \rho\varepsilon_1 + \varepsilon_2, \\
x_3 &= \alpha + \rho x_2 + \varepsilon_3 \\
&= \left(1+\rho+\rho^2\right)\alpha + \rho^3 x_0 + \rho^2 x_1 + \rho\varepsilon_2 + \varepsilon_3 \\
&= \alpha\sum_{t=0}^{2}\rho^t + \rho^3 x_0 + \sum_{t=1}^{3}\rho^{3-t}\varepsilon_t, \\
x_T &= \alpha\sum_{t=0}^{T-1}\rho^t + \rho^T x_0 + \sum_{t=1}^{T}\rho^{T-t}\varepsilon_t, \\
E\left[x_T\right] &= \alpha\sum_{t=0}^{T-1}\rho^t + \rho^T x_0, \\
T &\longrightarrow \infty \Longrightarrow E\left[x_T\right] = \frac{\alpha}{1-\rho},
\end{aligned}
$$

where we have used the standard result for a geometric series.

A.2 Derivation of equation (4)

If we wish to make model consistent forecasts conditional on the value of an observed variable x and the state variable s in any time period, we need both forecasts of the future state and for the expected growth rate in each state.

The probabilities of future states m steps ahead, is given by

$$\boldsymbol{\xi}_{t+m}|\boldsymbol{\xi}_t = \mathbf{P}^m \boldsymbol{\xi}_t$$

where $\boldsymbol{\xi}$ is a vector giving the probability of being in a specific state at a given time t. At an observed point in time, $\boldsymbol{\xi}$ assumes the values $[1,0]'$ or $[0,1]'$ depending on the observed state. As before we can use the properties of the AR(1)-process to find the m-step ahead prediction for each state, i.e.

$$E\left[x_{t+m}|\mathbf{x}_t, s\right] = \mu_s \sum_{i=0}^{m-1} \beta_s^i + \beta_s^m x_t,\ \ s = 1, 2, \tag{8}$$

where $\mathbf{x}_t$ denotes the entire realisation of x up until period t. If we, for each time-ahead period, stack these forecasts into a 2x1 vector $\mathbf{h}_{t,m}$, we can obtain the following expression for the forecasts of the [8]:

$$E\left[x_{t+m}|\mathbf{x}_t\right] = \mathbf{h}'_{t,m}\left(\boldsymbol{\xi}_{t+m}|\boldsymbol{\xi}_t\right)$$

The main result we need in this context however concerns the long run equilibrium of the modelled variables when there are several states to be modelled. From (8) it is evident that $\beta_s^m x_t$ tends to zero as m becomes large, and that $\mu_s \sum_{i=0}^{m-1} \beta_s^i$ simultaneously tends to $\frac{\mu_s}{1-\beta s}$. Thus we have a long-run $\mathbf{h}$ defined by

$$\mathbf{h}'_{LR} = \left[\frac{\mu_1}{1-\beta_1}, \frac{\mu_2}{1-\beta_2}\right]$$

In order to find the unconditional probabilities of the states we would need probabilities that do not change over time, i.e. $\boldsymbol{\xi} = \mathbf{P}\boldsymbol{\xi}$, under the restriction that the probabilities sum to unity, i.e. that $\boldsymbol{\iota}'\boldsymbol{\xi} = 1$. Note that the first condition implies that $(\mathbf{I} - \mathbf{P})\,\boldsymbol{\xi} = \mathbf{0}$ and that we therefore may write the

[8] This is also intuitively reasonable, but is demonstrated more rigorously in Hamilton(1994), pp. 694f.

problem as the following system

$$
\begin{aligned}
\begin{bmatrix} (\mathbf{I}-\mathbf{P}) \\ \boldsymbol{\iota}' \end{bmatrix} \boldsymbol{\xi} &= \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix} \\
\text{or by introducing } \mathbf{A} &= \begin{bmatrix} (\mathbf{I}-\mathbf{P}) \\ \boldsymbol{\iota}' \end{bmatrix} \Rightarrow \\
\mathbf{A}\boldsymbol{\xi} &= \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix} \Rightarrow \\
\boldsymbol{\xi} &= \left(\mathbf{A}'\mathbf{A}\right)^{-1} \mathbf{A}' \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}.
\end{aligned}
$$

Carrying out this matrix multiplication for the two-state case is somewhat tedious but straightforward and gives the following solution

$$
\boldsymbol{\xi} = \begin{bmatrix} (1-p_{22}) / (2-p_{11}-p_{22}) \\ (1-p_{11}) / (2-p_{11}-p_{22}) \end{bmatrix}.
$$

The long run equilibrium , x^* can therefore be computed as

$$
\begin{aligned}
x^* &= \mathbf{h}_{LR}' \boldsymbol{\xi} \\
&= \frac{\mu_1}{1-\beta_1} (1-p_{22}) / (2-p_{11}-p_{22}) + \frac{\mu_2}{1-\beta_2} (1-p_{11}) / (2-p_{11}-p_{22}).
\end{aligned}
$$