The SNDO's Simulation Model for Government Debt Analysis*

Pål Bergström, Anders Holmlund and Sara Lindberg†

May 2002

Abstract

The Swedish National Debt Office has developed a stochastic simulation model to analyze the long-term costs and risks of various conceivable central government debt portfolios. The model has been used in order to provide a quantitative basis for decisions on the Debt Office's annual proposed guidelines. This report presents a detailed description of the present version of the model and discusses some issues of implementation.

*This report, as well as the present version of the simulation model, draws heavily on and is partly identical to earlier work by the Debt Office, presented in Bergström & Holmlund (2000). The report has benefitted from comments and suggestions from several colleagues at the Debt Office, most notably Lars Hörngren. All errors are our own.

†Corresponding authors: Anders Holmlund and Sara Lindberg. The Swedish National Debt Office, SE-103 74 Stockholm, Sweden. E-mail: anders.holmlund@rgk.se and sara.lindberg@rgk.se

1 Introduction

The costs of central government debt are influenced by many factors. They are determined primarily by the size of the debt, but interest rates, exchange rates and inflation rates are also significant factors. Since borrowing occurs continuously and — on average — for long maturities, these factors will influence costs for a number of years to come. The costs during a given period are thus a complex function of the borrowing requirements, interest rates, exchange rates and issuance patterns during earlier periods. In order to assess government financial risks from a more systematic asset and liability management (ALM) perspective, it is also important to consider how the costs of central government debt co-vary with government revenues. A first step towards such an analysis is to relate all costs to GDP, which is thus being used as a measure of government net revenues.

In order to study the effects of portfolio choice on long-term costs and risks within a coherent framework, the Swedish National Debt Office has developed a stochastic simulation model. This model has been used in order to provide a quantitative basis for decisions on the Debt Office's annual proposed guidelines.[1] The object is to minimize the long-term costs of debt with due regard to risk. The development of the model has been made in two steps. The first version of the model was developed in 2000. A technical description of this model is given in Bergström & Holmlund (2000). The present version of the model was developed in 2001 and is the one presented in this report. This model is an extended version of the 2000 model, where the main amendment is the introduction of inflation-linked bonds. This has led to some major revisions of the stochastic processes generating the interest rates in the model.

A central issue when the simulation model was developed was how the variables which determine the costs of government debt co-vary with each other and how this interplay best should be modeled. One way to do this is

[1] See Swedish National Debt Office (2001) and Swedish National Debt Office (2002), available at www.rgk.se.

to introduce co-variation between the variables via the co-variance matrix for the disturbance terms (random components) of the variables. Another way is to model the variables within the framework of a macroeconomic model, taking directly into account the interplay among such variables as interest rates, growth and inflation. For the purpose of our model, we chose to adopt the latter approach, since we believe that this approach provides a clearer picture of the structural relations among the variables and thereby increases our understanding of the results.

In developing the simulation model, our ambition has been to keep the model simple and to minimize the number of variables to be modelled. Consequently, the model gives a highly stylized description of the economy and of central government debt management. In short, the model consists of two parts. The first part is a *macroeconomic model* which is used to simulate the development of a number of important macroeconomic and financial variables, such as GDP, inflation, and interest rates. The second part is a *strategy simulation part*, which is used to control how, in a given portfolio strategy, the government finances its period-to-period borrowing requirement and refinances maturing loans. This part also estimates what costs and risks are associated with different portfolio strategies, given the simulated economic course of events. In practice, we use the macroeconomic model to simulate 1,000 different economic paths over a 30 year period, in quarterly steps. These paths are then used as inputs in the strategy simulation model, where we confront various portfolio strategies — defined in terms of their debt composition and target duration — with the simulated economic paths of events and calculate the costs and risks associated with each strategy.

In the remainder of this report we present a detailed description of the present simulation model and discusses some issues of implementation. In this way, we hope to encourage comments and opinions on the simulation model from both researchers and other professionals in the field of government debt management. This is especially important since studies in this fields are still fairly scarce.[2] It is also worth pointing out that our model is still under construction. Therefore, one important object of this report is to describe the weaknesses of the model and the different aspects where the model could and should be improved.

The report is organized as follows. In the next section, we present the macroeconomic part of the simulation model, while the strategy part is presented in Section 3. In Section 4 we describe the methods we use to estimate the costs and risks of different portfolio strategies. In Section 5 we discuss some issues of implementation. Finally, Section 6 concludes with some thoughts about future directions of our analytical and modelling work. Note that the results of the simulation exercises are not presented here, since this report is primarily intended to document the construction and implementation of the model. Readers interested in the results are referred to the Debt Office's proposed guidelines for central government debt management for 2002 (see Swedish National Debt Office (2002)).

2 The macroeconomic simulation model

In this section we describe the macroeconomic simulation model which we use to simulate the development of a number of important financial and macroeconomic variables. The model is comparatively simple. It encompasses only three currency areas: Sweden, the United States, and the Euro zone. These are assumed to have stable economies, with cyclical swings between boom and recession roughly similar to those of the past thirty years, and where each respective central bank meets its annual inflation target, on average. In short, the model consists of six building blocks for each respective currency area. These building blocks model the *economic cycle regime*, *inflation*, *real*

[2] Recent studies in the academic field include Missale (1999), Missale (2000), and Missale (2001). The main results from our simulation model are qualitatively similar to some of the more general conclusions of these studies. Besides academic studies, practitioners at debt offices in many countries are also looking into the government debt strategy problem. Examples of this are Bolder (2002) and Danish National Bank (1998).

Figure 1: The model's variables and associations

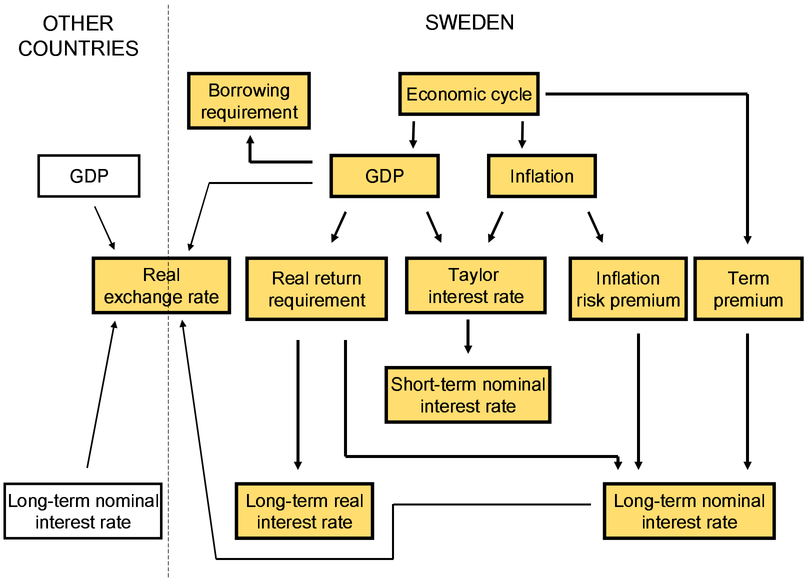

growth, *nominal short-term interest rate*, *nominal long-term interest rate*, and *exchange rates* for each respective currency area. For Sweden, there are two additional building blocks, one for the *real long-term interest rate* and one for the *net borrowing requirement*. Figure 1 provides a schematic picture of the model's variables and the associations between them.

In the following subsections, we describe each variable in more detail. But before we do that, it may be useful to present in more general terms the two modelling workhorses that we rely upon.

2.1 Stochastic processes

For a majority of the variables in the model, the modelling workhorse is a stationary *first-order autoregressive* ($AR(1)$) process. This is a simple but flexible way of modelling macroeconomic time series. The AR(1) process can for any variable y be written as

$$y_t = \alpha + \rho y_{t-1} + \varepsilon_t \tag{1}$$

where ρ is the autoregressive parameter which describes the effect of a unit change in y_{t-1} on y_t, and ε_t is a random shock. Throughout, ε_t is assumed to be normally and independently distributed with mean zero ($E[\varepsilon_t] = 0$), constant variance ($E[\varepsilon_t, \varepsilon_s] = \sigma_\varepsilon^2$ if $t = s$ and zero otherwise), and independent of y_{t-1} ($E[\varepsilon_t, y_{t-1}] = 0$).

To calculate the expected value of y s periods ahead one can use the following s-period-ahead forecast formula:

$$E[y_{t+s}] = \alpha \sum_{i=0}^{s-1} \rho^i + \rho^s y_t. \tag{2}$$

This formula is of particular interest in the case when s becomes very large. In that case the expected value of y can be interpreted as the long-run equilibrium of the series and can be used as such when calibrating the model. Letting s tend to infinity yields:[3]

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

A second feature of the model is that we allow the $AR(1)$ processes of some variables to change over the business cycle. Our way of modelling this is by using a *Markov chain switching regimes model*. The basic assumption of this model is that the $AR(1)$ process driving a variable y_t is influenced by an unobserved discrete random variable, s_t, which is called the *state* or *regime* that the process was in at time t. In practice, this means that the $AR(1)$ process has two different sets of parameter values depending on whether the economy is in a boom or a recession

$$y_t = \begin{cases} \alpha_B + \rho_B y_{t-1} + \varepsilon_t & if\ s = boom \\ \alpha_R + \rho_R y_{t-1} + \varepsilon_t & if\ s = recession \end{cases} \tag{4}$$

As a result, the variables have different expected values in these two states: $E[y_t | s = B] = \alpha_B / (1 - \rho_B)$ and $E[y_t | s = R] = \alpha_R / (1 - \rho_R)$.

[3] The derivation of this result is given in Appendix A.

Furthermore, the time series process driving the state variable is modelled as a 2-state Markow chain. It states that the probability that the state variable, s_t, equals some particular value j (in our case boom or recession) depends on the past only through the most recent value s_{t-1}:

$$P\{s_t = j \mid s_{t-1} = i, s_{t-2} = k, \ldots\} = P\{s_t = j \mid s_{t-1} = i\} = p_{ij}. \qquad (5)$$

For our setting with only two states, the transition probabilities, p_{ij}, are completely determined by a 2×2 matrix :

$$P = \begin{bmatrix} p_{BB} & p_{BR} \\ p_{RB} & p_{RR} \end{bmatrix}, \qquad (6)$$

where p_{BB} gives the probability that a boom state will be followed by a boom state, p_{BR} the probability that a boom state will be followed by a recession state, etc. Also note that $p_{BB} + p_{BR} + p_{RB} + p_{RR} = 1$.

The parameters of the $AR(1)$ processes, as well as the transition probabilities, can be estimated on historical data using the Hamilton (1989) procedure. When a model is parameterized more informally, one can also make use of the following relationship between transition probabilities, p_{ij}, and the expected duration of each state, D_i:[4]

$$D_B = \frac{1}{1 - p_{BB}}; \quad D_R = \frac{1}{1 - p_{RR}}. \qquad (7)$$

Given the parameterization of the two $AR(1)$ processes and the transition probabilities of the state variable, it is a straightforward exercise to simulate a regime switching $AR(1)$ process for any variable y_t. First one creates a series of error terms following the process $\varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right)$ for each $AR(1)$ process. Then one creates a series of state variables by generating a Bernoulli random number series on the basis of the transition matrix P. These two series can then be consecutively substituted into equation (4) in order to

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

calculate a series of "observed" values of y_t. It can be shown that the long-run equilibrium of this regime switching $AR(1)$ process is given by:

$$y^* = \left(\frac{\alpha_B}{1-\rho_B}\right)\left(\frac{1-p_{RR}}{2-p_{BB}-p_{RR}}\right) + \left(\frac{\alpha_R}{1-\rho_R}\right)\left(\frac{1-p_{BB}}{2-p_{BB}-p_{RR}}\right), \quad (8)$$

where $\alpha_j/\left(1-\rho_j\right)$ is the conditional mean of y_t in state j and $(1-p_{ii})/(2-p_{jj}-p_{ii})$ is the probability of being in state j. Thus, the long-run equilibrium of y_t is a weighted average of its conditional mean in each respective state. A proof of this result is outlined in Appendix A.

We now go on describing more specifically how the variables in the macroeconomic model are constructed. First we describe the structure of each time series process; then we present and comment on their parameterization.

2.2 Inflation

Inflation is modelled as an $AR(1)$ process without regime switching:

$$\pi_t = \alpha + \rho\pi_{t-1} + \varepsilon_t; \qquad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon\right). \quad (9)$$

The reason for this is that we assume that the central bank in each currency area is successful in its task of stabilizing inflation. The impact of economic states will therefore be entirely captured by the bank's reaction function for its key short-term interest rate. Consequently, inflation will be stable around the central bank's inflation target irrespective of economic developments.

The parameter values of equation (9) are presented in Table 1. Note that the simulation model is a quarterly model and that the parameters are adjusted to this. This implies that the observations must be aggregated over time in order to calculate the annual rate of inflation. The parameters are chosen so that the expected rate of inflation in each currency area is equal to the inflation target of the central bank. In other words, the central banks' inflation targets are assumed to be met on average. For Sweden and the

Table 1: Inflation

	Sweden	**EMU**	**US**
Annual average:	2.0%	1.5%	2.5%
Parameters			
α	0.00025	0.00019	0.00031
ρ	0.95	0.95	0.95
σ_ε	0.00062	0.00047	0.00078

euro zone, the inflation targets are consistent with the monetary policy rules declared by their central banks, even though the target is less explicit in the euro zone. For the US, the inflation target is assumed to be somewhat higher than for the other two economies, reflecting the US central bank's higher tolerance towards inflation.

Empirical estimates of equation (9) on historical data suggest that the inflation process is fairly persistent, with a ρ-parameter of about 0.95. The ρ-parameter is therefore set to 0.95 for all three economies. Furthermore, the estimated variance of the residuals corresponds to an annual standard deviation of inflation of 0.4 % for Sweden, 0.3 % for the euro zone, and 0.5 % for the US. The variances of the error terms in the simulation model are set accordingly.

2.3 Real GDP growth

Real GDP growth is assumed to follow a regime switching $AR(1)$ process:

$$g_t = \mu_s + \beta g_{t-1} + \varepsilon_t; \qquad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon\right); \quad s = B, R, \tag{10}$$

where g_t is the *percentage rate of growth* of real GDP and B and R represent the two states *boom* and *recession*. The change in regimes is captured by the constant μ_s, which indicates μ_B when the economy is in a boom and indicates μ_R when the economy is in a recession. We could also have modelled differences in the persistence of growth between the two regimes by using a state-dependent β-parameter, but for the sake of simplicity we only allow the

Table 2: Real GDP growth

	Sweden	**EMU**	**US**
Annual average:	2.5%	2.5%	3.1%
Parameters			
μ_B	0.00045	0.00042	0.00050
μ_R	−0.00028	−0.00016	−0.00024
β	0.95	0.95	0.95
σ_ε	0.00094	0.00094	0.00094
p_{BB}	0.95	0.95	0.95
p_{RR}	0.80	0.80	0.75

constant term to change across regimes. Note that the shifts between the two states become more gradual the closer the parameter β is to unity.

The basis for the parameterization of the growth equations is empirical data. The parameter values are presented in Table 2. The main feature is that the long-run real growth rate is higher in the US than in Sweden and the euro zone. It is also assumed that recessions are shorter in the US. The only difference between the growth processes in Sweden and the euro zone is that economic cycles are assumed to be more pronounced in Sweden, since it seems reasonable to assume that a small individual economy should experience bigger swings in economic cycles than a conglomerate of countries. The speed of adjustment from one period to another is assumed to be identical for all three economies.

2.4 Short-term nominal interest rate

Short-term nominal interest rates are modelled on the basis of a stylized monetary policy rule that the central bank is assumed to follow, the so-called Taylor rule.[5] This rule links the central bank's key interest rate to inflation and real GDP. The rule states that the central bank should raise its key rate if the expected inflation exceeds the targeted level, or if capacity

[5] In the strategy modelling part of the model, the short-term interest rate is interpreted as the three-month interest rate.

utilization in the economy exceeds its long-term sustainable level:

$$i_t^T = r^* + \tilde{\pi}_t^e + \theta\left(\tilde{\pi}_t^e - \tilde{\pi}^*\right) + \lambda\left(\ln Y_t - \ln Y_t^*\right). \qquad (11)$$

Here i_t^T is the Taylor interest rate, r^* is the equilibrium real interest rate, $\tilde{\pi}_t^e$ is the expected *annual* inflation rate, $\tilde{\pi}^*$ is the targeted annual inflation rate, and $(\ln Y_t - \ln Y_t^*)$ is the output gap, where Y_t is real GDP and Y_t^* is potential real GDP. The expected rate of inflation, $\tilde{\pi}_t^e$, is modelled under the assumption of adaptive expectations. This means that $\tilde{\pi}_t^e$ is assumed to be equal to the average annual rate of inflation during the last four quarters. Real GDP is calculated from its initial value and simulated growth path as $Y_t = Y_0 \prod_{j=1}^{t}(1 + g_j)$. Finally, potential real GDP is generated from real GDP by means of the so-called *Hodrick-Prescott (HP) filter*.[6]

For simulated series of $\tilde{\pi}_t^e$ and Y_t, equation (11) is used to construct a series of the Taylor interest rate. In the long term, $\tilde{\pi}_t^e = \tilde{\pi}^*$, $Y_t = Y_t^*$, and, hence, $i_t^T = r^* + \tilde{\pi}^*$. In the short term, however, the Taylor rate may swing sharply from one period to another. Therefore, the central bank is not assumed to follow the Taylor rule completely, but is assumed to adjust its key rate with a certain time lag. That is, the central bank practices some kind of interest rate smoothing. The short-term nominal interest rate is thus modelled as an $AR(1)$ process that adjusts gradually to the Taylor rate, giving short-term rates a more realistic pattern of movement:

$$i_t = \kappa + \varphi i_{t-1} - \zeta\left(i_{t-1} - i_{t-1}^T\right) + \varepsilon_t; \qquad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon\right) \qquad (12)$$

Technically, this is an error correction version of the standard $AR(1)$ process, where the speed of adjustment of the short rate towards the Taylor rate is

[6] The HP-filter is a smoothing method that is widely used among empirical macroeconomists to obtain a smooth estimate of the long-term trend component of a series. Technically, the HP-filter is a two-sided linear filter that computes the smoothed series Y_t^* of Y_t by minimizing the variance of Y_t around Y_t^*, subject to a penalty that constrains the second difference of Y_t^*. That is, the HP-filter chooses to minimize: $\sum_{t=1}^{T}(y_t - y_t^*)^2 + \lambda \sum_{t=2}^{T-1}\left[\left(y_{t+1}^* - y_t^*\right) - \left(y_t^* - y_{t-1}^*\right)\right]^2$.

Table 3: Short-term nominal interest rate

	Sweden	**EMU**	**US**
Average:	5.0%	4.5%	5.5%
Parameters			
θ	0.5	0.5	0.5
λ	0.5	0.5	0.5
κ	0.00250	0.00225	0.00275
φ	0.95	0.95	0.95
ζ	0.1	0.1	0.1
σ_ε	0.00145	0.00109	0.00109

governed by the ζ-parameter.

The parameters of the short-term nominal interest rates are presented in Table 3. The parameterization is chosen so that, in equilibrium, the short-term interest rate in each economy is equal to the Taylor rate in that economy. Differences in the equilibrium interest rate between the economies are due to the economies having different inflation targets. The Taylor rate process, as well as the speed of adjustment parameter, are both assumed to be identical for all three economies. The error term of the Swedish short-term interest rate is assumed to have a somewhat higher variance than the error terms of the other two economies, once again reflecting the smaller size of the Swedish economy.

2.5 Long-term interest rates

2.5.1 Real return requirement

Nominal long-term interest rates and yields on long-term inflation-linked bonds are determined on the basis of a real return requirement. The basic idea is that there is a real return requirement underlying any investment decision of an investor. The real return requirement, $\tilde{r}_t$, is assumed to be dependent on the capacity utilization in the economy as well as on its earlier values. It is modelled as an augmented $AR(1)$ process, where the output gap is included as an exogenous variable:

Table 4: Real return requirment

	Sweden	**EMU**	**US**
Average:	3.0%	3.0%	3.0%
Parameters			
ϖ	0.003	0.003	0.003
ϕ	0.9	0.9	0.9
ϑ	0.1	0.1	0.1

$$\tilde{r}_t = \varpi + \phi \tilde{r}_{t-1} + \vartheta \left(\ln Y_t - \ln Y_t^* \right). \tag{13}$$

The parameters of the real return requirement are presented in Table 4. They are set to the same values for all countries. The real return requirement is assumed to increase when GDP exceeds potential GDP and decrease when GDP falls below of potential GDP. Note that the real return requirement does not contain any disturbance term, other than the one included in GDP.

2.5.2 Yield on long-term inflation-linked bonds

The yield on inflation-linked bonds, r_t, is generated only for Sweden. It is determined by the sum of the real return requirement ($\tilde{r}_t$), a liquidity premium (lp_t), and an error term:

$$r_t = \tilde{r}_t + lp_t + \varepsilon_t; \qquad \varepsilon_t \sim NID \left(0, \sigma_\varepsilon \right). \tag{14}$$

The parameterization is presented in Table 5. The liquidity premium, lp_t, is set to 0.5 per cent in all time periods. Thus, the yield on inflation-linked bonds will exceed the real return requirement in the economy, on average.

Table 5: Yield on long-term inflation-linked bonds

	Sweden	**EMU**	**US**
Average:	3.5%	–	–
Parameters			
lp_t	0.005	–	–
σ_ε	0.00008	–	–

2.5.3 Long-term nominal interest rate

Nominal long-term interest rates, i_t^L, are created as the sum of the real return requirement ($\tilde{r}_t$), the expected annual inflation rate in the long term ($\tilde{\pi}^*$), an inflation risk premium (ip_t), a term premium (tp_t), and an error term (ε_t):

$$i_t^L = \tilde{r}_t + \tilde{\pi}^* + ip_t + tp_t + \varepsilon_t; \qquad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon\right). \tag{15}$$

Note that in the long term the expected annual inflation rate is equal to the central bank's inflation target. The expected annual inflation rate is therefore set to $\tilde{\pi}^*$ in this equation. The inflation risk premium, ip_t, is determined by earlier values of inflation. It is assumed to be higher if inflation has exceeded the central bank's inflation target during the last four quarters, and/or if inflation, measured on a quarterly basis, has varied a lot:

$$ip_t = \tau + \xi \max\left[\left(\tilde{\pi}_t - \tilde{\pi}^*\right), 0\right] + \chi \sum_{s=1}^{4} \left(\pi_{t-s} - \pi^*\right)^2. \tag{16}$$

The term premium, tp_t, is assumed to follow a regime switching $AR(1)$ process, with a higher expected value in booms than in recessions:

$$tp_t = \eta_s + \varsigma tp_{t-1}; \qquad s = B, R. \tag{17}$$

The parameterization of the long-term nominal interest rate, the inflation risk premium, and the term premium are presented in Table 6.

Table 6: Long-term nominal interest rate

	Sweden	**EMU**	**US**
Average:	5.9%	5.5%	6.6%
Parameters			
σ_ε	0.0016	0.0016	0.0016
Inflation risk premium			
Average:	0.4%	0.4%	0.4%
Parameters			
τ	0.0025	0.0025	0.0025
ξ	0.2	0.2	0.2
χ	0.2	0.2	0.2
Term premium			
Average:	0.6%	0.6%	0.7%
Parameters			
η_B	0.002	0.002	0.002
η_R	−0.002	−0.002	−0.002
ς	0.8	0.8	0.8

2.6 Exchange rates

Real exchange rates are modelled as augmented $AR(1)$ processes:

$$E_t^{i,j} = \tau + \psi E_{t-1}^{i,j} - \nu \left(g_t^i - g_t^j\right) - \omega \left(i_t^{Li} - i_t^{Lj}\right) + \varepsilon_t, \tag{18}$$

where $E_t^{i,j}$ is the real exchange rate between economies i and j, g_t^i is the growth rate in country i, g_t^j is the growth rate in country j, i_t^{Li} is the long-term interest rate in country i, i_t^{Lj} is the long-term interest rate in country j, and ε_t is an error term distributed as $NID\left(0, \sigma_\varepsilon\right)$. The ν-parameter captures short-term effects of differences in rates of growth. If growth is stronger in one of the economies, its currency tends to appreciate against the other economy. In equilibrium, the presence of this term will induce a trend in the real exchange rate if $g^{*i} \neq g^{*j}$. This effect is consistent with the literature on real exchange rates, which tells us that in equilibrium real exchanges rates may trend according to differences in productivity. The ω-parameter captures

Table 7: Real exchange rates

	EUR/SEK	**USD/SEK**
Starting value:	8.0%	9.0%
Parameters		
τ	0.09	0.10
ψ	0.99	0.99
ν	2.5	2.5
ω	0.25	0.25
σ_ε	0.07053	0.07053

differences in long-term nominal interest rates. The effects are similar to the ones discussed above.

Nominal exchange rates, $e_t^{i,j}$, are created by adding or subtracting differences in inflation rates between the countries in question. Specifically, $e_t^{i,j} = e_{t-1}^{i,j}\left(1 + \log E_t^{i,j} - \log E_{t-1}^{i,j} + \pi_t^i - \pi_t^j\right)$. Long-term differences in inflation will induce the same type of long-term trends in nominal exchange rates as differences in growth rates and interest rates do in real exchange rates.

The processes governing real exchange rates are very persistent and the autoregressive parameter, ψ, is therefore set to 0.99 for both EUR/SEK and USD/SEK (see Table 7). This means that the degree of mean reversion in real exchange rates is quite low. The adjustment parameters ν and ω are set on an *ad-hoc* basis and are assumed to be equal across the two real exchange rates. This is also the case for the variance term. The intercept, τ, is chosen so that real exchange rates, in equilibrium and in the absence of growth and interest rate differences, are equal to the exchange rates that prevailed when the simulation model was parameterized (May 2001).

Note that the basic parameterization makes no assumption about either real or nominal exchange rate trends, except those that follow from cross-country differences in potential GDP, long-term interest rates, and inflation. Given these assumptions, the structure and parameterization of the model implies a certain strengthening of the krona against the dollar in nominal terms, while the krona weakens against the euro.

2.7 Net borrowing requirement

The specification of the net borrowing requirement of the Swedish central government is highly stylized. It is based on the official fiscal policy target of a two per cent surplus in public finances over an economic cycle, plus our assumption that the pension system will show a surplus that exceeds these two per cent. Taken together, this means that the government needs to *borrow* when the economy is in equilibrium. The *equilibrium level* for the net borrowing requirement in the period t can therefore be expressed as

$$B_t = \gamma \left(1 + g^* + \pi^*\right)^t Y_0,$$

where γ is a positive, constant percentage and $\left(1 + g^* + \pi^*\right)^t Y_0$ is the equilibrium level for nominal GDP in period t.

A measure of economic activity, $(g_t - g^*)$, is then added to this equation to capture how capture how the government's *actual* net borrowing requirement varies with the economic cycle:

$$B_t = \gamma \left(1 + g^* + \pi^*\right)^t Y_0 - \delta \left(g_t - g^*\right) + \varepsilon_t; \qquad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon\right),$$

where $(g_t - g^*)$ is the difference between real growth and equilibrium real growth, and ε_t is an error term representing any random fluctuation not explicitly included here. The actual net borrowing requirement during a given period will therefore exceed or fall short of its equilibrium level, depending on whether the pace of economic growth is faster or slower than the equilibrium growth rate. In equilibrium, however, $E\left[g_t\right] = g^*$ and, hence, $E\left[B_t\right] = \gamma \left(1 + g^* + \pi^*\right)^t Y_0$, which means that despite the fluctuations, the target borrowing requirement will be met on average.

The parameterization of the net borrowing requirement is presented in Table 8. The parameter γ is chosen so that, in equilibrium, the net borrowing requirement is equal to 0.5 per cent of GDP per annum, reflecting our stylized assumption about the division of the targeted surplus in public finances

Table 8: Net borrowing requirement

	Sweden	**EMU**	**US**
Average:	0.5%	–	–
Parameters			
γ	0.00125	–	–
δ	2200	–	–
σ_{ε}	0.05	–	–

between central government and other parts of the public sector. This implies that central government debt will increase over time in nominal terms, but decrease as a share of GDP, since GDP itself will increase by more than 0.5 per cent on average. The parameter δ, which captures the sensitivity of government finances to business cycles fluctuations, are set to 2200.

3 The strategy simulation model

The goal of the simulation exercise is to compute the costs and risks associated with different debt strategies over a long period of time. In addition to the simulation of economic and financial variables we therefore need a tool to simulate different debt management strategies. The task of the strategy simulation is to roll-over the debt portfolio in each simulation step, thereby refinancing maturing debt in a way that makes the portfolio meet certain criteria. In each step, the strategy simulation determines what amounts should be borrowed in nominal domestic bonds, foreign currency bonds and inflation-linked bonds, as well as the maturity distribution of each debt type. The strategy simulation also keeps track of all individual cash flows, and calculates costs and risks. This section describes the strategy simulation part of the model. First, a couple of general topics about strategy simulations are discussed. Then the implementation of the strategy simulation in the SNDO model is described in more detail.

3.1 General topics

Strategy simulations of central government debt management raises two initial questions. First, in what terms should debt strategies be defined? Second, which debt portfolio should be used as the starting portfolio in each strategy simulation? Both these questions are briefly addressed below. For a more elaborate discussion, see Bergström & Holmlund (2000).

3.1.1 Formulating debt strategies

In principle, there are two main alternatives for the formulation of strategies. The first is to define the strategy as a refinancing scheme. The second alternative is to define a strategy in portfolio terms, e.g. with targets for the allocation and duration of the debt portfolio.

In the SNDO model, the latter approach is used. The strategies analyzed are defined in terms of duration and shares of nominal, inflation-linked and foreign currency debt. This is not necessarily because the portfolio approach is superior to the refinancing scheme approach. Indeed, the portfolio approach has some less appealing effects, primarily concerning foreign currency and inflation-linked debt, that the first approach overcomes. The main reason for using the portfolio approach, however, is that this is the way the Debt Office usually characterizes its debt portfolio, both internally and externally.[7]

3.1.2 The starting portfolio

At first sight it may seem logical to start each strategy simulation with the government's actual debt portfolio. However, having all strategies start in the same portfolio makes the costs of different strategies rather similar during the first parts of the simulation, which in turn makes the differences between strategies less clear cut. Since the aim of the simulations is to analyze long-term differences in costs and risks, rather than the transition from one

[7]It should be noted that the guidelines are not entirely portfolio based, since there is a target for the level of amortization of the foreign currency debt, rather than for its share.

portfolio to another, we therefore let all strategies start in a portfolio fulfilling the specified strategy right from the beginning. In this way, we can be sure that differences in cost between any two strategies are due to different debt structures, and not on the transactions needed to shift from one portfolio to another.

3.2 Strategy simulations

3.2.1 Strategies

The goal of the analysis is to examine the costs and risks of different portfolio strategies in rough terms. For that purpose it is sufficient to have a small number of strategies, although they should be clearly differentiated and span a relatively large space. In the initial simulations, nine portfolio strategies were examined. They are displayed in Table 9. In these simulations the share of inflation-linked bonds varied between 0, 10 and 20 percent. In a second set of simulations these intervals were increased to 0, 50 and 100 percent. The shares of nominal domestic and foreign currency debt in total government were varied accordingly. Whenever there was foreign currency debt in the portfolio, it was comprised of 70 percent euro and 30 percent US dollars, which is roughly the same as in the actual portfolio. The duration target for the nominal domestic debt was varied between 2, 3 and 4 years, while the duration of the foreign currency and the inflation-linked debt was set at 2 and 10 years, respectively.

Table 9: Portfolio strategies in base scenario

Shares of total debt			Duration, nominal krona debt		
			2	3	4
0% IL	75 % SEK	25% FX	Strategy 1	Strategy 2	Strategy 3
10% IL	65 % SEK	25% FX	Strategy 4	Strategy 5	Strategy 6
20% IL	55 % SEK	25% FX	Strategy 7	Strategy 8	Strategy 9

3.2.2 Treatment of foreign currency and inflation-linked bonds

In the model, foreign currency bonds and inflation-linked bonds are treated in the same way. All cash flows from foreign currency bonds are converted to Swedish krona using the appropriate FX rate. Likewise, the real cash flows from an inflation-linked bond are converted to nominal Swedish krona using the simulated CPI. In other words, CPI is viewed as the exchange rate between nominal and real money. In the following, the terms foreign currency and FX rate will sometimes be used both when referring to foreign currency and inflation-linked debt.

For an inflation-linked bond, almost all of the inflation compensation is paid out at maturity. This is equivalent with how the Swedish inflation-linked bonds work in reality, but can nevertheless lead to large peaks in costs for inflation-linked debt in the model.

3.2.3 The initial portfolio

Technically, the initial portfolio is held in a two dimensional matrix. This matrix has two columns for each debt type, one containing the debt principals and one containing the coupon payments. Since we are modelling portfolios with four different debt types, the matrix has eight columns in total. Moreover, the matrix has 120 rows, corresponding to 120 quarterly maturity buckets. The amounts in the first row of the matrix mature in the first quarter (step) of the simulation run, the amounts in the second mature in the second quarter and so on.

At the start of each simulation run the total outstanding debt is distributed over all maturity buckets in a way that is consistent with the duration and allocation targets of the simulated strategy. The amount of outstanding debt at the start of the simulation is set to 1180 billion SEK, approximately the same as the actual size of the debt. The initial coupon of the debt is set according to the initial yield curve of each debt type. It is thus implicitly assumed that all outstanding debt at the start of the simulation run has been raised just prior to the start of the simulation.

3.3 The strategy simulation step-by-step

In each simulation run, the portfolio is rolled over 30 years in quarterly steps, i.e. 120 times. Each roll-over consists of a number of steps. These are described in detail below.

3.3.1 Calculating coupon payments

In each simulation step, the first number calculated is the sum of all coupon payments due in that period. This number is calculated simply as the coupon payment in each currency, times the simulated FX rates for that particular period. The coupons are found in columns 2, 4, 6 and 8 of the portfolio matrix in Figure 2.

Figure 2: Calculating the coupon payments

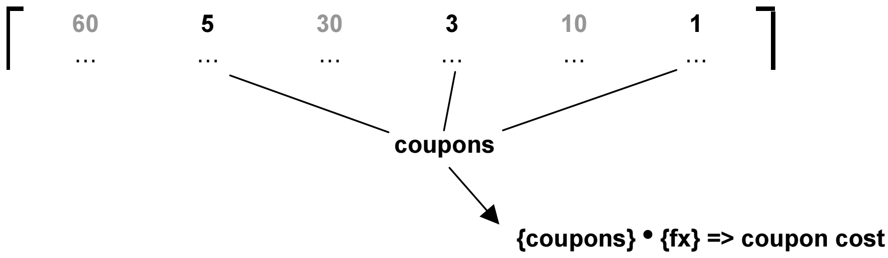

3.3.2 Calculating gross borrowing requirements

The borrowing requirement for the specific period is calculated in a similar way: the amounts maturing are found in columns 1, 3, 5 and 7 of the portfolio matrix in Figure 3, and they are converted to nominal Swedish krona at the simulated FX rates. This gives the borrowing requirement stemming from the debt itself. There is also a net borrowing requirement arising from the government's budget in the simulated economy. This budget deficit/surplus is added to the borrowing requirement from maturing bonds, giving the total financing need. Finally, the total financing need is distributed over the four currencies according to the allocation target of the strategy.

Figure 3: Calculating gross borrowing requirments

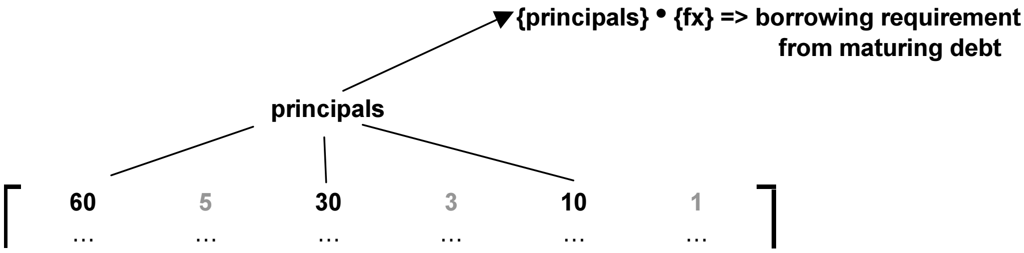

3.3.3 Financing gross borrowing requirements

The next step in the strategy simulation is to let the shortest bonds mature. This is done simply by 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 as it would look after the maturities of the current period have occurred, but before any new bonds have been issued.

The borrowing need must be financed in a way that makes the portfolio meet the duration target, if possible. To be able to determine at which duration point we should borrow, we need to know the *market value* and *duration* of the debt outstanding in each of the four currencies. To calculate these we need four sets of discount factors, one for each currency. In the economic/financial simulation, we simulate only two points on the yield curve, the three-months months rate and the ten-year rate.[8] To get the discount factors we therefore linearly interpolate yields for each maturity bucket, and then convert them to discount factors in the standard way:

$$df_t = \frac{1}{(1 + r_t)^t}. \tag{19}$$

Since we have the cash flows of the portfolios organized in maturity buckets, and know the time to maturity of each bucket, it is now straightforward to calculate the market value, mv_t, of each part of the debt (Swedish krona,

[8] For inflation-linked bonds, we simulate the five and thirty year points on the yield curve.

euro, dollars and inflation-linked) as

$$mv_t = \sum cf_t \cdot df_t, \tag{20}$$

where t denotes the time (in quarters) to maturity of each maturity bucket and runs from 1 to 120; cf_t denotes the cash flow in bucket t and df_t is the discount factor associated with maturity t. With the same notation, the duration of each part of the debt can be calculated as[9]

$$duration = \frac{\sum cf_t \cdot df_t \cdot t}{mv_t}. \tag{21}$$

Note that the duration of the different parts of the debt are not aggregated, since the duration targets are specified per currency, and not for the portfolio as a whole.

Given the market value and duration of the debt outstanding before new funding, it now possible to determine at which duration point we should finance the present borrowing need. 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:

$$d^* = \frac{d^0 mv^0 + d^a br}{mv^0 + br} \Longrightarrow d^a = \frac{d^* \left(mv^0 + br\right) - d^0 mv^0}{br}. \tag{22}$$

The entire 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 the nominal amount. In nominal Swedish krona, euros and dollars par bonds with one or ten year maturities

[9] 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 to measures is small: with an upward sloping yield curve, Fisher-Weill method will give a slightly lower duration than the Macauley method.

are used. For inflation-linked debt, five and thirty years maturities are used. Given the discount factors calculated from the simulated interest rates, the coupon c of a par bond maturing in N years can be calculated using the following equation:

$$100 = c\sum_{i=1}^{N} df_i + 100df_N \Longrightarrow c = \frac{(1 - df_N)\,100}{\sum_{i=1}^{N} df_i} \tag{23}$$

The duration of the par bonds is calculated using equation (21).

For each currency we now have the duration and coupons of par bonds maturing 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 usually done using both par bonds in each currency. 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 superindices S and L denote the short-term and long-term par bond, respectively. The sum of nominal amounts in the two bonds must equal the borrowing requirement br. The amounts that should be allocated to short-term and long-term bonds, N^S and $N^L = \left(br - N^S\right)$, are then given by:

$$d^* = \frac{d^S N^S + d^L\left(br - N^S\right)}{br} \Longrightarrow N^S = \frac{\left(d^* - d^L\right) br}{\left(d^S - d^L\right)} \tag{24}$$

and

$$N^L = \left(br - \frac{\left(d^* - d^L\right) br}{\left(d^S - d^L\right)}\right) \Longrightarrow N^L = \frac{\left(d^S - d^*\right) br}{\left(d^S - d^L\right)}. \tag{25}$$

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.

3.3.4 Bond repurchases

In most cases the gross borrowing requirement is larger than zero – we need to borrow in order to refinance maturing bonds and/or a budget deficit. However, sometimes it may happen that the gross borrowing requirement is less than zero. This happens when the simulated economy shows a budget surplus that is 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 only be made among those bonds that are currently in the portfolio. Furthermore, individual cash flows of a bond cannot be bought back separately, but must be handled as a whole. In extreme cases, it may happen that all of the outstanding debt is bought back. When no bonds are left to buy back, any remaining surpluses are deposited in the cash market for one quarter at the time. For a more detailed description of how bond repurchases are handled in the model, see Bergström & Holmlund (2000).

4 Cost and risk measures

In the SNDO simulation model, the costs of debt are related to (simulated) nominal GDP in order to get a measure of the real costs debt. Risk is defined as the variation of these costs and is measured in two ways. This section gives a detailed description of how costs are calculated in the model, how they are aggregated over the simulation horizon, and how risks are defined and calculated.

4.1 Nominal costs

Nominal costs are measured on a cash flow basis, meaning that costs only occur when money is paid out. Thus, unrealized mark-to-market effects from varying interest rates or FX rates are not included. The reason for this is that we are only concerned about the costs that actually affect the government's budget. Since the debt is long term and bonds are normally left outstanding until maturity, swings in mark-to-market value are in most cases not materialized. Therefore, mark-to-market costs are viewed as less relevant for the decision on the long term composition of the debt.

In the model, three events affect the costs of a strategy:

1. *Coupon payments.* The cost for the strategy increases with the coupon payment in local currency times the simulated FX rate.

2. *Redemption of foreign currency loans and inflation-linked bonds.* At maturity the difference between the amount received in base currency at the time the bond was issued and the base currency amount needed to redeem the bond is added to costs. For foreign currency bonds, this equals the difference between the FX rates at the two points in time, times the nominal amount. This effect can either increase or decrease costs, depending on how exchange rates have developed. For inflation-linked bonds the inflation compensation on the nominal amount of the bond, i.e. the accumulated inflation between issuance and maturity, is paid out at maturity. This effect normally increases costs, since inflation is practically always positive.

3. *Repurchases of existing bonds.* When coupon bonds are bought back, mark-to-market gains or losses are realized. Although the focus of the SNDO simulation exercise is not on mark-to-market costs, it would be inappropriate not to take these realized 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. The mark-to-market cost is not periodised in any way, but affects the debt costs by the full amount in the period in which it occurs. The effect on costs can be both positive and negative, depending on the relation between the coupon of the bond bought back and the present level of interest rates.

4.2 Real costs

Real costs are defined as the ratio of nominal costs to nominal GDP. To calculate the annual real cost of debt we do as follows. The nominal cost of a certain strategy is calculated in each simulation period. At the end of each simulation run, we thus have 120 quarterly cost figures. These are converted to annual figures as the sum of four quarterly costs, resulting in 30 annual cost figures. In order to obtain annual real costs of this strategy, these annual nominal cost figures are then divided by annual nominal GDP. For each step in the simulation, GDP is generated at an annual rate. In order to get GDP for a specific year we must therefore take the average of four annual rates. To be exact, annual real costs, C_i^Y, are calculated as:

$$C_i^Y = \frac{\sum_{j=1}^{4} C_{i,j}}{\frac{1}{4}\sum_{j=1}^{4} Y_{i,j}}; \qquad i = 1, 2, \ldots, 30, \tag{26}$$

where $C_{i,j}$ is the cost in quarter j of year i and $Y_{i,j}$ is the annual rate of nominal GDP in the same quarter and year.

For each simulation run the average annual real cost, $\bar{C}^Y$, of the strategy is also calculated. At the end of our 1,000 simulation runs we thus have 1,000 estimates of the average annual real cost of the strategy. The expected cost of the strategy is finally defined as the median (the 50th percentile value) of the simulated distribution of $\bar{C}^Y$:

$$E\left[\bar{C}^Y_{strategy}\right] = P_{50}\left(\bar{C}^Y\right) \tag{27}$$

The reason for why we use the 50th percentile instead of the mean is that this corresponds better to our risk measures, which is described below.

4.3 Risk measures

Risk is measured in two dimensions. In the first dimension, risk is measured as the variability of possible outcomes over a long period of time. We refer to this risk dimension as *scenario risk*. In the other dimension, risk is measured as the variability of costs between different years for a given economic development. We refer to this as *time series risk*. One might say that scenario risk expresses how bad things can go given that we do not know which one of the 1,000 simulated economic paths will be realized, whereas time series risk expresses how much annual costs can be expected to vary over time for an arbitrary economic course of events. It is not evident which of these risk measures is the more interesting one. In the long term, scenario risk is probably the risk that the government is more eager to control. In a shorter perspective, however, large swings in debt service costs from one year to another may be just as undesirable.

4.3.1 Scenario risk

Scenario risk is calculated as follows. In each simulation run the average annual real cost over 30 years, $\bar{C}^Y$, is calculated as described above, giving a distribution of 1,000 outcomes for the average cost of the strategy. The scenario risk then measures the width of this simulated distribution. Specifically, it is calculated as the relative distance between the 95th and 50th percentiles of the simulated distribution:

$$R_{SC} = \frac{P_{95}\left(\bar{C}^Y\right)}{P_{50}\left(\bar{C}^Y\right)} - 1 \qquad (28)$$

The reason we use this relative percentile distance is that the percentile distance in absolute terms tends to be bigger, the higher the expected cost is. With an absolute measure, high cost strategies therefore also tend to look riskier, even if the relative variation in cost may be much lower.

Figure 4 shows a graphical illustration of the scenario risk measure. The oscillating lines represent simulated paths for the annual cost to GDP ratio. The three dotted horizontal lines are the corresponding averages over 30 years. In this picture, the scenario risk would be given by $R_{SC} = \left(\bar{C}_1/\bar{C}_2\right) - 1$.

Figure 4: Scenario risk

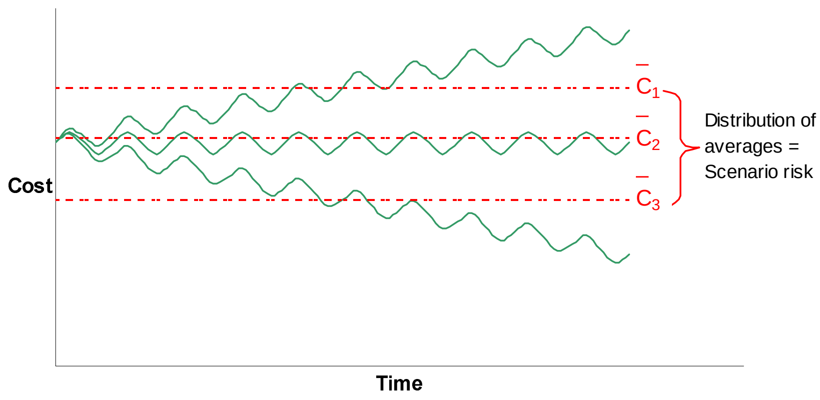

4.3.2 Time series risk

Time series risk is calculated as follows. In each simulation run annual real costs, C_i^Y, are calculated as described above. A straight line is fitted to this simulated cost path using ordinary least squares (OLS), giving a vector of estimated costs along a linear trend, L_i^Y. For each simulation run k, a vector of absolute deviations of costs from the linear trend is then calculated as

$$A_k = \left| C_i^Y - L_i^Y \right|_k .$$

The time series risk for a strategy is finally defined as the average relative distance between the 95th and 50th percentiles for the absolute deviation in each simulation run:

$$R_{TS} = \frac{1}{1000} \sum_{k=1}^{1000} \left(\frac{P_{95}\left(A_k\right)}{P_{50}\left(A_k\right)} - 1 \right).$$

The reason we measure the variability of costs as deviations from a linear trend, and not in annual real costs directly, is that the latter would make strategies generating a trend in costs look worse than they actually are. For instance, a strategy generating a stable downward trend in costs would have a high relative percentile distance simply because it covers a larger range of annual costs, rather than because annual cost fluctuates a lot. Graphically, time series risk can be illustrated as in Figure 5.

Figure 5: Time series risk

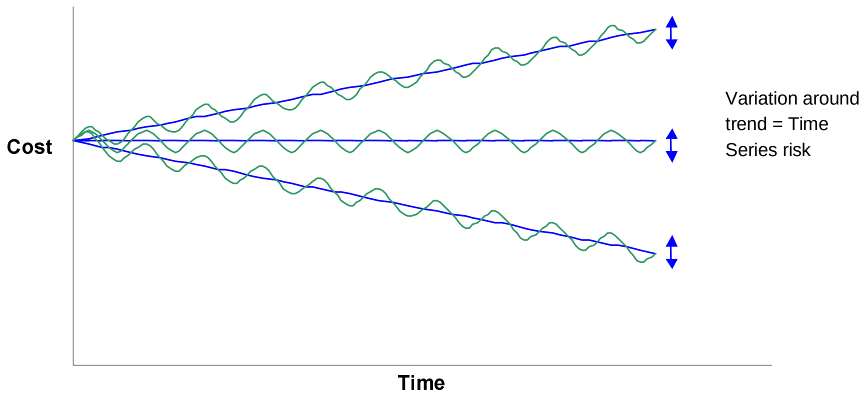

4.4 Simulated cost paths

The cost paths which are actually generated by the model are unfortunately not as typical and clear-cut as the ones in the figures above. Instead, the cost paths are generally downward sloping (reflecting the fact GDP grows at a higher pace than government debt) and not very volatile. Moreover, the paths from portfolios with different shares of inflation-linked debt are often fairly similar. Figure 6 displays three simulated cost paths of portfolios with 0, 50 and 100 per cent inflation-linked debt, respectively (the remainder of debt being nominal Swedish krona debt).

Figure 6: Cost paths for portfolios with 0, 50 and 100 percent inflation-linked debt (replication number 421, base parametrisation)

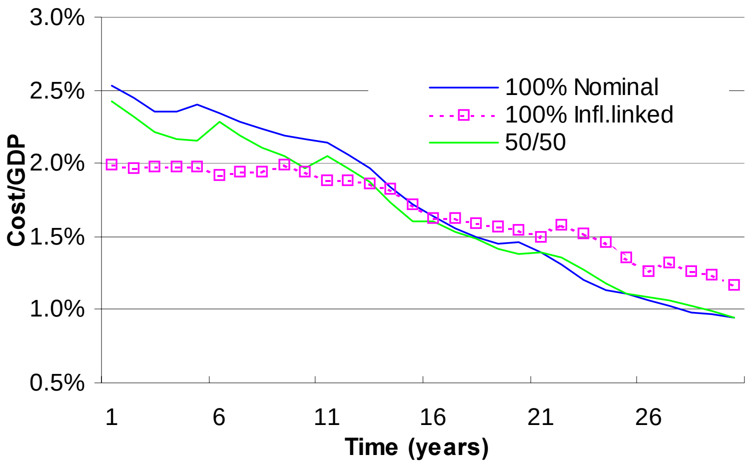

5 Implementation and timing issues

5.1 Implementation

The SNDO model is implemented in Microsoft Excel and Visual Basic for Applications (VBA). This programming language offers a simple and flexible environment, and knowledge about it is widespread in the financial community. The main drawback of using it is that it is slower than many other programming languages. However, much of the slowness of Excel lies in the spreadsheet part of the application, whereas the VBA language is reasonably fast. Since all the calculations in the SNDO simulation model are done in VBA, the time needed to run the model can be kept at a manageable level. It is also very easy to make use of external code resources in VBA through COM technology, e.g. for intensive calculations. In the current implementation we have used Matlab in this way for some smaller parts of the simulation.

The simulation algorithms are very loop intensive. For instance, in every step of the strategy simulation there a several loops through a cash flow matrix. Loops like these could be handled in a very elegant and efficient way with a vector based language like Matlab. Matlab also has a lot of built-in functions that can be utilized for random number generation, interpolation, curve fitting etc.

For future development it would therefore be very interesting to rewrite some of the most time-consuming parts of the program in Matlab and use it through the above mentioned interface to Excel. In this way, we could use the graphical interface strengths of Excel, while having Matlab's stronger calculation engine behind.

5.2 Timing

The simulations were run on PIII 733 MHz / 384 MB desktop computers. A typical simulation with a thousand runs for one strategy over 30 years with quarterly steps takes about 25 minutes. The 12 strategies that were analyzed in this years work thus took approximately five hours to run.

It is interesting to analyze in which parts of the code most time is spent. Perhaps contrary to intuition, random number generation and simulation of economic and financial variables does not take up a major part of the simulation time. Instead, more than 90 per cent of the simulation time is spent handling the portfolio strategy and computing cost and risks. This reflects the rather detailed nature of the strategy simulation. Table 10 summarizes the accumulated time spent in different parts of the simulation.

The main reason that so much of the time is spent in the portfolio parts of the simulation is that these routines are called much more often than the routines in the economic modelling. For example, duration figures and debt composition for the 30-year portfolio are calculated twice in each of the 120 quarterly steps in every replication, whereas the economic development is only simulated once for that same replication.

5.3 Random number generation

Uniform random numbers are generated with VBA:s built-in random number generator. The uniform random numbers (u) are transformed into normal

Table 10: Time spent in different parts of the simulation

Part of simulation		**% of time**	
Generate normal random numbers		0.4%	
Generate economic regimes		0.0%	
Simulate economic development		4.2%	
Handle portfolio and calculate costs & risks		94.0%	
of which	Calculate yield curves		12.1%
	Roll portfolio		18.7%
	Compute portfolio duration and allocation		48.1%
	Allocate borrowing & Update portfolio		15.2%
	Other		5.9%
Other		1.3%	

deviates (z) using the simple standard algorithm:

$$z = \sum_{i=1}^{12} u_i - 6 \tag{29}$$

It should be noted that the algorithm used in the VBA random number generator is not the same as the one used in Excels worksheet function $RAND()$. Even though the VBA generator is better than the corresponding worksheet function, we are aware that the VBA random number generator may not be state-of-the-art. It is a linear congruential generator of the form[10]

$$u_{i+1} = (u_i \cdot a + c) \text{ mod } 2^{24}, \tag{30}$$

where $a = 114,0671,485$ and $c = 12,820,163$. The main drawback of this generator for our simulation exercise is that the sequence of random numbers it produces is too short compared to the number we use. The period of the VBA generator is $16,777,216$, which is of the same order as the number of uniform random numbers we need for a thousand simulation runs. According to Press, Teukolsky, Vetterling & Flannery (1992), a rule of thumb is that the

[10] See Microsoft Knowledge Base article Q231847, available at http://support.microsoft.com/.

period length should be at least 20 times the number of random numbers used by the application. Although unclear how or if this affects our results, another random number generator should probably be used for future simulations.

6 Directions for future work

Over the last few years, the Debt Office has in its annual proposals for guidelines for the management of central government debt covered the important dimensions of government debt characteristics, i.e. structure and maturity. This does not mean that the Debt Office's analytical and modelling work has been completed. Further steps should be taken towards a more systematic ALM approach to government debt and its structure. Among other things, this should include improved descriptions of the structural and cyclical factors that determine the borrowing requirement — and thus government debt. One important question in this context is how the primary borrowing requirement is affected by developments in the economy, since the interplay between interest rates on the government debt and the primary borrowing requirement is what determines the fluctuations in the total borrowing requirement. Today's simulation model does not include such feedback in the net borrowing requirement equation. As we plan our continued development work, we welcome opinions and comments on the model from both researchers and practitioners in the field of government debt management.

References

Bergström, P. & Holmlund, A. (2000). A simulation model framework for government debt analysis. Technical Report, Riksgäldskontoret (Swedish National Debt Office).

Bolder, D. J. (2002). Towards a more complete debt strategy framework. Bank of Canada Working Paper, forthcoming.

Danish National Bank (1998). The interest and refinancing risk on domestic government debt - Cost-at-Risk. Discussion Paper.

Hamilton, J. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* .

Hamilton, J. (1994). *Time Series Analysis*, Princton University Press., New Jersey.

Kim, C.-J. & Nelson, C. (1999). *State Space Models with Regime Switching*, MIT Press, Cambridge, Massachusets.

Microsoft (2001). INFO: How visual basic generates pseudo-random numbers for the RND function. Microsoft Knowledge Base, article Q231847. Available at http://support.microsoft.com/.

Missale, A. (1999). *Public Debt Management*, Oxford University Press, Oxford.

Missale, A. (2000). How should the debt be managed? supporting the stability pact, *in* B. M. Brunila, A. & F. Franco (eds), *The Stability and Growth Pact - The Architecture of Fiscal Policy in EMU.*

Missale, A. (2001). Optimal debt management with a stability and growth pact, *Public Finance and Management* **1**.

Press, W., Teukolsky, S., Vetterling, W. & Flannery, B. (1992). *Numerical Recipes in C*, Cambridge University Press.

Swedish National Debt Office (2001). Central government debt management - proposed guidelines.

Swedish National Debt Office (2002). Central government debt management - proposed guidelines.

A Appendix

This appendix presents the derivations of some of the results appearing in the main text.

A.1 Derivation of Equation (3)

Consider the following $AR(1)$ process of y_t:

$$y_t = \alpha + \rho y_{t-1} + \varepsilon_t; \qquad \rho < |1|, \quad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right)$$

The unconditional mean of y_t can be derived as follows:

$$\begin{aligned}
y_t &= \alpha + \rho y_{t-1} + \varepsilon_t \\
&= \alpha + \rho\left(\alpha + \rho y_{t-2} + \varepsilon_{t-1}\right) + \varepsilon_t \\
&= (1+\rho)\,\alpha + \rho^2 y_{t-2} + \varepsilon_t + \rho\varepsilon_{t-1} \\
&= (1+\rho)\,\alpha + \rho^2\left(\alpha + \rho y_{t-3} + \varepsilon_{t-2}\right) + \varepsilon_t + \rho\varepsilon_{t-1} \\
&= \left(1+\rho+\rho^2\right)\alpha + \rho^3 y_{t-3} + \varepsilon_t + \rho\varepsilon_{t-1} + \rho^2\varepsilon_{t-2} \\
&\vdots \\
&= \left(1+\rho+\ldots+\rho^{t-1}\right)\alpha + \rho^t y_0 + \varepsilon_t + \rho\varepsilon_{t-1} + \ldots + \rho^{t-1}\varepsilon_{t-(t-1)} \\
&= \alpha\sum_{s=0}^{t-1}\rho^s + \rho^t y_0 + \sum_{s=0}^{t-1}\rho^s\varepsilon_{t-s} \\
E\left[y_t\right] &= \alpha\sum_{s=0}^{t-1}\rho^s + \rho^t y_0 \\
Let\ t &\rightarrow \infty,\ then, \\
E\left[y_t\right] &= \frac{\alpha}{1-\rho}
\end{aligned}$$

using the standard result for a geometric series.

A.2 Derivation of Equation (8)

Suppose we want to make model consistent forecasts of a time series y_t,following a 2-state $AR(1)$ process:

$$y_t = \begin{cases} \alpha_1 + \rho_1 y_{t-1} + \varepsilon_t \\ \alpha_2 + \rho_2 y_{t-1} + \varepsilon_t \end{cases} ; \qquad \rho < |1|, \quad \varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right)$$

$$P\{s_t = j \mid s_{t-1} = i\} = p_{ij}$$

We then need forecasts on both the future state of the process, s, as well on the expected value of y_t 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.,

$$E\left[y_{t+m}|\mathbf{y}_t, s\right] = \alpha_s \sum_{i=0}^{m-1} \rho_s^i + \rho_s^m y_t; \qquad s = 1, 2, \tag{31}$$

where $\mathbf{y}_t$ denotes the entire realization of y up until period t. By stacking each forecast for all periods ahead into a 2×1 vector $\mathbf{h}_{t,m}$, we will obtain the following expression for the forecast:[11]

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

The main result we need here, however, concerns the long run equilibrium of y. From (31) it is evident that $\rho_s^m y_t$ tends to zero as m tends to infinity,

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

and that $\alpha_s \sum_{i=0}^{m-1} \rho_s^i$ simultaneously tends to $\alpha_s / (1-\rho_s)$. Thus, the long run $\mathbf{h}$ is defined by:

$$\mathbf{h}'_{LR} = \left[\frac{\alpha_1}{1-\rho_1}, \frac{\alpha_2}{1-\rho_2} \right].$$

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. $\boldsymbol{\iota}'\boldsymbol{\xi} = 1$. Note that the first condition implies that $(\mathbf{I} - \mathbf{P})\boldsymbol{\xi} = \mathbf{0}$ and that we therefore may write the problem as the following system

$$\mathbf{A}\boldsymbol{\xi} = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}$$

where $\mathbf{A} = [(\mathbf{I} - \mathbf{P}), \boldsymbol{\iota}]'$. Solving for $\boldsymbol{\xi}$ yields

$$\boldsymbol{\xi} = (\mathbf{A}'\mathbf{A})^{-1} \mathbf{A}' \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}.$$

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

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

Finally, the long run equilibrium, y^*, can be calculated as

$$\begin{aligned} y^* &= \mathbf{h}'_{LR}\boldsymbol{\xi} \\ &= \left(\frac{\alpha_1}{1-\rho_1}\right)\left(\frac{1-p_{22}}{2-p_{11}-p_{22}}\right) + \left(\frac{\alpha_2}{1-\rho_2}\right)\left(\frac{1-p_{11}}{2-p_{11}-p_{22}}\right). \end{aligned}$$