

A Two-Factor, Stochastic Programming Model of Danish Mortgage-Backed Securities

Søren S. Nielsen* and Rolf Poulsen†

June 11, 2002

Abstract

Danish mortgage loans have several features that make them interesting: Short-term revolving adjustable-rate mortgages are available, as well as fixed-rate, 10-, 20- or 30-year annuities that contain embedded options (call and delivery options). The decisions faced by a mortgagor are therefore non-trivial, both in terms of deciding on an initial mortgage, and in terms of managing (rebalancing) it optimally.

We propose a two-factor, arbitrage-free interest-rate model, calibrated to observable security prices, and implement on top of it a multi-stage, stochastic optimization program with the purpose of optimally composing and managing a typical mortgage loan. We model accurately both fixed and proportional transaction costs as well as tax effects. Risk attitudes are addressed through utility functions and through worst-case (min-max) optimization. The model is solved in up to 9 stages, having 19,683 scenarios. Numerical results, which were obtained using standard soft- and hardware, indicate that the primary determinant in choosing between adjustable-rate and fixed-rate loans is the short-long interest rate differential (i.e., term structure steepness), but volatility also matters. Refinancing activity is influenced by volatility and, of course, transaction costs.

Key words: Term structure modeling, stochastic programming, mortgage backed securities. **JEL Classification:** C44, E43, G11, G21.

*Informatics and Mathematical Modeling, Technical University of Denmark, Bldg. 321, DK-2800 Lyngby, Denmark. sn@imm.dtu.dk, www.imm.dtu.dk/~sn

†Department of Statistics and Operations Research, University of Copenhagen, Universitetsparken 5, DK-2100 Copenhagen, Denmark. rolf@math.ku.dk, www.math.ku.dk/~rolf

1 Introduction

In Denmark there is considerable public interest in the market for mortgage backed securities. Spurred on by the product innovation and advice of financial institutions, many home-owners manage their mortgage debt very actively. The individual mortgagor faces a number of non-trivial decisions. He has to decide whether to use adjustable rate loans where the debt is refinanced on a yearly basis (or every 3 or 5 years) or the more traditional fixed rate 20- or 30-year annuities, or some combination hereof. The annuities have embedded options: There is a call-feature because the mortgagor can repay the remaining principal at any time, and there is a delivery option because he can buy back his debt at market value (a feature not present in mortgages in the U.S.). We use a stochastic, two-factor term structure model (calibrated to market prices and volatilities of non-callable bonds and to prices of callable bonds) and formulate the mortgagor's problem of optimal debt management as a multistage stochastic programming model that is solved with standard software. The formulation as a stochastic programming model makes it possible to incorporate both the inherently path-dependent aspects of mortgage management, and the optimal policy aspects. In addition, it makes it easy to model investor specific risk attitudes and tax rates, as well as both fixed and proportional transaction costs.

Results indicate that the model behaves reasonably and accurately replicates actually observed mortgagor behavior. In particular, the model actively utilizes both "down-" and "up-" mortgage rebalancings in the fixed-rate loans, i.e., the embedded call and delivery options, and it uses both adjustable-rate and fixed-rate loans realistically. We find that the choice between using long (fixed-rate) or short (adjustable-rate) bonds depends largely on the spread between long and short rates, but also on volatilities. These effects are hard to capture with the one-factor models often employed.

The problems of mortgage prepayment behavior has been studied, e.g., by Richard and Roll (1989), and by Kang and Zenios (1992), and of mortgage portfolio management (from the investor's side) by e.g., Zenios (1995) and many others, but our studies are the first that address the mortgagor's side of the problem. The topics presented in this paper were, with a preliminary implementation, also addressed in Nielsen and Poulsen (2002). However, we here present significant further developments: The interest rate model is further refined and implemented, its estimation to market data and the calculation of callable bond market prices are revised, all numerical results are new, and we further investigate effects of fixed transaction costs as well as the min-max risk-averse objective function. In addition, the GAMS implementation has been streamlined, leading to significantly faster solution times.

The outline of the rest of the paper is as follows. In Section 2 we give a more detailed description of the mortgage market, with particular emphasis on the products available to the individual mortgagor and the choices/trade-offs he faces. Section 3 introduces a stochastic interest rate model. We use a Gaussian two-factor model for zero coupon bond yields formulated in the Heath-Jarrow-Morton framework. In this section we also propose a simple, but very operational, regression approach for determining prices of callable mortgage backed bonds from zero coupon bond yields. In Section 4 the problem that the mortgagor from Section 2 faces and the dynamic model for interest rates and bond prices from Section 3 are combined in a multistage stochastic programming formulation. Section 5 presents extensive experiments with the model and numerical results. Finally, Section 6 concludes the paper and outlines topics for future research.

2 The Danish Mortgage Market and the Individual Mortgagor

The interest rate policy of the Danish Central Bank usually mimics that of the European Central Bank's short rates, and the yields on government bonds also closely follow those in "Euroland", with the addition of a spread of 25 to 50 bp. On several occasions the Danes have voted "No" to joining the single European currency, the Euro, but still the exchange rate is quite stable (between DKK 7.40 and 7.50 will buy you 1 Euro.)

Though small in absolute terms, the Danish mortgage market¹ has some interesting features. Mortgages have historically been financed by 20- or 30-year fixed-rate bonds that were issued through intermediaries (first only dedicated mortgage companies, since 1970 also banks) to a quite liquid market. The mortgagor can repay the remaining principal at any time,² in other words the bonds are *callable*. When interest rates drop, the mortgagor can issue new debt at the lower rates, typically in the form of a new fixed-rate (callable) bond with lower coupon rate, using the proceeds from this to pay off the old mortgage, which is "called" at par. Fixed-rate bond prices actually do increase to above 100 without being called. One such reason is that it is not possible or worthwhile for the individual mortgagor to be a highly efficient investor ("optimal behavior" is ill-defined), another is transaction

¹Figures from late 2000: Measured by market value the size of the Danish bond market was DKK $2 \cdot 10^{12}$ (and a tiny bit more by face value); roughly 50% are mortgages, 30% government bonds, 8 % inflation-linked bonds, and the rest are more or less exotic bonds that only constitute 3% of the turnover.

²Actually such extraordinary prepayments can only take place on coupon payment dates (of which there are typically 4 per year) and two months' prior notification has to be given.

costs. But there is also a typical options aspect: Even when an option is in-the-money it may be optimal to hold it, because it may become even better yet.

Conversely, if interest rates increase, the mortgagor can buy back the now relatively cheap mortgages in the market, again funding this by issuing a new mortgage, now with a higher coupon rate. This will make the installments larger, but reduce the principal on the mortgage, and the mortgagor has positioned himself better if interest rates subsequently drop. This was a strategy widely recommended by financial institutions from mid-1999 to mid-2000, because there was a perception that “rates are high” (since they had recently gone up) and “they will come down” (in particular because the common belief was that Denmark would soon join the Euro). An effective refinancing strategy therefore involves refinancing both when rates decrease and increase, in either case issuing bonds as close to par as possible (mortgagors are by law only allowed to issue bonds priced at or below par).

A key question is, by how much should rates change before refinancing is optimal? Danish banks now offer so-called mortgage-watch programs, where they alert their customers to opportunities. Of course the banks make a sizeable profit (there is a fixed cost to the bank of typically DKK 1,500-2,000 and a variable cost of 0.15% of the price of the new issue, in addition to taxes) when a customer refinances, so it is in their interest to induce refinancing.

To the purchaser of these bonds, the credit risk is very low because the lending institution pools and insures the issue. Only some systemic collapse, affecting a major part of the economy, would present a risk that could hardly be avoided anyway. But there are other issues that make the purchaser side interesting, in particular modeling the prepayment behaviour of the mortgagors in the bond issue. Note the asymmetry; the mortgagor takes into account very specifically his own characteristics, whereas the purchaser buys the “average” mortgagor. Ideally, one should consider both sides in detail and then arrive at (model) market prices by some equilibrium argument. However, our model only addresses the side of the issuer in detail, the overall market behaviour is exogenous, but modeled as a stochastic system in a statistically plausible way.

In the mid-90'ties, adjustable-rate mortgages (ARMs) in the form of revolving, short-term loans were introduced, following a legislative change. These loans were first offered by Realkredit Danmark under the trademark FlexLån, then (reluctantly) by most other intermediaries. The simplest one is the F1, whereby the complete outstanding principal is refinanced every year by January 1st at the prevailing 1-year rate; similarly there are F2, F3, up to F10 loans. Another option is the P-loans, for instance the P25,0, where 25% of the debt is refinanced every year at the 4-year rate, many other fractions up to 50% exist. These ARMs all share two characteristics: They depend on the short end of the yield curve,

and they are very vulnerable to increases in short rates. In a "normal", positively sloped term structure they have some appeal compared to the long-term loans, and the market for these loans, after some hesitation, is now very large. Figure 1 shows the Danish short and long rates during the period 1997 to early 2001. Although ARMs had an obvious appeal during the early years, there was a significant narrowing (flattening) in 2000 (due to the uncertainty of the Danish referendum Sep. 28 whether to join the Euro; we did not), making the situation far from obvious for the individual mortgagor.

3 The Two-Factor Interest Model and Prices of Mortgage Backed Bonds

Our interest-rate model has two components: A classic term structure model for the stochastic movement of zero-coupon bond prices (ZCB prices; $P(t, T)$) and a model that links the ZC-yield curve (simply referred to as *the term structure*) to prices of mortgage backed callable bonds, or securities (simply referred to by MBSs).

3.1 A Two-Factor Term Structure Model

It is convenient to work with a model that can take today's observed term structure directly as input. This is most easily done when we work in the Heath-Jarrow-Morton-framework (see Heath, Jarrow and Morton (1992)). Models of this type are usually formulated in terms of instantaneous forward rates (although a specification of the dynamics of all ZCB prices or yields is theoretically equivalent, see Björk (1998, Chapter 15) for a lucid exposition of different approaches to continuous-time interest rate modelling),

$$f(t, T) = -\frac{\partial \ln P(t, T)}{\partial T},$$

i.e. $f(t, T)$ is the forward rate seen at time t for an extremely short term loan at time T ,³ and $P(t, T)$ is the price at time t of a ZCB that matures at T .

We use the Gaussian volatility specification

$$df(t, T) = \alpha(t, T)dt + \sigma_1 dW_1 + \sigma_2 e^{-\lambda(T-t)/2} dW_2.$$

where W_1, W_2 are (uncorrelated) Brownian motions under some probability measure \mathbb{P} (the actual/objective/statistical/physical measure), and σ_1, σ_2 and λ are positive constants. In

³When interest rates are quoted on a continuously compounded, yearly basis, then the simple forward rate seen from time t for a loan between times T and $T + h$ is $\ln(P(t, T)/P(t, T + h))/h$. The $f(t, T)$ rate is the " $h \rightarrow 0$ " limit of these simple forward rates.



Figure 1: Danish ZC-yields between September 1997 and January 2001.

this model, changes in the term structure are caused by two factors. The first factor, σ_1 , uniformly shifts rates of all maturities, while the second factor, σ_2 , affects short rates more than long rates. The second factor can also be thought of as a spread between long and short rates. The very short rate volatility is $\sigma_1 + \sigma_2$, the long rate volatility is σ_1 , and λ controls how rapidly the volatility flattens out. If $\lambda = 0$, then rates of all maturities are perfectly correlated, positive λ -values mean that short rates are less-than-perfectly correlated with long rates, and a “very high” λ -value means that the short rate “has a life of its own” (to some extent at least; there is a (long, short)-correlation lower bound of $\sigma_1/(\sigma_1 + \sigma_2)$).

This two-factor model allows for a wider range of term structure deformations than the (effectively) parallel shifts of a one-factor model. Figure 1 depicts the Danish 1Y-, 10Y-, and 30Y-ZC-yields⁴ for the period September 1997 to January 2001. Clearly, these would be poorly fitted by a model that allowed only parallel shifts. The form of the volatility is specified such that it is easy to incorporate the empirically robust fact that long rates are

⁴On any given day, the ZC-yield curve was estimated from prices of Danish government bonds by the non-parametric smoothing technique described in Tanggaard (1997). The data were kindly supplied by ScanRate.

less volatile than short rates; a fact caused by mean-reversion.

For the model to be arbitrage-free, a risk-neutral probability measure Q (or: martingale measure) must exist. This implies the existence of a 2-dimensional stochastic process ϕ of risk-premia such that the \mathbb{P} -drifts of the forward rates must obey the so-called HJM-drift condition:

$$\alpha(t, T) = \sigma_f^\top(t, T) \int_t^T \sigma_f(t, s) ds - \sigma_f^\top(t, T) \phi(t).$$

Assuming constant risk-premia $\phi^\top(t) = (\phi_1, \phi_2)$ we have for $s < t < T$ that

$$\begin{aligned} f(t, T) &= f(s, T) + \sigma_1^2(t-s)(T-(t+s)/2) \\ &\quad - \frac{2\sigma_2^2}{\lambda^2} \left[(e^{-\lambda(T-t)} - e^{-\lambda(T-s)}) - 2(e^{-\lambda(T-t)/2} - e^{-\lambda(T-s)/2}) \right] \\ &\quad + \phi_1 \sigma_1(t-s) + \phi_2 \frac{2\sigma_2}{\lambda} (e^{-\lambda(T-t)/2} - e^{-\lambda(T-s)/2}) \\ &\quad + \sigma_1(W_1(t) - W_1(s)) - \sigma_2 e^{-\lambda T/2} \int_s^t e^{\lambda u/2} dW_2(u). \end{aligned} \quad (1)$$

Note that in fact we have a Markovian structure, in that the entire forward rate curve can be represented in terms of a deterministic function and the stochastic variable

$$\left(\begin{array}{c} \sigma_1 \times \sqrt{t-s} \times n_1 \\ \sigma_2 \times e^{-\lambda T/2} \times \sqrt{\frac{e^{\lambda t} - e^{\lambda s}}{\lambda}} \times n_2 \end{array} \right), \quad (2)$$

where n_1 and n_2 are independent and standard normal. For practical purposes we recast (2) in two ways; in terms of ZCB prices, and in terms of dynamics of ZC-yields. Since $P(t, T) = \exp(-\int_t^T f(t, u) du)$, we find by integrating (2) wrt. T , changing signs and exponentiating that (not quite as elegant as $E = mc^2$):

$$\begin{aligned} P(t, T) &= \frac{P(s, T)}{P(s, t)}. \quad (3) \\ &\exp \left(-\sigma_1^2(t-s)(T-t)(T-s)/2 - \frac{2\sigma_2^2}{\lambda^3} (e^{-\lambda(T-t)} - 1 - e^{-\lambda(T-s)} + e^{-\lambda(t-s)}) \right. \\ &\quad + \frac{8\sigma_2}{\lambda^3} (e^{-\lambda(T-t)/2} - 1 - e^{-\lambda(T-s)/2} + e^{-\lambda(t-s)/2}) + \phi_1 \sigma_1(t-s)(T-t) \\ &\quad - \phi_2 \frac{2\sigma_2}{\lambda^2} (e^{-\lambda(T-t)/2} - 1 - e^{-\lambda(T-s)/2} + e^{-\lambda(t-s)/2}) - n_1 \cdot \sigma_1(T-t) \sqrt{t-s} \\ &\quad \left. - n_2 \cdot \frac{2\sigma_2}{\lambda \sqrt{\lambda}} \sqrt{e^{-\lambda(T-t)} + 1 - 2e^{-\lambda(T-t)/2} - e^{-\lambda(T-s)} - e^{-\lambda(t-s)} + 2e^{-\lambda((T+t)/2-s)}} \right). \end{aligned}$$

The ZC-yield with time τ to maturity is defined by $y(t, \tau) = -\ln P(t, t+\tau)/\tau$, so we find the ZC-yield dynamics to be

$$dy(t, \tau) = \dots dt + \sigma_1 dW_1 + \frac{2\sigma_2}{\tau \lambda} (1 - e^{-\lambda \tau/2}) dW_2. \quad (4)$$

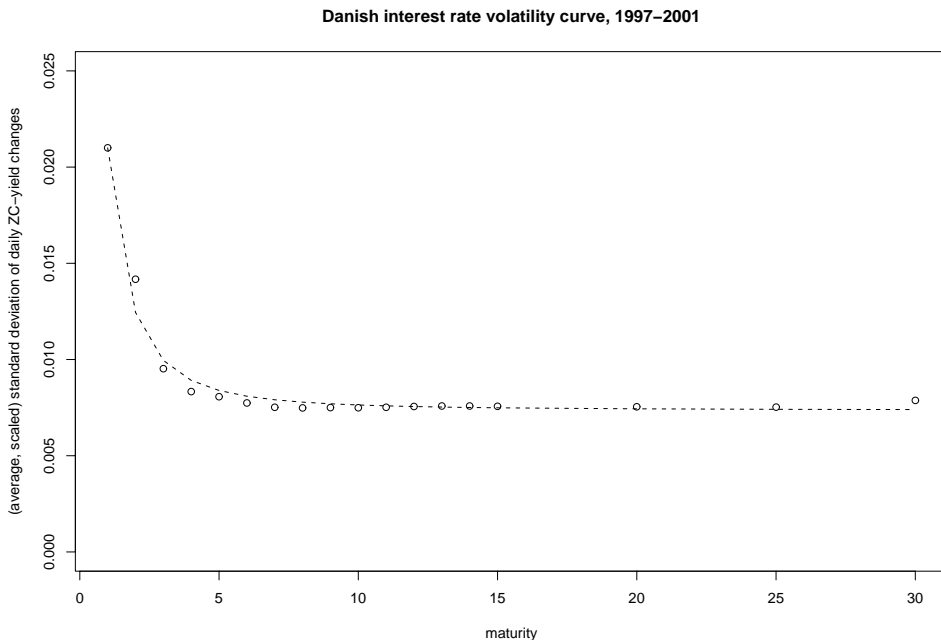


Figure 2: Theoretical (dotted line) and observed (o) standard deviations of Danish ZC-yields, when the Gaussian two-factor model is fitted by least squares til daily ZC-yield changes from September 1997 to January 2001.

3.1.1 Estimation of Parameters

We want to estimate the parameters σ_1 , σ_2 , λ , ϕ_1 and ϕ_2 . The former 3 are (primarily) volatility related parameters, so they can reliably be estimated from fairly-frequently-sampled historical data (without the need for an immensely long time period from first to last observation.) Suppose we have time series observations (at t_i 's that are $\Delta t = 1$ day or 1 week apart) on ZC-yields⁵ for maturities τ_i (1Y, 2Y, \dots , 30Y; ZC-yields of maturities shorter 1Y have considerable idiosyncratic noise and are excluded). From (4) we see that

$$\text{std. dev.} \left(\frac{\Delta y(t_i, \tau_j)}{\sqrt{\Delta t}} \right) \approx \sqrt{\sigma_1^2 + \frac{4\sigma_2^2}{\tau_j^2 \lambda^2} (1 - e^{-\lambda \tau_j / 2})^2},$$

⁵Term structure information is typically supplied in the form of a number of points on the ZC-yield curve, not in terms of instantaneous forward rates. It is advisable to perform analytical integration (as we did to arrive at (4)) rather than numerical differentiation (to get “observed” instantaneous forward rates.)

Estimates of ZCB curve parameters in equation (4)

$\hat{\sigma}_1$	$\hat{\sigma}_2$	$\hat{\lambda}$
0.00736	0.0782	7.80
$\hat{\phi}_1$	$\hat{\phi}_2$	
0.00325	-0.0498	

Coefficients in the MBS price pricing function

a	b	c
0.0219	1.85	81.8

Figure 3: Parameter estimates. The MBS-curve becomes flat for present values above 124.82. The maximal model price of a MBS is 101.6. Note that this is close to “100 + proportional cost”.

and when Δt is small the dt -term in (4) is small compared to the dW -terms, so parameters can be estimated solely from observed standard deviations of yield changes. We simply fit (using a least squares criterion) these theoretical standard deviations to the observed standard deviations of ZC-yields. Using daily ZC-yield changes in the 1997-2001 period we find the estimates shown in to upper panel of Table 3. The observed standard deviation and the fitted curve are shown in Figure 2. We see a fairly good fit; in factor analysis terms the model explains about 90% of the variation in the data. The risk premia (the ϕ 's) are estimated at the values that produce the best fit on average (from September 1997 and onward and with volatility parameters fixed) to the slope of the term structure at the 1Y and 25Y year points.

3.1.2 Discretization of the Model

From (2) it is easy to see that we essentially only need a discretization of two independent, standard normal variables. We want to do that in trinomial fashion, which schematically looks like this:

$$\begin{array}{r}
\begin{array}{c} 1/3 \\ 1/3 \\ 1/3 \end{array} \left(\begin{array}{c} n_1^u \\ n_2^u \end{array} \right) \xrightarrow{\text{by (4) with } t=s+\Delta \text{ and } T=s+\Delta+\tau} (P_u(s+\Delta, s+\Delta+\tau))_{\tau \geq 0} \\
(P(s, s+\tau))_{\tau \geq 0} \left(\begin{array}{c} n_1^m \\ n_2^m \end{array} \right) \xrightarrow{\text{by (4)}} (P_m(s+\Delta, s+\Delta+\tau))_{\tau \geq 0} \\
\left(\begin{array}{c} n_1^d \\ n_2^d \end{array} \right) \xrightarrow{\text{by (4)}} (P_d(s+\Delta, s+\Delta+\tau))_{\tau \geq 0}
\end{array}$$

It is easily checked that if we use the values

$$\left(\begin{array}{c} 1/\sqrt{2} \\ \sqrt{3/2} \end{array} \right), \left(\begin{array}{c} -2/\sqrt{2} \\ 0 \end{array} \right), \text{ and } \left(\begin{array}{c} 1/\sqrt{2} \\ -\sqrt{3/2} \end{array} \right)$$

for the (n_1, n_2) “up”, “middle” and “down” states, then the first two moments fit uncorrelated $N(0, 1)$ -variables.

The lattice is recombining but the actual implementation is not – because the decisions are path-dependent. The resulting tree corresponds to a N -stage, stochastic optimization program with decision points at $(t_0, t_1, \dots, t_{N-1})$ and a final (trivial) horizon decision at t_N . The tree has a total of $\sum_{n=0}^N 3^n$ nodes in its non-recombining version.

3.2 Market Prices of Callable Mortgage Backed Bonds

Of course market participants know full well that MBSs are callable. Hence the bonds do not trade at prices that are the payments of the non-callable annuity discounted by ZCB rates, but at lower prices to reflect the value of the embedded option

$$\text{MBS} = \text{non-callable bond} - \text{call-option}$$

But which prices *do* the MBSs trade at? Or in other words, how do we get from a dynamic model of the ZC-yields to a model for MBS-prices? There are two “schools” in the literature when it comes to this. The first tries to price the embedded call-option (which is of American or Bermudan type, to complicate matters further) by standard “no-arbitrage techniques” in stochastic interest rates. This approach can be augmented (as in the classic paper Stanton (1995)) by introducing borrower heterogeneity (wrt. costs) and “bounded rationality” (markets are “only checked now and again”). While theoretically well-founded, the method does not produce overly satisfactory results when confronted with market data. The other “school” focuses on “empirical prepayment”. Here, a statistical model of what

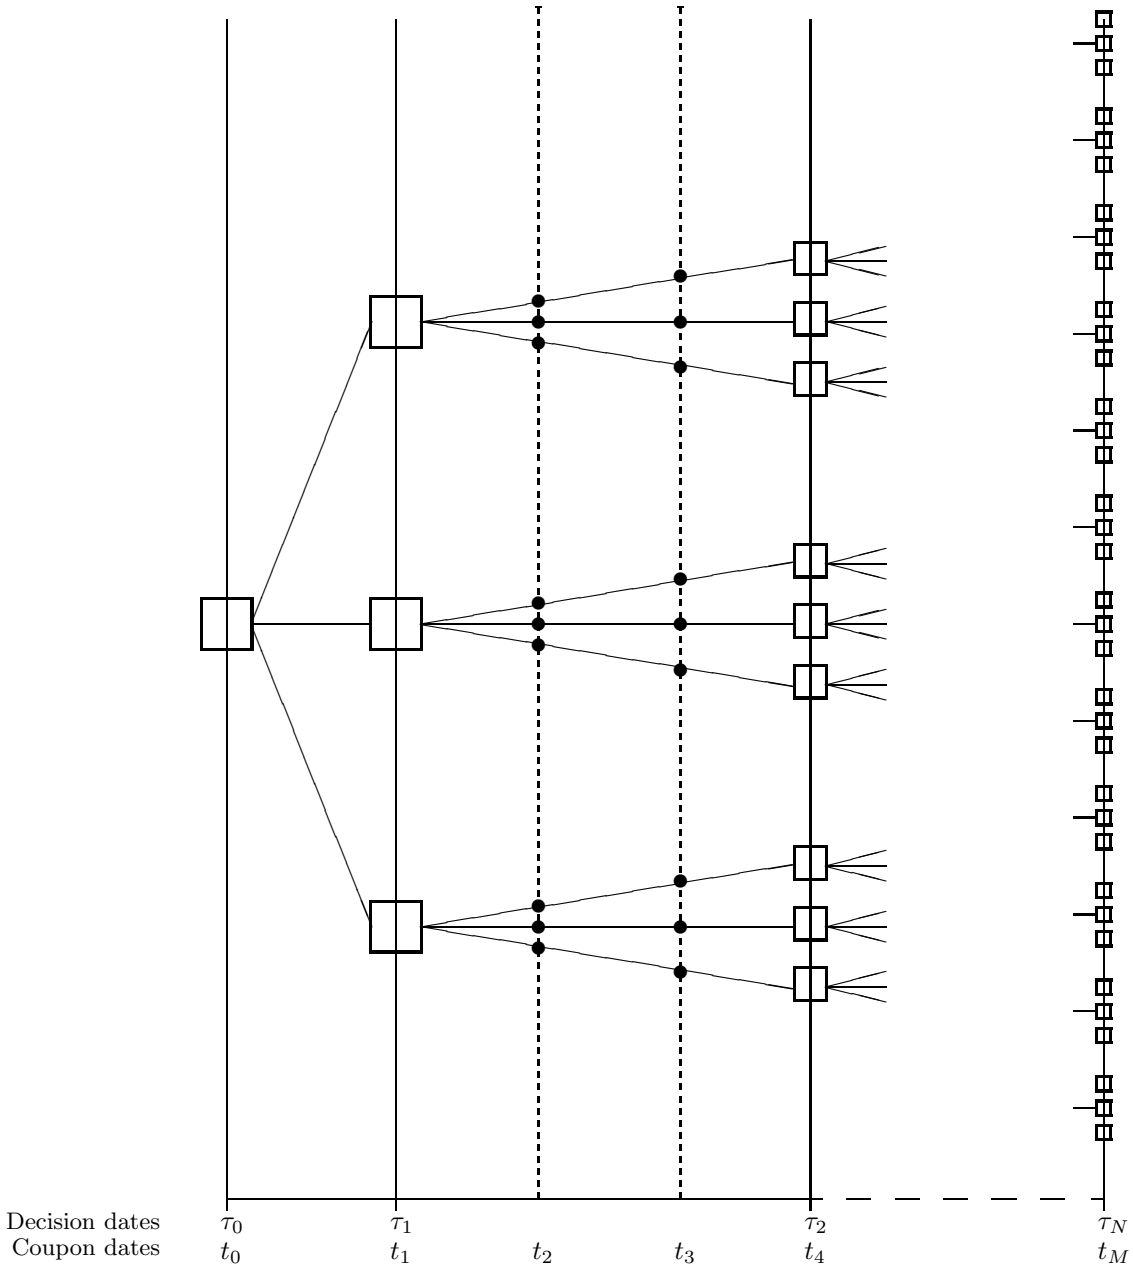


Figure 4: The decision tree. The square boxes (\square) where we can actively make a decision (in other words the k 's that we sum over in (5)). At each of these nodes we have a complete term structure and a universe of MBSs. The bullet points (\bullet) are other dates at which cash-flows occur.

PRICING OF CALLABLE BONDS




Figure 5: Pricing of callable MBSs: Model prices and observed prices. The parameters are estimated by fitting the fully drawn curve (of the form $a(x - c)^b$) to observed 30Y prices (o's). The dotted curve and the x's are model prices and observed prices of 20Y bonds. Note that these model prices are “out-of-sample” predictions.

causes people to prepay is set up (think of this as a regression of observed prepayments on “whatever could be perceived as relevant and then some”) and this model is then linked to market prices of MBSs (typically through so-called option adjusted spreads). This method is quite data-intensive, but still does not produce overwhelmingly accurate price predictions.

This motivates our use of a simple model to produce MBS prices in the stochastic programming set-up. The model is easy to estimate and work with, performs reasonably well out-of-sample and overall “doesn't do silly things”. The story goes as follow:

Let t denote a given day and consider an n -year MBS. We find the vector of cashflows of the corresponding non-callable bond (say $cf(t + \tau_i)$ at date $t + \tau_i$) and calculate its price

$$PV^{nY}(t, cf) = \sum_i cf(t + \tau_i) \exp^{-y(t, \tau_i)\tau_i} .$$

This price must be a very important influence on the price a MBS. So important, in fact, that we are tempted to look for functions, f_n (one function for each maturity), such that the price of the callable bond is

$$P_{MBS}^{nY} = f_n(PV^{nY}).$$

And what should f_n look like? If PV is far below 100, the embedded option is far out of the money, and the two bonds cost about the same. On the other hand, nobody would want to pay “significantly” more than 100 for the MBS; whoever sells the bond can buy it back immediately for 100. So we should have an upper limit for the price of a MBS; presumably somewhere around “100 + cost incurred upon prepaying”. Further, the shorter the time to maturity of the underlying bond, the closer the price of the MBS is to $\min(PV, 100 + \text{costs})$. The f_n -functions given below possess the desired properties (look at Figure 5 to get an impression of their graphs) :

$$f_{30}(x) = \begin{cases} x & \text{for } x \leq c \\ x - a(x - c)^b & \text{for } c \leq x \leq c + (ab)^{\frac{1}{1-b}} \\ c + (ab)^{\frac{1}{1-b}} - a(ab)^{\frac{b}{1-b}} & \text{for } x \geq (ab)^{\frac{1}{1-b}} \end{cases},$$

$$f_0(x) = \min\left(x, c + (ab)^{\frac{1}{1-b}} - a(ab)^{\frac{b}{1-b}}\right), \text{ and}$$

$$f_m(x) = \frac{m}{30}f_{30}(x) + \frac{30-m}{30}f_0(x).$$

We estimate the parameters a , b and c from market data: Given observed term structures at specific dates (between 1997 and 2001 in this case) and a guess of (a, b, c) , model-predictions of MBS-prices can be calculated. These are compared to market prices, and we estimate (a, b, c) as the values that give the best fit (with a least squares criterion). The results of this analysis is given in Figure 5, where the “model” and “market” prices are shown, and in Table 3 where the estimated coefficients are given.

A few remarks:

- The lower and upper cut-off points are implicitly estimated. We get an upper limit of 101.6, i.e. an “implied cost estimate” of 1.6%, which is very reasonable.
- By law, the borrower is only allowed to issue MBSs when the price is below 100, so he doesn’t particularly care, how the MBS price behaves once it gets above 100.
- We only use 30Y bonds to estimate parameters. The f_0 -function (giving the price of a MBS that is just about to mature) is estimated as $\min(x, 100 + \text{implied cost estimate})$, and the other f_n ’s are determined by linear interpolation between f_0 and f_{30} . This

means that by comparing “model” and “market” prices of 20Y MBSs (these are also fairly liquid, although not the primary choice of private mortgagors), we get a simple out-of-sample test of the procedure. The result of this is also shown in Figure 5. The model tends to underprice the 20Y bonds, but for the purposes of this paper, the accuracy is certainly acceptable.

4 The Multistage Stochastic Programming Model

We now present a multistage, stochastic optimization model which implements the mortgagor’s problems of initially establishing an optimal portfolio (of bonds), and also manages this portfolio optimally. As is typical of investment problems, the “portfolio” in question will consist of just a single instrument in the risk-neutral case, but the model allows for risk-aversion by a suitable choice of objective. At this point we take for given the structure of the scenario tree (i.e., a N -stage, trinomial tree), and that there is at each node in the tree a term structure, and also a finite set of loan types (bonds) available. Bond (cash-flow) present values as well as their market prices are determined for each node using the term structure. The stochastic optimization model itself is completely detached from the interest rate/term structure generating model used. This model allows, for instance, a jump diffusion or a mean-reversion as the underlying process structure.

The objective of the model is to maximize some measure of expected payments, suitably discounted over time, and is discussed in Section 4.2. Below we discuss the constraints of the model (balance equations), whose primary purpose is to keep track of remaining principal versus payments, under consideration of possible rebalancing, and across nodes in the stochastic tree.

We recall the relationship between time periods, decision points, tree nodes and their indices: The stochastic tree covers N time periods, or stages, with decisions at times τ_n , $n = 0, 1, 2, N - 1$ (the “decision” at the horizon, τ_N , is trivial). Typically, $\tau_N = 30$ years; $\tau_0 = 0$. The set of tree nodes, of which there are 3^n at level n (corresponding to time τ_n), is indexed by k , $k = 1, \dots, K = \sum_{n=0}^N 3^n$.

4.1 Principal Balance Equations

The balance equations manage outstanding principal amounts of each (potential) bond (loan type), between pairs of tree nodes, (a, k) , where node a is the *predecessor* of node k . The only node without a predecessor is the root node. At each node k there is a finite universe

of traded bonds, indexed by i , and given the term structure at the node we can find market prices of these, I_i^k .

Let:

DF_t^k be the discount factor of node k 's term structure for a loan maturing at time $\tau_k + t$, i.e., t years after the node time;

PV_i^k be the *present value* of bond i at node k , calculated using DF_t^k on the bond's cash flow. This would be the fair (ex-coupon) price of the bond were it not callable;

I_i^k be the market price of bond i at node k . The bond can be *sold* at this price (if it is at or below par), and can be *called* at $\min(I_i^k, 100)$. It is calculated as explained in Section 3.2;

c_i be bond i 's coupon rate (which is 0 for ARMs).

We need the following **decision variables**:

x_i^k is the outstanding principal amount of bond i at node k ; for a node-ancestor pair (a, k) , the balance equations given below serve to link together x_i^a and x_i^k ;

s_i^k is the amount of bond i sold at node k (to raise funds of DKK 1,000,000, bonds can only be sold if $I_i^k \leq 100$)

pp_i^k is the scheduled payment of principal at node k , given by the annuity formulas in Section 4.1.2;

ip_i^k is the scheduled payment of interest at node k , also given in Section 4.1.2;

q_i^k is the extraordinary pre-payment of principal when bonds priced above par are called at par (using the embedded call option);

p_i^k is the amount of bonds priced below par, purchased at the market price I_i^k and delivered towards the outstanding loan (using the embedded delivery option).

In the absence of refinancing or prepayments, principal is paid down by regularly scheduled payments, pp_i^k . However, let us clarify the difference between pp_i^k and q_i^k : Whenever interest rates increase sufficiently, it is optimal to "convert up", i.e., purchasing back one's own bonds in the market, thus reducing the outstanding debt (this happens when the bond is cheap, priced below par), which is modeled by p_i^k . This purchase is financed by issuing new, higher-coupon bonds. Similarly, when interest rates decrease sufficiently, bonds whose

prices are now above par can be called at par ("converting down"), which is modelled by q_i^k .

We are now in a position to state our primary **Balance Constraint**: For bond i the principal amounts of debt at nodes (a, k) satisfy

$$x_i^a - pp_i^k - q_i^k + s_i^k = x_i^k + p_i^k$$

There is also a cash constraint enforcing that no cash is added to or removed from the system between time 0 and 30 (except scheduled payments). An additional constraint in effect for F -year ARMs, i.e., those that are refinanced every F years (1, 2, 3 or 5), is that we are not allowed to prepay them during those F years; in real life, these loans can be bought out only at a relatively high cost, or if the underlying asset (house) is sold.

4.1.1 Initial and terminal conditions

We must raise a nominal DKK 1,000,000 (or any other, predetermined amount) at the outset; everything must be paid back after at most 30 years, the maximum horizon allowed for personal mortgages. This condition will be satisfied because the annuitized principal repayment schedule (Section 4.1.2) must be adhered to.

4.1.2 Annuity Calculations

The total mortgage installments (principal and interest) are given by the standard *annuity formula*:

$$A_i = \frac{c_i \cdot (1 + c_i)^n}{(1 + c_i)^n - 1}$$

so that this amount is paid in every period (of n periods) per unit face value, given a coupon rate of c_i on the loan (for $c_i = 0$ we have $A_i = 1/n$). Of this, repayment of principal in the j th period, $j = 1, \dots, n$, is $(1 + c_i)^{-(n-j+1)}$, and the rest is interest.

Define the per unit principal and interest payments during the period $[\tau_a, \dots, \tau_k - 1]$ as:

$$PP_i^a = A_i \sum_{t=\tau_a}^{\tau_k-1} (1 + c_i)^{t-m_i}$$

$$IP_i^a = A_i \sum_{t=\tau_a}^{\tau_k-1} (1 - (1 + c_i)^{t-m_i});$$

where m_i is the maturity year of bond i . Principal and interest payments accrued during this period are paid at time τ_k , which is enforced by the constraints

$$pp_i^k = PP_i^a \cdot x_i^a,$$

$$ip_i^k = IP_i^a \cdot x_i^a.$$

Note that by keeping track of principal and interest payments separately, the model allows proper *tax treatment* of interest payments (in Denmark, mortgage interest payments are tax-deductible at the current rate of 32.13% whereas principal payments are not. Other interest payments are fully deductible at the marginal tax rate. The Danish top marginal tax rate is currently 61%, we repeat, 61%).

4.1.3 Total Payment Calculation

The total payment (cash out) at node k is then:

$$T^k = \sum_i (pp_i^k + q_i^k + (1 - \xi) \cdot ip_i^k - (1 - \beta) \cdot I_i^k \cdot s_i^k + I_i^k \cdot p_i^k)$$

where ξ is the tax rate and β is the (proportional) transaction cost on sales (and hence on refinancing). Proportional transaction costs are incorporated by setting the parameter β to the appropriate percentage cost of issuing (selling) bonds. Fixed costs can be modeled using binary variables; we do this in Section 5.7.

4.2 Objective Functions

We propose as the model's basic objective to minimize *expected present value* of life-time expenses:

$$\text{Minimize } z = \sum_{k=0}^K \pi^k \cdot \delta_t \cdot T^k \quad (5)$$

where:

$k = 0, \dots, K$ are tree nodes indices,

δ_t : Discount factor for payments at time t (e.g., the initial yield curve plus spread),

π^k : The probability of state (node) k ; $\pi^k = 3^{-n}$ if node k is at level n ,

T^k : Total payment at node k .

This objective function corresponds to an unlikely *risk-neutral* investor. In real life, by contrast, home owners tend to be very risk-averse, given that a mortgage is typically the largest single investment one will ever make. In practice, this risk-adversion manifests itself by a strong reluctance to fully finance a mortgage by ARMs and could thus be modeled by

excluding such loans from the loan universe. A more rigorous approach is to use a *utility function*, \mathcal{U} , concave in its second parameter, and then

$$\text{Maximize } z_{\mathcal{U}} = \sum_{k=0}^K \pi^k \cdot \mathcal{U}(\delta_t, T^k); \quad (6)$$

where \mathcal{U} also appropriately accounts for intertemporal utility comparisons. In Section 5.5, for instance, we use the objective

$$\text{Maximize } z_{\log} = \sum_{k=0}^K \pi^k \cdot \log(\delta_t \cdot (B_{\tau_k} - T^k)), \quad (7)$$

where B_{τ_k} is a budget assumed to be available for payments at time τ_k , to maximize the expected present value of budget surplus.

Finally, we implement a worst-case, or min-max objective which seeks to *minimize the maximum* present value of life-time expenses across all scenarios, where a scenario is a path from the tree root to a leaf node. This objective corresponds to extreme risk-aversion.

5 Numerical Experiments

In this section, extensive numerical experiments are reported. The purpose of the experiments presented here is to validate the model, i.e., verify that it behaves reasonably, and also to demonstrate its feasibility with respect to solution times for large, non-trivial instances. With starting point in a "base case" model, we investigate the effects on the initial, optimal mortgages and on refinancing behavior of varying the number of stages, of incorporating risk-aversion, of changing the term structure steepness and volatilities, and of fixed and proportional transaction costs. The model is implemented in GAMS, using CPLEX as the linear programming solver.

5.1 Base Case Model

The experiments presented in this section use a common "base case" setup. We chose as our base case model parameters, $\sigma_1, \sigma_2, \lambda, \phi_1, \phi_2$ the volatilities estimated in Table 3. These parameters result in a relative volatility of the 1-year rate of 18.5%, and of the 30-year rate of 14.5%, with a correlation of 0.96. The initial yield curve used is from January 20, 2001, with a 1-year rate of 5.14%, a 30-year rate 5.83% and "flat" in between, though with a slight dip to 4.80% around year 3. Market prices of the fixed-rate, callable bonds are estimated as explained in Section 3.2.

Stages	Decision points, τ (years)	Tree Nodes	Scenarios	Solution Times	
				PC	Linux
3	(0, 1, 5, 30)	40	27	1	0.5
4	(0, 1, 5, 10, 30)	121	81	3	1
5	(0, 1, 2, 5, 10, 30)	364	243	7	4
6	(0, 1, 2, 5, 10, 20, 30)	1093	729	22	11
7	(0, 1, 2, 3, 5, 10, 20, 30)	3280	2187	174	57
8	(0, 1, 2, 3, 5, 10, 15, 20, 30)	9841	6561	1881	461
9	(0, 1, 2, 3, 5, 7, 10, 15, 20, 30)	29524	19683	21309	N/A

Table 1: The multistage, stochastic models: Number of stages, decision points, nodes in the resulting trinomial tree and number of scenarios (paths from root to leaves), and solution times in seconds on a Windows XP PC (800MHz Pentium, 128Mbytes RAM) and on a Linux machine (1GHz Pentium, 128 Mbytes RAM), using GAMS Rev 121 and CPLEX 7.0. The Linux machine had insufficient virtual memory to solve the largest instance.

The model is set up to accommodate a number from 3 to 9 stages, with the values of τ shown in Table 1, and for the base case we select 5 stages; this captures the compromise between model accuracy and speed of solution.

The universe of ARM- and fixed-rate loan types used is shown in Table 2. We include a representative mix of ARMs and 20- and 30-year fixed rate bonds with a range of coupon rates.

A proportional transaction cost of 1.5% is used; this rate is quite close to the actual refinancing costs for a DKK 1,000,000 loan. A part of this figure is actually a fixed cost (bank fees and taxes), and the effect of modeling this fixed cost is investigated in Section 5.7.

5.2 Base Case Results

The base case results, against which we subsequently perform comparisons along different dimensions, are given in Table 3. This is a fairly low-volatility scenario. We show the loans held at each node where a rebalancing occurs during the first 3 periods (years 0, 1, 2 and 5). For instance, it is seen that the initial, optimal portfolio is the 30-year 6% fixed-rate loan. In year 2, the model shifts (in the "middle-middle" branch) into the 30-year 4% etc. Note that scenario names are not directly related to the changes in interest rates (in the base case, the short and long rates tend to increase slightly in "Up"-scenarios, drop in "Middle"-scenarios,

Loan Type	Maturity	Annuity Period	Coupon
Adjustable-Rate (ARMs)	1, 2, 3, 4, 5	30	
Fixed-rate	10	10	4, 6, 8, 10, 12, 14
Fixed-rate	20	20	4, 6, 8, 10, 12, 14
Fixed-rate	30	30	4, 6, 8, 10, 12, 14

Table 2: The universe of 23 loan types used. Each loan type has a maturity date; by this date, all of the remaining, outstanding principal has to be payed back (or refinanced), and a period over which payments are annuitized, both are shown in years. Note that ARMs are annuitized over a significantly longer period than their maturities, since they are scheduled to be refinanced in a revolving fashion.

and drop (short) or increase (long) slightly in “Down”-scenarios).

In the base case, the model does not utilize the adjustable-rate loans. Fixed loans used range from 4% to 8%; both 20- and 30-year loans are used. Due to the somewhat limited universe of loan possibilities included, the model has no way to shift into ARMs after the initial decision. Also, since the objective is linear and therefore models a risk-neutral investor, the model will never hold mixed loan portfolios: informally, this is because, in an LP model, the marginal utility of a given loan is constant regardless of its level in the portfolio, hence, if a loan is present at all in an optimal portfolio, that loan alone is also optimal. Of course, if there are bounds constraints on loans proportions, mixed portfolios may still arise.

The interesting feature of these results is primarily that the model utilizes both the possibility of converting up and down, and that this behavior follows, primarily, the movements of the long rate. This clearly indicates that the possibility of converting down (absent in, for instance, the US market) has value and further studies will allow an estimation of this value.

During the first 3 periods (with 40 nodes), there were 13 rebalancing events (as seen in Table 3, and 75 during the next period (5-10 years, with 81 nodes). In other words, the probability that refinancing will become optimal from years 1 to 5 is 33% (or 6.7% per year), and from years 6 to 10 it is 93% (or 18.6% per year). These numbers are very reasonable compared to rough figures from the Danish National Bank, which indicate a refinancing rate of no less than 8% per year (or about than 2% per quarter) when interest rates are stable, more when they change (in either direction).

Note	τ	Short	Long	Loan Held
root	0	5.1	5.8	Fixed30-06
M-M	2	2.8	4.1	Fixed30-04
U-U-U	5	7.3	8.1	Fixed30-08
U-U-D	5	6.9	8.1	Fixed30-08
U-M-M	5	3.2	4.2	Fixed30-04
U-D-U	5	7.4	8.1	Fixed30-08
U-D-D	5	6.9	8.1	Fixed30-08
M-U-M	5	3.2	4.2	Fixed30-04
M-D-M	5	3.2	4.2	Fixed30-04
D-U-U	5	7.3	8.1	Fixed30-08
D-U-D	5	6.9	8.1	Fixed30-08
D-M-M	5	3.2	4.2	Fixed30-04
D-D-U	5	7.3	8.1	Fixed30-08
D-D-D	5	6.9	8.1	Fixed30-08

Table 3: Base case loans: Nodes where rebalancing occurs during the first 4 periods are shown. The model actively refinances in both directions, following movements in the long rate.

5.3 Solution Times

Solution times for the 3 - 9 stage models are shown in Table 1. These times are measured in seconds on a Windows portable PC and on a Linux machine. The models are expressed entirely using the General Algebraic Modeling System, GAMS, Brooke, Kendrick and Meer-
 aus (1992), and the LP solver used was CPLEX 7.0. Standard settings for CPLEX were used throughout. The 8-stage model contained 718,347 constraints and 1,161,239 variables, but a large part of these variables could be fixed to 0, which (by using the GAMS "hold-fixed=1" option) reduced the model size to 444,108 constraints and 465,309 variables. Of the times shown, CPLEX accounted for roughly 25%; the rest was spend by GAMS on model generation and set-up. In fact, the 9-stage model containing 29524 nodes, and about 1.3M constraints and 1.3M variables, was also solved, but was so time-consuming (about 6 hours on the PC) as to be impractical to consider further.

It is worthwhile to note that these stochastic programs of significant size were solved using standard software on off-the-shelf hardware. However, it is also clear that by the present state of hard- and software technologies, it would not be reasonable to attempt to

solve much larger instances. There are, however, special-purpose software that might allow the solution of significantly larger problems, such as IBM's OSL/SE which has also been integrated with GAMS. Another approach is to reduce the size of the stochastic program without significantly reducing its accuracy. This can be done through techniques such as sampling, bundling, or pruning. See Birge and Louveaux (1997) or Kall and Wallace (1994) for a general introduction to Stochastic Programming and its algorithms, and Censor and Zenios (1997) for a coverage of their solution through decomposition and parallel computing.

5.4 Varying the number of stages

The model was solved in versions having from 3 to 9 stages, using the base case term structure and volatility parameters. The results show a remarkable stability across the numbers of stages and hence scenarios. For any number of stages (with the exception of the 3-stage model), the optimal, initial loan is the fixed-rate, 30-year 6%. This is not unexpected: The term structure on the base date is very flat (short rate 5.14%, long rate 5.83%), which favors the fixed-rate loans. Also, in none of the models were there refinancing until stage 3, corresponding to 5 years for 3-5 stages and 2 years for more stages. Again, this is reasonable given the moderate volatilities of the base case parameters: Short rate volatility $\sigma_s = 18.5\%$, long rate volatility $\sigma_l = 14.5\%$, with correlation $\rho_{sl} = 0.96$. The 3-stage version had the ARM-5 as the optimal, initial investment; this seems to be too few stages to accurately model a 30-year period.

A clear conclusion is that in the present economic environment in Denmark, with a flat yield curve and moderate interest rate uncertainty, the optimal loan is a fixed-rate, 30-year loan, but, as we shall see later, ARMs definitely add value under more "usual" conditions.

5.5 The Risk-Averse Investor

A primary objection to the linear objective function is that it models an unlikely, risk-neutral investor. We solve in this section the 5-stage model using two objectives that address risk-aversion: An expected utility-maximization, and a worst-case minimization.

In the Expected Utility model we assume that the investor has a fixed budget, B_t , available at each point τ_t in the future, and wants to *maximize expected utility of surplus* at future times. Specifically, our objective is to:

$$\text{Maximize } z_{log} = \sum_{k=0}^K \pi^k \cdot \log(\delta \cdot (B_{\tau_k} - T^k)) \quad (8)$$

Of course, this is only one of many possible utility functions.

The results of this model clearly indicate risk-aversion. The initial portfolio is now a 10-year fixed-rate 6% loan instead of a 30-year loan, which is a relatively conservative loan type (since its duration and hence interest-rate sensitivity is low). In addition, many portfolios at later stages are diversified, which is also a clear example of “spreading your eggs”.

The non-linear programming solver Minos 5 was used for this experiment. The GAMS model was set up to solve the linear case first, which provides the non-linear solver with a feasible starting point. As a result, solution times were reasonable, and models with up to 7 stages could be solved.

As yet another example of implementing risk-averse behavior, we solved the worst-case, or min-max model which minimizes the *maximum present value* of life-time expenses across scenarios, where a *scenario* is a path of nodes from the root to a leaf node. This objective corresponds to extreme risk-aversion, and can be formulated as a linear program. See also Rustem, Becker and Marty (2000) on min-max portfolio strategies.

With this objective, the first-stage investment is now diversified into a three-way split containing two different ARMs and a fixed-rate loan! The split is 27.5% ARM-2, 28.3% ARM-5, and 44.2% Fixed30-06. At each of the 3 successor nodes (year 1), the model rebalances, each time into a new, diversified portfolio.

The results of this section show that a range of risk-aversion measures can be addressed within the framework of the model, and that the model behaves in a reasonable way in this respect. They also indicate that ARMs are of value to the highly risk-averse investor. In fact, when the ARMs are removed from the loans universe, the objective (expected present value of life-time expenditures) of the min-max model increases by 1.3% (the optimal portfolio changing to the Fixed30-06), giving an indication of the value of these instruments even in the current low-volatility and flat-yield-curve environment.

5.6 Yield Curve Steepness and Volatilities

It is expected that the steepness of the term structure, i.e., the difference in short and long rates, and the volatility of the term structure, will both have a significant impact on the optimal investment. To test this hypothesis, we perform the following experiments.

First, steepness of the term structure was varied, Table 4. We chose values for the short (1-year) and long (30-year) rates, then interpolate the intermediate values linearly. The values chosen cover the observed short and long rates during the period 1998-2001 in Figure 1, but more extreme values are also included. It is very clear that ARMs are favored whenever short rates are more than 2-3 percentage points below long rates, which is historically much more common in Denmark than the present, very flat term structure.

r_s	r_l	Root	2nd stage	Value of ARMs
0%	10%	ARM-5	U: ARM-5, M: ARM-5, D: ARM-5	8.7%
3%	6%	ARM-1	U: Fixed30-06, M: Fixed30-04, D: Fixed30-06	0.16%
3%	7%	ARM-2	U: Fixed30-06, M: ARM-2, D: Fixed30-06	0.94%
3%	10%	ARM-5	U: ARM-5, M: ARM-5, D: ARM-5	5.3%
3%	14%	ARM-5	U: ARM-5, M: ARM-5, D: ARM-5	11.3%
4%	6%	Fixed30-06	U: Fixed30-06, M: Fixed30-06, D: Fixed30-06	0%
4%	7%	ARM-1	U: Fixed30-06, M: Fixed30-06, D: Fixed30-06	0.81%
5.5%	6%	Fixed30-06	U: Fixed30-06, M: Fixed30-06, D: Fixed30-06	0%
7%	10%	ARM-2	U: ARM-2, M: ARM-2, D: ARM-2	4.1%
7%	14%	ARM-2	U: ARM-2, M: ARM-2, D: ARM-2	9.9%

Table 4: Effects of varying the yield curve steepness over the range of values observed 1998-2001 (Figure 1), and more extreme values. r_s, r_l are the short and long rates, respectively; intermediate rates were interpolated linearly. "Value of ARMs" shows the deterioration in objective value when the ARMs were excluded from the universe.

By excluding the ARMs from the loans universe we can estimate their value to the mortgagor. For short-long interest rate differentials below about 3 percentage points, their value is below 1%, but increases very rapidly as the curve steepens.

We next varied the interest model volatility parameters around their base values of $\sigma_1 = 0.0067, \sigma_2 = 0.0216, \lambda = 7.49$. These three parameters do not directly control the short and long volatilities or their correlation, but those quantities were calculated separately, and shown in Table 5 together with the optimal portfolios. It appears that ARMs are favored when interest rate volatilities are very low (not surprising), or when their correlation is low. The results are not unreasonable but also do not allow for very clear conclusions.

5.7 Fixed and Proportional Transaction Costs

Fixed costs: We have up till now used only proportional transaction costs of 1.5% on bond issues. In reality, the trading costs are composed of about 0.5% proportional costs and a fixed cost of several thousand DKK. To investigate whether modeling fixed costs for a loan of the magnitude used here (DKK 1M), we implemented a fixed-cost model using binary variables: Whenever a bond is issued, a binary variable is forced to become 1, which triggers a fixed cost per issue. When varying the fixed cost we also change the proportional cost so

Experiment	σ_s	σ_l	ρ_{sl}	Root Portfolio	Value of ARMs
Base case	18.5%	14.5%	0.96	Fixed30-06	0%
$\sigma_1 = \sigma_2 = 0$	0.0%	0.0%	1.00	ARM-1	0.18%
$\sigma_1 = 0$	5.4%	4.3%	1.00	ARM-1	0.16%
$\sigma_1 = 0.003$	9.6%	7.6%	0.84	Fixed30-06	0%
$\sigma_1 = 0.01$	26.9%	20.8%	0.98	Fixed30-06	0%
$\sigma_1 = 0.10$	218.3%	58.8%	1.00	Fixed30-04	0%
$\sigma_2 = 0.01$	18.0%	14.0%	0.99	Fixed30-06	0%
$\sigma_2 = 0.05$	21.4%	17.1%	0.83	Fixed30-06	0%
$\sigma_2 = 0.10$	29.8%	24.1%	0.62	Fixed30-06	0%
$\lambda = 0.10$	5.9%	9.5%	0.97	Fixed20-04	0%
$\lambda = 1.00$	25.4%	30.0%	0.59	ARM-5	0.36%
$\lambda = 10.0$	18.2%	14.2%	0.98	Fixed30-06	0%

Table 5: Effects of varying the interest model parameters, one at a time, from their base values of $\sigma_1 = 0.0067$, $\sigma_2 = 0.0216$, $\lambda = 7.49$. The short and long rate volatilities (σ_s, σ_l) and their correlation (ρ_{sl}) are shown, in addition to the optimal first-stage loan, and the "Value of ARMs".

that they together, for a DKK 1M loan, total DKK 15,000, or 1.5%.

The inclusion of fixed costs of DKK 10,000 per issue do not lead to a different solution to the base case. This is not too surprising, since the fixed cost at this level almost exactly replaces part of the proportional cost. If the fixed cost is increased to DKK 50,000 (which is unrealistically high but corresponds to the realistic cost for a loan of about DKK 200,000), there is no rebalancing until the very last period.

Solving the 5-stage mixed-integer model takes about 11 seconds on the Linux machine, compared to about 4 seconds for the LP case. Solving the larger mixed-integer models takes considerably longer: An attempt to solve the 6-stage model to optimality was aborted after 15 minutes, solving it to an integer gap of 0.5% took 44.2 seconds.

The min-max model of Section 5.5 was also solved using fixed transaction costs. Without fixed costs the initial portfolio was diversified into 3 instruments; with fixed costs of DKK 4,000, diversification decreased to 2 instruments, the initial loan becoming undiversified at DKK 10,000. These runs were solved to an integrality gap of 0.5%, but some still took several hours to solve. We could not solve the fixed-cost version of the utility model since it is a mixed-integer, non-linear model.

Proportional costs: Without fixed costs, the level of proportional transaction costs was increased from its base case value of 1.5%. At 10% there is no rebalancing until year 10 at stage 5 (the last stage where rebalancing can occur), where it occurs with 49% chance (40 of 81 nodes). At 20%, this probability drops to 20%. If transaction costs are removed altogether, the model rebalances vigorously: 100% at year 1, 22% at year 2, 52% at year 10, and 94% at year 20.

In conclusion, the effects of varying transaction costs are as expected, with less rebalancing the higher the transaction cost is. Including fixed costs do not change the results for the risk-neutral (undiversified) case, and at realistic levels (DKK 1,500-2,000) also do not seriously affect results for the risk-averse mortgagor, however, they do affect diversification at higher levels (or for smaller loans). Solving the mixed-integer min-max models is not feasible except for the smaller versions.

6 Conclusions

We have in this paper proposed a two-factor, no-arbitrage interest-rate model that effectively captures the (parallel or non-parallel) shifts of the term structure, and also suggested a multistage, stochastic optimization model built upon this interest-rate model. The application of the resulting, comprehensive portfolio management model to the Danish Mortgage-Backed Security market was demonstrated.

The model has certain advantages over such traditional approaches to pricing and portfolio management as Monte Carlo-simulation (where it is very hard to model American-style options features), recombining trees (which have a hard time modeling path-dependence), and Dynamic Programming (where a Markov property must be assumed, losing path-dependence, or a combinatorial explosion be suffered). Of these, our model is closest in nature to Dynamic Programming with an extended state-space, but still gives a very appealing trade-off between realism and size.

The stochastic model was implemented using standard modeling and optimization software, and was shown to be efficient in capturing the observed, real-life behavior of mortgagors. It accurately captures the effects of varying the term structure and transaction costs, and actively utilizes the embedded MBS options. It is also capable of modeling risk-aversion through the use of a concave utility function, or worst-case analysis. In addition, the model proved to be solvable on off-the-shelf hardware in a reasonable time, even for very large instances. However, combining risk-aversion and fixed transaction costs leads to models that are currently computationally intractable, but modeling fixed costs does not

seem crucial in realistic cases.

Future research currently planned based on this model includes operationalizing it for real-life portfolio management (including using the full, existing security universe), estimating the value of the various embedded options, and adapting it to be able to establish *optimal strategies* for MBS portfolio management, as opposed to step-by-step recommendations. In addition, it can be adapted to the (usually much larger) mortgage markets in other countries.

Acknowledgements. We thank two anonymous referees and several colleagues for helpful comments. We also thank ScanRate for data assistance, and GAMS Development Corporation for access to their software.

References

- [1] J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer Series in Operations Research, 1997.
- [2] T. Björk. *Arbitrage Theory in Continuous Time*. Oxford, 1998.
- [3] A. Brooke, D. Kendrick, and A. Meeraus. *GAMS: A User's Guide*. The Scientific Press, 2 edition, 1992.
- [4] Y. Censor and S.A. Zenios. *Parallel Optimization: Theory, Algorithms, and Applications*. Oxford University Press, New York, NY, 1997.
- [5] D. Heath, R. Jarrow, and A. Morton. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60(1):77–105, 1992.
- [6] P. Kall and S.W. Wallace. *Stochastic Programming*. John Wiley & Sons, New York, 1994.
- [7] P. Kang and S.A. Zenios. Complete prepayment models for mortgage-backed securities. *Management Science*, 38(11):1665 – 1685, Nov. 1992.
- [8] S.S. Nielsen and R. Poulsen. Planning your own debt. *European Financial Management*, 8(2):193–210, 2002.
- [9] S.F. Richard and R. Roll. Prepayments on fixed-rate mortgage-backed securities. *Journal of Portfolio Management*, 15(Spring):73–82, 1989.

- [10] B. Rustem, R.G. Becker, and W. Marty. Robust min-max portfolio strategies for rival forecast and risk scenarios. *Journal of Economic Dynamics and Control*, 24(11-12):1591–1621, 2000.
- [11] R. Stanton. Rational Prepayment and the Valuation of Mortgage-Backed Securities. *Review of Financial Studies*, 8(3):677–708, 1995.
- [12] C. Tanggaard. Nonparametric Smoothing of Yield Curves. *Review of Quantitative Finance and Accounting*, 9:251–269, 1997.
- [13] S.A. Zenios. Asset/liability management under uncertainty for fixed-income securities. *Annals of Operations Research*, 59:77–98, 1995. (Reprinted in *World Wide Asset and Liability Modeling*, W.T. Ziemba and J.M. Mulvey (editors), Cambridge University Press, 1998.).